Behavior Research Methods, Instruments, & Computers
2001, 33 (4), 443-456

A comparison of four methods for simulating the diffusion process

FRANCIS TUERLINCKX
University of Leuven, Leuven, Belgium

ERIC MARIS
University of Nijmegen, Nijmegen, the Netherlands

ROGER RATCLIFF
Northwestern University, Evanston, Illinois

and

PAUL DE BOECK
University of Leuven, Leuven, Belgium

Four methods for the simulation of the Wiener process with constant drift and variance are described. These four methods are (1) approximating the diffusion process by a random walk with very small time steps; (2) drawing directly from the joint density of responses and reaction time by means of a (possibly) repeated application of a rejection algorithm; (3) using a discrete approximation to the stochastic differential equation describing the diffusion process; and (4) a probability integral transform method approximating the inverse of the cumulative distribution function of the diffusion process. The four methods for simulating response probabilities and response times are compared on two criteria: simulation speed and accuracy of the simulation. It is concluded that the rejection-based and probability integral transform method perform best on both criteria, and that the stochastic differential approximation is worst. An important drawback of the rejection method is that it is applicable only to the Wiener process, whereas the probