

Validity conditions for stochastic chemical kinetics in diffusion-limited systems

Daniel T. Gillespie,^{1,a)} Linda R. Petzold,^{2,b)} and Effrosyni Seitaridou^{3,c)}

¹Dan T Gillespie Consulting, 30504 Cordoba Pl., Castaic, California 91384, USA

²Department of Mechanical Engineering, University of California Santa Barbara, Santa Barbara, California 93106, USA

³Department of Physics, Oxford College of Emory University, 100 Hamill Street, Oxford, Georgia 30054, USA

(Received 30 September 2013; accepted 16 January 2014; published online 6 February 2014)

The chemical master equation (CME) and the mathematically equivalent stochastic simulation algorithm (SSA) assume that the reactant molecules in a chemically reacting system are “dilute” and “well-mixed” throughout the containing volume. Here we clarify what those two conditions mean, and we show why their satisfaction is necessary in order for bimolecular reactions to physically occur in the manner assumed by the CME and the SSA. We prove that these conditions are closely connected, in that a system will stay well-mixed if and only if it is dilute. We explore the implications of these validity conditions for the reaction-diffusion (or spatially inhomogeneous) extensions of the CME and the SSA to systems whose containing volumes are not necessarily well-mixed, but can be partitioned into cubical subvolumes (voxels) that are. We show that the validity conditions, together with an additional condition that is needed to ensure the physical validity of the diffusion-induced jump probability rates of molecules between voxels, require the voxel edge length to have a strictly positive lower bound. We prove that if the voxel edge length is steadily decreased in a way that respects that lower bound, the average rate at which bimolecular reactions occur in the reaction-diffusion CME and SSA will remain constant, while the average rate of diffusive transfer reactions will increase as the inverse square of the voxel edge length. We conclude that even though the reaction-diffusion CME and SSA are inherently approximate, and cannot be made exact by shrinking the voxel size to zero, they should nevertheless be useful in many practical situations.

© 2014 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4863990>]

I. INTRODUCTION

Stochastic chemical kinetics is concerned with a system of molecules of N chemical species S_1, \dots, S_N undergoing M chemical reactions R_1, \dots, R_M inside some volume Ω at some temperature T . Its aim is to describe the behavior of $\mathbf{X}(t) \equiv (X_1(t), \dots, X_N(t))$, where $X_i(t)$ is the number of S_i molecules in Ω at time t . Its key premise is that there exists, for each reaction R_j , a *propensity function* $a_j(\mathbf{x})$ which satisfies

$$a_j(\mathbf{x}) dt \equiv \text{the probability, if } \mathbf{X}(t) = \mathbf{x}, \text{ that } R_j \text{ will fire} \\ \text{somewhere inside } \Omega \text{ in the next infinitesimal} \\ \text{time interval } [t, t + dt) \quad (j = 1, \dots, M). \quad (1)$$

It is further assumed that when an R_j reaction does occur, it changes the molecular population of species S_i by ν_{ij} , thus changing the system's state from \mathbf{x} to $\mathbf{x} + \mathbf{v}_j$ where $\mathbf{v}_j \equiv (\nu_{1j}, \dots, \nu_{Nj})$.

Under these assumptions, the laws of probability imply that the time evolution of $\mathbf{X}(t)$ can be described exactly in two equivalent ways: first via the chemical master

equation (CME),

$$\frac{\partial P(\mathbf{x}, t | \mathbf{x}_0, t_0)}{\partial t} = \sum_{j=1}^M [a_j(\mathbf{x} - \mathbf{v}_j) P(\mathbf{x} - \mathbf{v}_j, t | \mathbf{x}_0, t_0) \\ - a_j(\mathbf{x}) P(\mathbf{x}, t | \mathbf{x}_0, t_0)], \quad (2)$$

where $P(\mathbf{x}, t | \mathbf{x}_0, t_0)$ is the probability that $\mathbf{X}(t)$ will be equal to \mathbf{x} given that $\mathbf{X}(t_0) = \mathbf{x}_0$ for $t \geq t_0$; and second via the stochastic simulation algorithm (SSA):

- 1° In state \mathbf{x} at time t , generate two random numbers τ and j according to the joint probability density function $p(\tau, j | \mathbf{x}, t) = e^{-a_0(\mathbf{x})\tau} a_j(\mathbf{x})$, where $a_0(\mathbf{x}) \equiv \sum_{k=1}^M a_k(\mathbf{x})$.
- 2° Actualize the next reaction by replacing $t \leftarrow t + \tau$ and $\mathbf{x} \leftarrow \mathbf{x} + \mathbf{v}_j$.
- 3° Record the new (\mathbf{x}, t) . Return to 1°, or else end the simulation.

Since the CME and the SSA can each be derived via mathematically rigorous reasoning from the above definitions of the propensity functions a_j and the state-change vectors \mathbf{v}_j ,¹ it follows that the CME and the SSA are *logically equivalent* to each other. Anything that validates or invalidates or extends one will validate or invalidate or extend the other. In this paper, we are going to be concerned with validating and

^{a)}Electronic mail: gillespiedt@mailaps.org

^{b)}Electronic mail: petzold@cs.ucsb.edu

^{c)}Electronic mail: eseitar@emory.edu

extending the CME and the SSA, so we will often refer to them jointly as the CME/SSA.

The only problematic step in establishing the *physical* validity of the CME/SSA is proving that chemical reactions physically occur in the manner prescribed by hypothesis (1). There is a close relationship between the propensity function defined in (1) and the reaction “rate” that appears in the ordinary differential equations of traditional deterministic chemical kinetics. Early work in stochastic chemical kinetics tended to view the propensity function as a kind of *ad hoc* stochastic extension of the classical reaction rate, with the latter having the more rigorous physical justification. But in fact, the reaction rates turn out to be approximations of the propensity functions in the thermodynamic (large system) limit. That being the case, we cannot rigorously derive propensity functions from the reaction rates. Propensity functions can be reliably inferred only by looking to molecular physics to see how molecules actually behave. Doing that inevitably requires adopting a specific, usually idealized model for bimolecular chemical reactions. The model we shall adopt here is a fairly common one which assumes that: (i) the reacting molecules are hard spheres which move about as solute molecules in a sea of very many, much smaller, chemically inert solvent molecules; (ii) two solute molecules chemically react by first colliding with each other, their velocities at the instant of collision being distributed according to the Maxwell-Boltzmann distribution; and (iii) two colliding molecules will immediately undergo a specific bimolecular reaction R_j with probability q_j ($0 < q_j \leq 1$), a parameter which in principle is determined by the physics of the two molecules.

In Secs. II and III, we review the derivations of the propensity function from molecular physics, with the aim of illuminating and clarifying the conditions that must be satisfied if those derivations are to be valid. In Secs. IV–VI, we discuss the implications of these validity conditions for the reaction-diffusion extension of the CME/SSA to systems that are not completely homogeneous. In Sec. VII, we show that, provided its validity conditions are satisfied, the reaction-diffusion CME/SSA will give a plausible modeling of bimolecular chemical reactions in a solution. In Sec. VIII, we summarize our conclusions, and offer some comments on the advantages and limitations of the reaction-diffusion CME/SSA.

II. PHYSICAL JUSTIFICATION FOR THE PROPENSITY FUNCTION HYPOTHESIS

The implicit assumption in (1) that a chemical reaction is a physical event that occurs practically instantaneously dictates that every R_j in the CME/SSA will always be one of two types: either unimolecular or bimolecular. Reactions that are commonly called trimolecular or reversible practically always occur, at least in the cellular chemistry setting that we are primarily concerned with here, as a series of two or more unimolecular or bimolecular reactions. The rarity of true trimolecular reactions is a simple consequence of the rarity with which three molecules simultaneously collide with each other under well-mixed conditions; indeed, with hard-sphere molecules that virtually never happens.

Since quantum mechanics governs the way in which atoms arrange themselves into molecules, the dynamics of any unimolecular reaction $S_1 \rightarrow \dots$ is inherently stochastic. More specifically, quantum mechanics implies, at least on time scales of practical interest, that the probability that a particular S_1 molecule will undergo a unimolecular reaction in the next infinitesimal time dt will practically always be equal to some constant c_j multiplied by dt .² Summing the single-molecule reaction probability $c_j dt$ over all x_1 S_1 molecules in Ω gives, by the addition law of probability, Eq. (1) with $a_j(\mathbf{x}) = c_j x_1$.

A propensity function for the bimolecular reaction $S_1 + S_2 \rightarrow \dots$ has been shown to be justified by molecular physics in two cases: when the reactant molecules are a dilute well-mixed gas, and when they are solute molecules in a dilute, well-mixed solution. Although the latter case is the only one of relevance for cellular chemistry, the simpler derivation in the former case illustrates more clearly the necessity of the dilute and well-mixed requirements. The propensity function for a dilute well-mixed gas is³

$$a_j(x_1, x_2) = (\pi \sigma_{12}^2 \bar{v}_{12} q_j |\Omega|^{-1}) \cdot x_1 x_2 \quad (\text{dilute gas}). \quad (3)$$

Here, σ_{12} is the average distance between the centers of a pair of reactant molecules at collision (the sum of their radii for hard sphere molecules); $\bar{v}_{12} = \sqrt{(8k_B T)/(\pi m_{12})}$ is their average relative speed, with m_{12} their reduced mass, k_B Boltzmann’s constant, and T the absolute temperature of the system; and q_j is the probability that an S_1 – S_2 collision will result in an R_j reaction. The derivation of Eq. (3) goes as follows: $\pi \sigma_{12}^2 \cdot \bar{v}_{12} dt$ is the average “collision volume” that a randomly chosen S_2 molecule sweeps out relative to the center of a randomly chosen S_1 molecule in time dt . Dividing that collision volume by the system volume $|\Omega|$ gives, for reasons that will be elaborated below, the probability that the center of the S_1 molecule lies inside the collision volume, and hence the probability that the two molecules will collide in the next dt . That collision probability multiplied by the conditional probability q_j of a reaction given a collision yields, by the multiplication law of probability, the probability that the two molecules will react in the next dt . And finally, that single-pair reaction probability summed over all $x_1 x_2$ distinct reactant pairs gives, by the addition law of probability, the probability defined in Eq. (1), thus establishing Eq. (3). For the bimolecular reaction $2S_1 \rightarrow \dots$, the sum over all distinct reactant pairs would give a factor $\frac{1}{2} x_1(x_1 - 1)$ instead of $x_1 x_2$. A “collision energy threshold” reaction model yields for q_j the well-known Arrhenius factor; however, we need not assume anything specific here about q_j other than it will have some value between 0 and 1.

The critical step in the foregoing derivation is the assertion that the ratio of {the average collision volume $\pi \sigma_{12}^2 \cdot \bar{v}_{12} dt$ } to {the total system volume $|\Omega|$ } provides a valid estimate of the probability that the center of the S_1 molecule will lie inside the collision volume. This key assertion entails two assumptions: First, *the probability that a reactant molecule will be found inside any small subvolume of Ω is independent of where inside Ω that subvolume is*. That will be our definition of a “well-mixed” system. Note that this definition does *not* require that there be a perfectly regular placement

of the reactant molecules inside Ω , nor that there be a large number of those molecules.

The second assumption implicit in the aforementioned assertion is that all but a negligibly small fraction of the containing volume Ω is physically accessible to the center of the randomly chosen S_1 molecule. That of course assumes Ω to be sufficiently convex that this would be true in the *absence* of any other molecules. But the presence of other reactant molecules will inevitably occlude some of Ω . Just how that occluded volume should be quantitatively taken into account is far from clear. But it is clear that the argument leading to Eq. (3) will be valid only if *the volume occluded by the reactant molecules inside Ω is very small compared to $|\Omega|$* . That will be our definition of a “dilute” system. If we make the simplifying approximation that each S_1 molecule and S_2 molecule has a diameter that is the average of the two, namely, σ_{12} , then a rough estimate of the total volume that all those molecules occlude is

$$(x_1 + x_2) \cdot \frac{4}{3}\pi \left(\frac{1}{2}\sigma_{12}\right)^3 = \frac{\pi}{6}(x_1 + x_2)\sigma_{12}^3.$$

Therefore, the diluteness assumption is essentially the order-of-magnitude requirement

$$(x_1 + x_2)\sigma_{12}^3 \ll |\Omega|. \quad (4a)$$

Since $|\Omega|/(x_1 + x_2)$ is the average volume allotted to each reactant molecule inside Ω , then the cube root of that quantity estimates the average distance between the reactant molecules. Condition (4a) evidently requires that average distance to satisfy

$$\left(\frac{|\Omega|}{x_1 + x_2}\right)^{1/3} \gg \sigma_{12}. \quad (4b)$$

So an *equivalent* way of stating the diluteness requirement is to say that *the average distance between a pair of reactant molecules must be very large compared to their average diameter*.

Since the minimum value of the factor $(x_1 + x_2)$ in Eq. (4a) for a reaction to occur is 2 (when $x_1 = x_2 = 1$), then a *minimal practical requirement for diluteness* would appear to be that *the diameter $|\Omega|^{1/3}$ of the system must be at least an order of magnitude larger than the average diameter σ_{12} of a reactant molecule*. It might be thought that this diluteness requirement is simply the requirement that the reactant molecules be “points.” But that would not be a correct assessment. As can be seen by putting $\sigma_{12} = 0$ in the propensity function (3), two point molecules have zero probability of reacting with each other; therefore, satisfying the dilute conditions (4) by simply assuming the reactant molecules are points and setting $\sigma_{12} = 0$ is not viable option.

A convincing physics derivation of a bimolecular propensity function for the situation in which the reactant molecules are *solute molecules in a solution* eluded researchers for a long time. The fact that the standard diffusion equation implies that the average displacement of a solute molecule in time dt is proportional to \sqrt{dt} seemed to suggest, at least on the basis of the foregoing derivation of the dilute gas result, that the probability of a reaction between two solute molecules in the next dt might be proportional to \sqrt{dt} instead of dt . That would be totally inconsistent with hypothesis (1).

But in what might be described as a refined, corrected, and stochastically extended version of the analysis of Collins and Kimball,⁴ it was recently shown that if the reactant molecules are dilute and well-mixed in the senses defined above, then a propensity function for $S_1 + S_2 \rightarrow \dots$ as defined in (2) does exist and is given by²

$$a_j(x_1, x_2) = \left(\frac{4\pi\sigma_{12}^2 D_{12} \bar{v}_{12} q_j |\Omega|^{-1}}{4D_{12} + \sigma_{12} \bar{v}_{12} q_j}\right) \cdot x_1 x_2 \quad (\text{dilute solution}). \quad (5)$$

Here, D_{12} is the sum of the diffusion coefficients of the S_1 and S_2 molecules, and the other quantities are as previously defined. Note that the requirement for diluteness in a solution applies only to the reactant (solute) molecules, and not to the chemically inert solvent molecules. In the “fast-diffusion” regime $4D_{12} \gg \sigma_{12} \bar{v}_{12} q_j$, Eq. (5) reduces to the dilute gas result (3). At the opposite extreme $4D_{12} \ll \sigma_{12} \bar{v}_{12} q_j$, which is the “diffusion-limited” or “Smoluchowski” regime which typifies cellular systems, the factor in parentheses in Eq. (5) reduces to $4\pi\sigma_{12} D_{12} |\Omega|^{-1}$. This result corresponds to a well known deterministic rate result that can be obtained by adapting Smoluchowski’s famous analysis of colloidal coagulation.⁶ The more rigorous derivation of Eq. (5) actually makes use of Smoluchowski’s reasoning, but does so in a way that takes account of the often overlooked fact² that the standard diffusion equation on which the Smoluchowski analysis is based is physically incorrect on small length scales (we will return to this point in Sec. V). As in the case of the derivation of the dilute gas result (3), the derivation of Eq. (5) fails if the system is not dilute and well-mixed, for reasons that are basically the same as in the dilute gas case.

To summarize: Molecular physics provides the following justification for the CME/SSA when the reactant molecules are solute molecules in a solution with very many, much smaller, chemically inert solvent molecules:

- (i) A propensity function for the *unimolecular* reaction $S_1 \rightarrow \dots$ normally exists in the form $a_j(\mathbf{x}) = c_j x_1$, where c_j is independent of both \mathbf{x} and $|\Omega|$.
- (ii) A propensity function for the *bimolecular* reaction $S_1 + S_2 \rightarrow \dots$ normally exists *provided* the reactant (solute) molecules are dilute and well-mixed inside Ω . The propensity function then has the form $a_j(\mathbf{x}) = c_j x_1 x_2$, where c_j is independent of \mathbf{x} and inversely proportional to $|\Omega|$.

An examination of the derivation of Eq. (5)² reveals that the bimolecular reaction results in (ii) will *not* be true if the reactant molecules are not dilute and well-mixed. In those situations, there is at present no valid physics-based derivation of a propensity function, and hence no valid physics-based CME and SSA.

III. BEING DILUTE IS A NECESSARY AND SUFFICIENT CONDITION FOR STAYING WELL-MIXED

Diffusion not only serves to bring solute molecules together so that they can chemically react, it is also the main mechanism by which the system stirs itself. In order for the

bimolecular reaction $S_1 + S_2 \rightarrow \dots$ to occur in the well-mixed setting assumed by the propensity function hypothesis (1), the S_1 and S_2 molecules need to diffuse around for a while before they react with each other. We will now show that will happen if and only if the system is dilute.⁷

Consider a randomly chosen S_1 molecule. According to the standard theory of diffusion, this molecule will move an average net distance $\sqrt{2D_1\tau}$ in a time τ . So since the diameter of Ω is on the order of $|\Omega|^{1/3}$, the time needed for the S_1 molecule to randomly reposition itself inside Ω will be roughly the time τ_d defined by $\sqrt{2D_1\tau_d} = |\Omega|^{1/3}$:

$$\tau_d = \frac{|\Omega|^{2/3}}{2D_1}. \quad (6)$$

According to Eqs. (1) and (5), the probability that this S_1 molecule will react with one of the x_2 S_2 molecules inside Ω in the next infinitesimal dt is

$$\left(\frac{4\pi\sigma_{12}^2 D_{12} \bar{v}_{12} q_j |\Omega|^{-1} x_2}{4D_{12} + \sigma_{12} \bar{v}_{12} q_j} \right) \cdot dt.$$

But in the “diffusion-limited” regime assumed by Eq. (6), which as noted in connection with Eq. (5) is defined by $4D_{12} \ll \sigma_{12} \bar{v}_{12} q_j$, this probability reduces to

$$(4\pi\sigma_{12} D_{12} |\Omega|^{-1} x_2) \cdot dt.$$

From this it is easy to show that the average time τ_r before the S_1 molecule will react with one of the x_2 S_2 molecules inside Ω in the next infinitesimal dt is

$$\tau_r = (4\pi\sigma_{12} D_{12} |\Omega|^{-1} x_2)^{-1}. \quad (7)$$

What is needed to secure the well-mixed condition for this reaction is that {the average time required for the S_1 molecule to become randomly repositioned inside Ω } be much less than {the average time before the S_1 molecule reacts with one of the x_2 S_2 molecules inside Ω }, i.e., $\tau_d \ll \tau_r$. Substituting into this condition from Eqs. (6) and (7) yields the requirement

$$2\pi\sigma_{12}(D_{12}/D_1)x_2 \ll |\Omega|^{1/3}. \quad (8)$$

Since D_1 and D_2 are typically the same order of magnitude, we can further approximate $D_{12}/D_1 \approx 2$, and so conclude that the well-mixed condition requires

$$4\pi x_2 \sigma_{12} \ll |\Omega|^{1/3}. \quad (9)$$

In words, the sum of the diameters of x_2 (and also by symmetry x_1) “average” reactant molecules must be much less than the diameter of the system. From a practical point of view, this is basically the diluteness condition (4b). Thus we have shown that the reactant molecules in a diffusion-limited system will remain well-stirred if and only if they are dilute. We note in passing that essentially this same result can also be established in the ideal gas regime, $4D_{12} \gg \sigma_{12} \bar{v}_{12} q_j$.⁸

This is an intuitively plausible result that seems not to be widely appreciated: If the system is “dilute,” so that the average distance between reactant molecules is large compared to their average diameter, then a reactant molecule will usually have to wander around for a relatively long time before it chances to collide with another reactant molecule. That wandering around makes the system well-mixed. But if the reactant molecules crowd each other, then they will likely find

reacting partners before they have wandered very far, and consequently they will react in a system that is not well-mixed.

IV. THE REACTION-DIFFUSION CME/SSA

If the system is not well-mixed in Ω , but can be considered approximately well-mixed inside each subvolume or “voxel” Ω_k of some partitioning $\{\Omega_1, \dots, \Omega_K\}$ of Ω , then one can try the following: Model the chemical reactions as events occurring wholly inside single voxels as is done in the standard CME/SSA, and model the diffusional movement of reactant molecules between adjacent voxels as concurrent “diffusive transfer reactions” in a way that is consistent with the standard diffusion equation. This is the strategy of what is commonly called the “reaction-diffusion master equation (RDME)” and the “spatial SSA” (or rather less commonly, the “spatial CME” and the “reaction-diffusion SSA”). We shall refer to these jointly as the “reaction-diffusion CME/SSA.” The original paper on this strategy was the 1976 paper of Gardiner *et al.*⁹ As was done in that paper, we will take the voxels to be *cubes* of edge length h , since that greatly simplifies the mathematics of the diffusive transfer reactions.

Parceling the chemical reactions out to the K voxels has the effect of replacing the M nominal reactions $\{R_j\}$ by KM reactions $\{R_{jk}\}$, where R_{jk} is reaction R_j inside voxel Ω_k . The propensity function a_{jk} for R_{jk} is the propensity function a_j for R_j , but now referred to the voxel Ω_k and regarded as a function of $\mathbf{x}_k = (x_{1k}, \dots, x_{Nk})$, where x_{ik} is the current number of S_i molecules in Ω_k . If R_j is a unimolecular reaction, the coefficient c_{jk} in the R_{jk} propensity function will be identical to the coefficient c_j for R_j . If R_j is a bimolecular reaction, the coefficient c_{jk} will be the factor in parentheses in Eq. (5), except the factor $|\Omega|^{-1}$ there must be replaced by h^{-3} , the reciprocal of the voxel volume. In all cases, the state-change vector ν_{jk} for R_{jk} is the ν_j for R_j , but confined to the space of \mathbf{x}_k .

The reaction-diffusion CME/SSA aims to model the movement of an S_i molecule inside Ω in accordance with the standard Einstein diffusion equation

$$\frac{\partial p(\mathbf{r}, t)}{\partial t} = D_i \nabla_{\mathbf{r}}^2 p(\mathbf{r}, t), \quad (10)$$

where p is the position PDF of a single S_i molecule, and D_i is the molecule’s diffusion coefficient. But since in the reaction-diffusion CME/SSA we do not know the position \mathbf{r} of any reactant molecule—we know only that the molecule is equally likely to be anywhere inside a particular voxel—we must be content to model only the transfer of some S_i molecule from its current voxel Ω_k to an adjacent voxel Ω_l . The reaction-diffusion CME/SSA accomplishes this by positing a “diffusive transfer reaction” $R_{ikl}^{(d)}$, which has propensity function

$$a_{ikl}^{(d)}(\mathbf{x}_k) = \frac{D_i}{h^2} x_{ik}. \quad (11)$$

Thus, the reaction-diffusion CME/SSA assumes that $a_{ikl}^{(d)}(\mathbf{x}_k) dt$ is the probability that an S_i molecule in Ω_k will move to Ω_l in the next dt . We will examine in detail the justification for this assumption in Sec. VI. The state-change vector $\nu_{ikl}^{(d)}$ for reaction $R_{ikl}^{(d)}$ simply decreases x_{ik} by 1 and increases x_{il} by 1.

Since the state jumps induced by diffusional transfers of reactant molecules between voxels have the same mathematical character as the state jumps induced by chemical reactions (both are described by propensity functions and state-change vectors), then the reaction-diffusion CME/SSA is just the standard CME/SSA described in Sec. I, but with the following re-interpretation of its symbols: the N -dimensional state vector $\mathbf{x} = \{x_i\}$ is now regarded as the KN -dimensional state vector $\{x_{ik}\}$; and the M propensity functions $\{a_j\}$ and their associated state-change vectors $\{\mathbf{v}_j\}$ are now regarded as those for the KM chemical reactions $\{R_{jk}\}$ and the $2NB$ diffusive transfer reactions $\{R_{ikl}^d\}$, where B is the total number of boundaries between adjacent voxels. The factor $2N$ here comes from the fact that a molecule of any of the N species can cross the boundary in either direction.

It is important to understand that the voxel strategy of the reaction-diffusion CME/SSA is *inherently approximate*. This can be seen in two ways. First is the artificiality of its assumption that the position PDFs of the reactant molecules are uniform inside each voxel and change discontinuously at the boundaries between voxels. Since the boundaries between voxels are imaginary constructs, no real chemical system will ever present itself in this way. Second is the artificiality of its assumption that each reactant molecule is always wholly inside a single voxel. Even if we were to stipulate (as we usually do) that a reactant molecule is “in” the voxel that contains its center of mass, the molecule’s non-zero size can extend its physical presence to neighboring voxels. Since these artificialities are most severe near the boundaries of the voxels, and least severe near the centers of the voxels, it follows that the reaction-diffusion CME/SSA describes the system *only to resolution* h . Increased accuracy can thus be obtained only by decreasing h .

But there is a price to be paid for making h smaller. For example, halving h increases the total number of voxels by a factor of $2^3 = 8$. That causes factor-of-8 increases in the number of state variables $\{x_{ik}\}$, and in the number of chemical reaction channels $\{R_{jk}\}$, and (at least approximately) in the number of diffusive transfer reaction channels $\{R_{ikl}^d\}$. Furthermore, as we will see later, halving h will also increase the average number of diffusive transfer events that will occur between successive chemical reaction events by a factor of $(1/2)^{-2} = 4$. In view of these substantial computational penalties for adopting a finer spatial resolution, it would appear that in making practical computations we should always choose h as large as possible. We will have more to say about choosing h later.

An often voiced concern with the reaction-diffusion CME/SSA arises in connection with its rule that bimolecular chemical reactions can occur only between molecules that are in the same voxel. That would seem to allow a reaction between two molecules that are relatively far apart in the same voxel, yet not allow reactions between two molecules that are very close together but separated by a voxel boundary. However, there is an answer to this concern: the hypothesis that reactant molecules are “well-mixed” inside any voxel *precludes* us from making positional distinctions between molecules in the same voxel. In particular, a reactant molecule that is “near a boundary of its voxel” is *not* a reactant molecule that

is “uniformly distributed inside its voxel”, and the latter is the *only* kind of reactant molecule that the reaction-diffusion CME/SSA is able to say anything about.

V. DIFFUSIVE TRANSFER REACTIONS

The reaction-diffusion CME/SSA transfers an S_i molecule from its present voxel Ω_k to an adjacent voxel Ω_l in accordance with the propensity function (11). In the Appendix, we give a concise proof of the fact that *in the limit* $h \rightarrow 0$ this molecule transfer strategy *exactly* replicates the dynamics prescribed by the standard diffusion equation (10). Given that result, one might be tempted to conclude that, at least in the absence of chemical reactions, the reaction-diffusion CME/RRA becomes exact in the limit $h \rightarrow 0$. But that is not true.

To see why, consider a single S_i molecule in some interior voxel. Equation (11) stipulates that the probability that this molecule will jump to a particular one of its adjacent voxels in the next infinitesimal time dt is $(D_i/h^2)dt$. Then by the addition law of probability, the probability that the molecule will, in the next dt , jump to either of the two adjacent voxels along a particular Cartesian axis, say the x -axis, is $2(D_i/h^2)dt$. That implies that the time it takes the molecule to leave its present voxel for either of those adjacent voxels is an exponential random variable with mean $h^2/(2D_i)$. Thus, the *average time* it takes the molecule to move to a distance h along the x -axis is $h^2/(2D_i)$. Therefore, the *average speed* of the molecule along the x -axis while making that voxel-to-voxel transition is

$$\bar{s}_{\text{vv}} = \frac{\text{avg dist moved}}{\text{avg time}} = \frac{h}{h^2/(2D_i)} = \frac{2D_i}{h}. \quad (12)$$

Note that this voxel transition speed will be *less* than the molecule’s *average instantaneous speed* \bar{s} along the x -axis,

$$\bar{s}_{\text{vv}} < \bar{s}, \quad (13)$$

since, owing to the irregular back and forth motion of the diffusing molecule, the total distance it travels along the x -axis during that voxel transition will be larger than h .

Equations (12) and (13) imply the disturbing result that the molecule’s average instantaneous speed \bar{s} along any Cartesian axis, as well as the average speed \bar{s}_{vv} with which it moves from voxel to voxel in a given Cartesian direction, both go to *infinity* as $h \rightarrow 0$. It is important to understand that this result is *not* a harmless reflection of the fact that the Maxwell-Boltzmann velocity distribution allows molecular speeds to be unbounded. The Maxwell-Boltzmann distribution says that the velocity along any Cartesian axis of a particle of mass m at absolute temperature T will be a normal random variable with mean zero and variance $k_B T/m$. That implies that the root-mean-square velocity of our S_i molecule along any Cartesian axis is $\sqrt{k_B T/m_i}$, and to a factor of order unity that is the “average speed” of the S_i molecule along the x -axis. So even though the Maxwell-Boltzmann distribution does allow molecular speeds that are arbitrarily large, its *average speed* is *finite*. In contrast, the *average speed* \bar{s} predicted by the diffusional jumping rule (11) *grows without bound* as

$h \rightarrow 0$. There are at least two other ways of demonstrating this unphysical prediction of the standard diffusion equation.¹⁰

Since, as we have just seen, the average instantaneous x -axis speed of an S_i molecule moving according to the voxel-hopping rule of the reaction-diffusion CME/SSA is always *greater* than $2D_i/h$, and since statistical thermodynamics stipulates that the average instantaneous x -axis speed of a molecule of mass m_i at temperature T is approximately $\sqrt{k_B T/m_i}$, then a *minimal* condition for the reaction-diffusion CME/SSA to conform to the requirements of statistical thermodynamics is $\sqrt{k_B T/m_i} > 2D_i/h$. Solving this for h , we conclude that a *minimal* condition for the reaction-diffusion CME/SSA to be *physically acceptable* is

$$h > 2D_i \sqrt{\frac{m_i}{k_B T}}.$$

In fact, an analysis based on Langevin's more accurate theory of Brownian molecular diffusion reveals that the voxel-hopping rule (11) will be accurate *only* if this inequality is *strongly* satisfied, i.e., only if¹¹

$$h \gg D_i \sqrt{\frac{m_i}{k_B T}}. \quad (14)$$

To summarize: In the limit $h \rightarrow 0$, the voxel-hopping rule (11) converges exactly to the behavior predicted by the diffusion equation. But the diffusion equation, despite its name, does *not* accurately describe the physical motion of a molecule undergoing simple Brownian diffusion *on all length scales*; more specifically, it fails on length scales smaller than that prescribed by condition (14), where the molecule actually moves ballistically. If the voxel size h is steadily decreased below the value in (14), the voxel-hopping rule's modeling of a diffusing molecule's *physical* motion becomes more and more inaccurate, and eventually catastrophically inaccurate. For example, if h is reduced below $2D_i/c$ where c is the speed of light, then Eq. (12) implies that S_i molecules will be moving from voxel to voxel at an average speed greater than c . Even aside from the strictures of special relativity, it should be obvious that there is no physically plausible mechanism by which the velocity of a solute molecule could be continually and rapidly switched between "very fast in one direction" to "very fast in the opposite direction" by collisions with the surrounding much smaller and much less massive solvent molecules. We conclude that if the diffusional motion of molecules in the reaction-diffusion CME/SSA is to remain correct from the perspective of *physics*, then h must obey the lower bound condition (14).

VI. LOWER BOUNDS ON h

Since smaller values of h (i) improve the accuracy of the key assumption that the PDFs of the reactant molecules do not vary appreciably over any one voxel and (ii) cause the diffusive transfer propensity function (11) to more accurately model the Einstein diffusion equation, it is tempting to conclude that the accuracy of the reaction-diffusion CME/SSA can be made as great as desired simply by taking h sufficiently small. However, in Sec. V we saw that there is a caveat to (ii), in that if h is smaller than allowed by condition (14) then the

molecular motion prescribed by the diffusive transfer propensity function (11) will be physically incorrect. This means that we are not free to take h arbitrarily small, and we certainly cannot take the limit $h \rightarrow 0$.

There is, furthermore, *another* lower bound on h which in practice is usually more restrictive than (14). It arises from the presumption that the CME/SSA correctly describes bimolecular reactions occurring *inside individual voxels*. In order for that to be true, the reactant molecules for all bimolecular reactions must be *dilute and well-mixed inside each voxel*. As shown in Secs. II and III for the bimolecular reaction $S_1 + S_2 \rightarrow \dots$ inside a volume Ω , being dilute and staying well-mixed requires satisfaction of conditions (4) and (9), conditions which in all practical circumstances amount to the requirement that $|\Omega|^{1/3}$ be very much larger than σ_{12} . Thus, the reaction-diffusion CME/SSA requires that *the voxel edge length must be much larger than σ_{12}* :

$$h \gg \sigma_{12}. \quad (15)$$

Calculations using parameter values roughly typical of protein molecules in water at room temperature suggest that the h -bound in (15) will be several thousand times larger than the h -bound in (14). In any case, both conditions (14) and (15) must be satisfied if the reaction-diffusion CME/SSA is to be *physically* valid.

VII. BEHAVIOR OF THE AVERAGE REACTION RATES FOR SMALL VOXELS

It has been argued elsewhere¹² that, in the limit $h \rightarrow 0$, bimolecular reactions stop occurring in the reaction-diffusion CME/SSA. In view of the positive lower bounds on h in Eqs. (14) and (15), it might be objected that the limit $h \rightarrow 0$ is not allowed. But let us suppose that h is decreased in a *restricted* way that does *not* violate conditions (14) and (15). In that case it is fair to ask, does the average rate at which bimolecular reactions occur inside Ω decrease along with h ? In this section, we will show that it does not.

For this demonstration, we will consider a system that is well-mixed throughout its entire volume Ω . We begin by examining how chemical and diffusive transfer reactions *in a single voxel* behave as h is made smaller and smaller, but not smaller than allowed by conditions (14) and (15). Since the decrease in h is not accompanied by any change in either Ω or the total number x_i of S_i molecules inside Ω , the average density of the S_i molecules in Ω will remain constant at the value $x_i/|\Omega|$. The average number of S_i molecules inside any one voxel will therefore be $(x_i/|\Omega|) \cdot h^3 \equiv \bar{x}_i$, and thus will decrease with h like h^3 . Notice that if h is taken small enough, we could have $\bar{x}_i \ll 1$; however, satisfaction of the general diluteness condition ensures that that will not be due to the voxel being too small to accommodate more than one reactant molecule, but simply to there being many more voxels than reactant molecules in Ω .

The average propensity function for the *unimolecular chemical* reaction $S_1 \rightarrow \dots$ inside a single voxel has the form $c\bar{x}_1$, where c is independent of h . Since $\bar{x}_1 \propto h^3$, the average propensity function for a unimolecular chemical reaction in a single voxel will therefore decrease with h like h^3 .

The average propensity function for the *bimolecular chemical* reaction $S_1 + S_2 \rightarrow \dots$ inside a single voxel has the form $c\bar{x}_1\bar{x}_2$ where $c \propto h^{-3}$. Here we have invoked the proviso that h be large enough to satisfy condition (15)—see point (ii) at the end of Sec. II. Since $\bar{x}_i \propto h^3$, the average propensity function for a bimolecular chemical reaction in a single voxel will therefore decrease with h like $h^{-3}h^3h^3 = h^3$. Note that this is the same small- h behavior as for a unimolecular chemical reaction.

By Eq. (11), the average propensity function for the *diffusive transfer* reaction in which an S_i molecule leaves a given voxel for a particular adjacent voxel is $(D_i/h^2)\bar{x}_i$, where D_i is independent of h . Since $\bar{x}_i \propto h^3$, the average propensity function for a diffusive transfer of an S_i molecule from a particular voxel to a particular adjacent voxel will therefore decrease with h like $h^{-2}h^3 = h$.

To summarize, the average propensity function for any *chemical* reaction inside a particular voxel will decrease with h like h^3 , while the average propensity function for any *diffusive transfer* reaction from a particular voxel to a particular adjacent voxel will decrease with h like h . However, our goal is to assess the behavior of these reactions over the entire volume Ω . To do that, we will use the fact that propensity functions are additive; because, the probability for an R_1 reaction or an R_2 reaction to occur in the next infinitesimal time dt is, by the addition law of probability,¹³

$$a_1dt + a_2dt = (a_1 + a_2)dt.$$

As a consequence of propensities being additive, the average propensity function for a *chemical* reaction *anywhere* inside Ω will be the sum of all the single voxel chemical reaction propensities. We can estimate that sum as the product of {the average single-voxel chemical reaction propensity}, which we have just seen is $\propto h^3$, times $\{M \times$ the total number of voxels}. Since the total number of voxels is $|\Omega|/h^3$, we conclude that the h -dependence of the average propensity function for a chemical reaction anywhere inside Ω will be $\propto h^3 \times h^{-3} = h^0$, i.e., the average propensity function for a *chemical* reaction anywhere inside Ω will be *independent* of h .

What about diffusive transfer reactions? The average propensity function for those will be the product of {the average single-molecule diffusive transfer propensity}, which we found above is $\propto h$, times $\{2 \times N \times$ the total number of interfaces between adjacent voxels}; the factor $2 \times N$ here accounts for the fact that at each such interface, a molecule of any of the N species can cross in either direction. An *interior* voxel will have 6 such interface boundaries, each of which is shared with one other voxel. Since the ratio of the number of voxels on the boundary of Ω to the number of interior voxels will approach zero with h , the total number of interfaces between adjacent voxels will be approximately $6 \times \frac{1}{2} \times$ the total number of voxels $|\Omega|/h^3$. The h -dependence of the total number of interfaces between adjacent voxels is therefore approximately $\propto h^{-3}$. Thus, the h -dependence of the average propensity function for a *diffusive transfer reaction* anywhere inside Ω will be approximately $\propto h \times h^{-3} = h^{-2}$.

Since the reciprocal of the propensity function of a reaction gives the average time between those reactions, we have

thus proved the following: If h is decreased in a way that respects the lower bounds (14) and (15), then the average time between *chemical* reactions (unimolecular or bimolecular) inside Ω does not change, while the average time between *diffusive transfer* reactions decreases approximately like h^2 .

VIII. SUMMARY AND CONCLUSIONS

The physical validity of the CME/SSA model of chemical kinetics hinges solely on there being a sound basis in molecular physics for the propensity function hypothesis (1). We began by showing that if the reactant molecules are hard-sphere solute molecules in solution with very many much smaller chemically inert solvent molecules, then the propensity function hypothesis can be justified from molecular physics only if the reactant molecules of all bimolecular reactions are *dilute* and *well-mixed* inside the containing volume Ω . Here, “dilute” means that the total volume occluded by the reactant molecules (but not the solvent molecules) is only a very small fraction of $|\Omega|$, or equivalently that the average distance between two reactant molecules is very large compared to their diameters. “Well-mixed” means that the probability of finding the center of a reactant molecule inside any small subvolume of Ω is independent of the location of that subvolume. We presented an argument showing that ordinary diffusion will suffice to maintain the well-mixed condition if and only if the diluteness condition is satisfied, i.e., the system will stay well-mixed if and only if it is dilute.

We next examined the reaction-diffusion extension of the CME/SSA. It assumes that even if the reactant molecules are not well-mixed inside Ω , we can partition Ω into a set of cubic voxels $\Omega_1, \dots, \Omega_K$, each of edge length h , such that the reactant molecules are approximately well-mixed inside each voxel. Chemical reactions are then viewed as occurring wholly inside single voxels in accordance with the usual propensity functions of the CME/SSA, while the diffusion of reactant molecules between adjacent voxels is modeled as diffusive transfer reactions with propensity functions of the form (11). The obvious physical artificiality of the two assumptions (i) that the distributions of the reactant molecules are perfectly uniform inside each voxel and change discontinuously at voxel boundaries, and (ii) that each reactant molecule always lies wholly inside a single voxel, makes it clear that this is an inherently approximate, coarse grained description, its spatial resolution being the voxel edge length h .

Although a smaller value of h will yield a more finely resolved description, we showed that *physics* considerations impose two lower bounds on h which prevent us from taking h arbitrarily small. First is the lower bound stipulated by condition (14). It ensures that the propensity function (11) for diffusive transfer reactions is physically accurate; taking h smaller than allowed by condition (14) will result in movement of the reactant molecules between voxels that is unphysical, despite being consistent with the predictions of the standard diffusion equation. Second is the lower bound on h stipulated by condition (15). It ensures that bimolecular reactions occurring inside voxels can be described using propensity functions of the form (5); taking h smaller than allowed by condition (15) will make it impossible for two reactant molecules to be “dilute”

inside a single voxel, as is required to derive the bimolecular propensity function in a voxel. The lower bound (15), namely, $h \gg \sigma_{12}$, will usually be the controlling one, since it will typically be several orders of magnitude larger than the lower bound (14).

We finally showed that if a chemical system is well-mixed inside its full volume Ω and we decrease h in a way that respects the lower bounds (14) and (15), then contrary to what might be inferred from a recent analysis,¹² the average rate at which both unimolecular and bimolecular chemical reactions occur inside Ω does not change. That, of course, is exactly what we should expect: the rate at which chemical reactions occur inside Ω should not be affected by how finely we subdivide Ω into imaginary voxels. We also showed that the average rate at which *diffusive transfer* reactions occur inside Ω increases with decreasing h approximately like h^{-2} . So, if we imagine a time-line on which all reactions are recorded by placing a dot at the instant they occur, then as h gets smaller, the “chemical reaction dots” will on average stay as they are, but the “diffusive transfer dots” will become more numerous, filling in the spaces between the chemical reaction dots.

At least three earlier works^{14–16} have suggested modifications to the rate constant in the bimolecular propensity function (5) in an attempt to allow the reaction-diffusion CME/SSA to be used with very small values of h . These three analyses differ in the specific reactions they consider, the basic assumptions they make, the inference logic they use, and the specific results they obtain; thus it is very difficult to compare them with each other or with our work here. Fange *et al.*¹⁴ appear to push the lower bound on h all the way to zero. That would contradict the lower bound (14), which we claim arises whenever the diffusional hopping rule (11) is used. But the results of both Erban and Chapman¹⁵ and Hellander *et al.*¹⁶ do imply positive lower bounds on h . Hellander *et al.*¹⁶ have suggested that those two lower bounds more or less agree, and they concluded that no h -dependent modification of the bimolecular rate constant in Eq. (5) will allow the reaction-diffusion CME/SSA to be correct unless $h > \pi\sigma_{12}$. If that is true, then it would appear that such modifications cannot yield much in the way additional latitude for h beyond our restriction of $h \gg \sigma_{12}$ on the conventional reaction-diffusion CME/SSA. We have made no attempt in our paper to find a generalization of the bimolecular propensity function (5) that would make it correct when the volume occluded by all the reactant molecules in a single voxel is not negligibly small compared to the voxel volume $|\Omega| = h^3$. But since that occluded volume will depend on the numbers x_{1k}, \dots, x_{Nk} of reactant molecules that are currently in voxel k , then it would seem that any physics-based correction to the propensity function (5) will inevitably also change its dependence on the voxel population variables x_{1k}, \dots, x_{Nk} to something considerably more complicated than the simple factor $x_{1k}x_{2k}$.

In the quest for a lower bound on h , we should bear in mind that a *practical* lower bound might exist that is *larger* than the lower bound $h \gg \sigma_{12}$, owing to the fact that the computational complexity of the reaction-diffusion CME/SSA increases so rapidly with decreasing h . For example, merely halving h increases by a factor of 8 both the number of state variables and the number of reaction channels, and also in-

creases by a factor of 4 the average number of diffusive transfer events that occur between successive chemical reaction events. So even if decreasing h is desirable, it might not be feasible. From a practical standpoint, the goal in using the reaction-diffusion CME/SSA should be to use a value for h that is just small enough to capture the spatial non-uniformities in the system. Anything smaller will only make the calculation more difficult.

Our conclusion that the reaction-diffusion CME/SSA is inherently approximate, and cannot be made arbitrarily accurate by taking h arbitrarily small, immediately prompts the following question: Are there finer-scale models that can be used in those regions of space where one requires greater resolution than can be provided by the reaction-diffusion CME/SSA? There is of course molecular dynamics, which meticulously tracks the movement of not only all the reactant (solute) molecules but also all the solvent (usually water) molecules; however, that is such a computationally intensive enterprise that it is rarely seen as a practical option. A less detailed but more feasible approach would be to simulate the movement of only the reactant (solute) molecules. But there are serious difficulties in doing that too: It is true that we can simulate the unimpeded x -displacement of a molecule, as prescribed by the Einstein diffusion equation (A3), from time t to a later time $t + \Delta t$ with the formula

$$x_{t+\Delta t} = x_t + n\sqrt{2D\Delta t}, \quad (16)$$

where n is a sample of the normal random variable with mean 0 and variance 1; and similarly for the y - and z -displacements. The problem is that we have no way of knowing where the molecule went *between* times t and $t + \Delta t$, and therefore no way of knowing whether it might have reacted with another solute molecule during that Δt interval. Efforts to find out by taking Δt progressively smaller will initially be hampered by the inherently fractal nature of the trajectory described by Eq. (16), and will ultimately be thwarted by the fact that Eq. (16) will be *physically incorrect* unless $\Delta t \gg \tau$, where $\tau = mD/(k_B T)$ and m is the molecule’s mass.¹⁰

One way around this problem might be to use, instead of the stepping algorithm (16) which is prescribed by the Einstein model of diffusion, the stepping algorithm that is prescribed by the more accurate Langevin model of diffusion.¹⁷ The latter agrees with the former for $\Delta t \gg \tau$, but unlike the former it properly segues to ballistic motion as Δt is reduced below τ . At least for *sufficiently small* Δt , the Langevin stepping algorithm would solve the problem of not knowing where the molecules went between times t and $t + \Delta t$; because, for $\Delta t \ll \tau$, the predicted trajectory will be a nearly straight line, and a smooth interpolation will therefore be warranted. But using such a small Δt will be extremely time consuming.

A promising alternative molecule-tracking approach is the extended Green’s Function Reaction Dynamics (eGFRD) method of Takahashi, Tănase-Nicola, and ten Wolde.¹⁸ It corral each pair of solute molecules by surrounding the pair with an imaginary absorbing surface, so that absorption of either molecule on that surface will signal that the corral has been breached. eGFRD then solves the two-molecule Einstein diffusion equation to obtain the earliest of (i) the times to

absorption of each molecule, and (ii) the time to a bimolecular reaction between the molecules. It avoids difficulties arising from the small-scale deficiencies of the Einstein diffusion equation by imposing appropriate boundary conditions on the solution of that equation, the most critical of which is a special “radiation” boundary condition² at the collision surface of the two molecules if they are able to react with each other. The overall procedure allows one to advance the system from one reaction to the next without skipping over any reaction. A caveat is that the present eGFRD method appears to be computationally efficient only if the system is fairly dilute.¹⁸ As we have shown in this paper, dilute systems are typically amenable to the even more computationally efficient CME/SSA. The development of efficient, fine-scale simulation strategies is an important on-going effort, but a more detailed discussion of that topic is beyond the scope of this paper.

ACKNOWLEDGMENTS

D.T.G. was funded by the University of California, Santa Barbara under professional services agreement 130401A40, pursuant to NIH award R01-EB014877-01. L.R.P. was funded by NSF award DMS-1001012, ICB award W911NF-09-0001 from the U.S. Army Research Office, NIBIB of the NIH under award R01-EB014877-01, and U.S. DOE award DE-SC0008975. The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of these agencies.

APPENDIX: PROOF THAT THE D/h^2 VOXEL-JUMP PROBABILITY RATE AGREES WITH THE EINSTEIN DIFFUSION EQUATION IN THE LIMIT $h \rightarrow 0$

Consider a system composed of a single molecule inside a right cylindrical volume of length L and constant cross sectional area A . Let the cylinder’s axis coincide with the x -axis, and let the cylinder’s end faces be at $x = 0$ and $x = L$. Subdivide this volume into $K = L/h$ right cylindrical voxels, each of length h and cross sectional area A , by means of planes through the points $x_k \equiv k \cdot h$ ($k = 1, \dots, K$). Number the voxels so that voxel k is the one occupying the x -axis interval $[x_{k-1}, x_k] \equiv [x_k - h, x_k]$, where $x_0 = 0$. Assume that the molecule, in any voxel, will jump to a particular adjacent voxel in the next infinitesimal time interval $[t, t + dt]$ with probability $(D/h^2)dt$.

Let $q(k, t)$ be the probability that the solute molecule is in voxel k at time t . Then it follows from our assumption and the addition and multiplication laws of probability that, for any interior voxel $k \in [2, K - 1]$,

$$\begin{aligned} q(k, t + dt) = & q(k + 1, t) \cdot \left(\frac{D}{h^2}\right) dt \\ & + q(k - 1, t) \cdot \left(\frac{D}{h^2}\right) dt \\ & + q(k, t) \cdot \left[1 - 2\left(\frac{D}{h^2}\right) dt\right] + o(dt). \end{aligned} \quad (\text{A1})$$

Here, the first term on the right is the probability that the solute molecule is in voxel $k + 1$ at time t and then jumps to voxel k in the next dt ; the second term is the probability that the molecule is in voxel $k - 1$ at time t and then jumps to voxel k in the next dt ; the third term is the probability that the molecule is in voxel k at time t and does not jump across either boundary of that voxel in the next dt ; and the last term, which satisfies $o(dt)/dt \rightarrow 0$ as $dt \rightarrow 0$, recognizes that the probabilities for all routes to voxel k at time $t + dt$ via other voxels at time t are of higher order than 1 in dt . If we subtract $q(k, t)$ from both sides of Eq. (A1), divide through by dt , and then take the limit $dt \rightarrow 0$, we obtain the following exact time evolution equation for $q(k, t)$ for any interior voxel k :

$$\frac{\partial q(k, t)}{\partial t} = \left(\frac{D}{h^2}\right) [q(k - 1, t) - 2q(k, t) + q(k + 1, t)]. \quad (\text{A2})$$

The standard (Einstein) diffusion equation for the probability density function $p(x, t)$ for the x -coordinate of a solute molecule with diffusion coefficient D is

$$\frac{\partial p(x, t)}{\partial t} = D \frac{\partial^2 p(x, t)}{\partial x^2}. \quad (\text{A3})$$

Since by definition $p(x, t) \cdot A dx$ gives the probability that the x -coordinate of the solute molecule will at time t be in the infinitesimal interval dx at x , it follows by the addition law of probability that the two functions p and q are related by

$$q(k, t) = \int_{x_{k-h}}^{x_k} p(x, t) A dx \quad (k = 1, \dots, K). \quad (\text{A4})$$

For sufficiently small h , Eq. (A4) can be approximated

$$q(k, t) \doteq p(x_k, t) Ah \quad (\text{if } h \text{ is small}), \quad (\text{A5})$$

an approximation that becomes *exact* in the limit $h \rightarrow 0$. Substituting Eq. (A5) into Eq. (A2), dividing through by Ah , and then using the fact that $x_{k \pm 1} = x_k \pm h$, we convert equation (A2) for the function q into an equation for the function p :

$$\frac{\partial p(x_k, t)}{\partial t} \doteq \left(\frac{D}{h^2}\right) [p(x_k - h, t) - 2p(x_k, t) + p(x_k + h, t)]. \quad (\text{A6})$$

This equation has been derived from Eqs. (A2) and Eq. (A5), and since the former is exact while the latter becomes exact in the limit $h \rightarrow 0$, then Eq. (A6) becomes exact in the limit $h \rightarrow 0$. Since further

$$\lim_{h \rightarrow 0} \frac{p(x_k - h, t) - 2p(x_k, t) + p(x_k + h, t)}{h^2} = \frac{\partial^2 p(x_k, t)}{\partial x_k^2},$$

we conclude that the $h \rightarrow 0$ limit of Eq. (A6) is exactly the Einstein diffusion equation (A3). Thus we have proved that the molecular motion produced by the voxel-hopping probability rate D/h^2 converges, in the limit $h \rightarrow 0$, to the motion predicted by the standard diffusion equation (A3).

¹Derivations of the CME and the SSA from the definitions of the propensity functions and the state-change vectors, derivations which are mathematically rigorous in that they invoke only the laws of probability, can be found in the following two publications: Secs. 2 and 3 of D. Gillespie, “Stochastic chemical kinetics” in *Handbook of Materials Modeling*, edited by S. Yip (Springer, 2005), pp. 1735–1752; and Secs. 2.2 and 2.3 of D.

Gillespie, “Simulation methods in systems biology,” in *Formal Methods for Computational Systems Biology*, edited by M. Bernardo, P. Degano, and G. Zavattaro (Springer, 2008), pp. 125–167.

²This fact arises from a general result in quantum mechanics which is widely known as “Fermi’s Golden Rule” (although it is actually due to Dirac): For a wide class of energy-conserving transitions in atomic and molecular physics which are induced by a weak perturbation, the probability that the transition will occur in a time interval Δt will, to a good approximation, be linear in Δt , provided Δt is neither too small nor too large. The lower bound on Δt implied by the proviso is much smaller than what would be considered infinitesimally small on time scales typical of chemical reaction events. This seminal result also describes radioactive decay. But the quantum mechanical proof of this result is far from trivial, see, e.g., F. Mandl, *Quantum Mechanics* (Wiley, 1992), Chap. 9.

³D. Gillespie, *J. Comput. Phys.* **22**, 403 (1976); *J. Phys. Chem.* **81**, 2340 (1977); *Physica A* **188**, 404 (1992).

⁴F. Collins and G. Kimball, *J. Colloid Sci.* **4**, 425 (1949).

⁵D. Gillespie, *J. Chem. Phys.* **131**, 164109 (2009). An equivalent but logically neater derivation of Eq. (5) is given in D. Gillespie and E. Seitaridou, *Simple Brownian Diffusion: An Introduction to the Standard Theories* (Oxford University Press, 2012), Secs. 3.7 and 4.8. A version of Eq. (5) that appears more frequently in current literature can be obtained by multiplying the numerator and denominator of Eq. (5) by $\pi\sigma_{12}$ and then defining $k \equiv \pi\sigma_{12}^2\bar{v}_{12}q_j$:

$$a_j(x_1, x_2) = \left(\frac{4\pi\sigma_{12}D_{12}k}{4\pi\sigma_{12}D_{12} + k} \right) \cdot |\Omega|^{-1}x_1x_2.$$

The quantity in parentheses on the right becomes, in the thermodynamic limit of infinitely large Ω and molecular populations, the “rate constant” for reaction R_j in traditional deterministic chemical kinetics. It was originally obtained in Ref. 4 by imposing on the standard diffusion equation a “radiation boundary condition” with an *ad hoc* parameter k . The advantage of the more recent derivation of Eq. (5) cited above is that it uses physical reasoning which makes it unnecessary to overtly postulate a radiation boundary condition, and it provides an explicit value for k . Note that k , which is sometimes called the “microscopic association rate,” is in fact just the reaction rate constant associated with the dilute gas propensity function in Eq. (3). A rigorous derivation of the radiation boundary condition from the Kramers equation in the Langevin model of Brownian motion has been given by D. Bicout, A. Berezhkovskii, and A. Szabo, *J. Chem. Phys.* **114**, 2293 (2001). The 2009 derivation of Eq. (5) is based on the same underlying physics as the Bicout *et al.* derivation (i.e., ballistic motion on the smallest spatiotemporal scales), but it is arguably mathematically simpler and physically more transparent.

⁶M. Smoluchowski, *Z. Phys. Chem.* **92**, 129 (1917).

⁷The overall logic in the following argument is not new; see, e.g., S. A. Isaacson, *SIAM J. Appl. Math.* **70**, 77 (2009); R. Grima and S. Schnell, *Essays Biochem.* **45**, 41 (2008). However, the specific conclusions we will draw from those arguments, namely, Eqs. (8) and (9) are so far as we can tell new.

⁸In the ideal gas regime, we have $4D_{12} \gg \sigma_{12}\bar{v}_{12}q_j$, and the bimolecular propensity function (5) reduces to (3). From (3), it follows that the mean time to the next reaction of an S_1 molecule with any one of x_2 S_2 molecules in Ω is $\tau_r = (\pi\sigma_{12}^2\bar{v}_{12}q_j|\Omega|^{-1}x_2)^{-1}$. Since the molecules are now moving *ballistically*, the average time it takes the S_1 molecule to explore the volume Ω is given, not by (6), but rather by $\tau_b = |\Omega|^{1/3}/\bar{v}_1$, where \bar{v}_1 is the average speed of an S_1 molecule. The requirement for the reaction

to occur under well-mixed conditions is thus $\tau_b \ll \tau_r$. With the foregoing formulas for τ_b and τ_r , this becomes the requirement $\sqrt{\pi}(\bar{v}_{12}/\bar{v}_1)q_jx_2 \cdot \sigma_{12} \ll |\Omega|^{1/3}$. This condition can always be satisfied if the probability q_j , that an S_1 – S_2 collision will produce a reaction, is sufficiently close to 0. But if q_j is not $\ll 1$, then since $\bar{v}_{12}/\bar{v}_1 \geq 1$, the condition will be satisfied only if $\sqrt{\pi}x_2\sigma_{12} \ll |\Omega|^{1/3}$, i.e., only if the reactant molecules are dilute inside Ω .

⁹C. Gardiner, K. McNeil, D. Walls, and I. Matheson, *J. Stat. Phys.* **14**, 307 (1976). It is interesting to note that this seminal paper uses the terminology “diffusion-reaction master equation.”

¹⁰D. Gillespie and E. Seitaridou, *Simple Brownian Diffusion: An Introduction to the Standard Theories* (Oxford University Press, 2012), see Secs. 4.2–4.6.

¹¹Two different derivations of condition (14) can be found in Secs. 5.6 and 9.4 of the book cited in Ref. 10, as revised and corrected in the downloadable errata at <http://ukcatalogue.oup.com/product/9780199664504.do#>. These (revised) sections also make it clear that if condition (14) is *not* satisfied, then the movement of reactant molecules between voxels cannot be accurately described by *any* propensity function.

¹²See the paper by S. A. Isaacson cited in Ref. 7.

¹³The fact that dt here is an *infinitesimal* ensures that the occurrence of more than one reaction firing in dt will be so rare that those firings are, for all practical purposes, “mutually exclusive.” That is crucial for invoking the addition law of probability.

¹⁴D. Fange, O. Berg, P. Sjöberg, and J. Elf, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 19820 (2010).

¹⁵R. Erban and S. Chapman, *Phys. Biol.* **6**, 046001 (2009).

¹⁶S. Hellander, A. Hellander, and L. Petzold, *Phys. Rev. E* **85**, 042901 (2012).

¹⁷The stepping algorithm implied by the Langevin model of molecular diffusion is derived in the book cited in Ref. 9, and is given explicitly in that book’s Eq. (9.22). A generic version of this algorithm was first presented in a more general mathematical context in D. Gillespie, *Phys. Rev. E* **54**, 2084 (1996). The Einstein algorithm (16) evidently updates the position of the solute molecule from time t to time $t + \Delta t$ through a single formula of the form,

$$x_{t+\Delta t} = x_t + g(D, \Delta t) \cdot n;$$

here, g is a simple function of D and Δt , and n is a “unit normal” random number—i.e., a sample of the normal random variable with mean 0 and variance 1. In contrast, the Langevin stepping algorithm updates the position *and the velocity* of the solute molecule using *two coupled* formulas of the form

$$\begin{cases} v_{t+\Delta t} = f_1(\tau, \Delta t) \cdot v_t + f_2(D, \tau, \Delta t) \cdot n_\alpha, \\ x_{t+\Delta t} = x_t + g_1(\tau, \Delta t) \cdot v_t + g_2(D, \tau, \Delta t) \cdot n_\alpha + g_3(D, \tau, \Delta t) \cdot n_\beta. \end{cases}$$

Here, $\tau = mD/(k_B T)$, f_i , and g_i are explicit (and in some cases rather complicated) functions of their respective arguments; and n_α and n_β are two independent unit normal random numbers. Importantly, the Langevin updating formulas are physically accurate for *all* $\Delta t > 0$, whereas the Einstein updating formula is physically accurate only if $\Delta t \gg \tau$. For $\Delta t \ll \tau$, the Langevin updating formula for x reduces (approximately) to $x_{t+\Delta t} = x_t + v_t \cdot \Delta t$, and in that case there is no question about what the value of $x_{t'}$ was for any $t' \in (t, t + \Delta t)$; it was $x_{t'} = x_t + v_t \cdot (t' - t)$.

¹⁸K. Takahashi, S. Tănase-Nicola and P. R. ten Wolde, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 2473 (2010), see its Supplemental Information section.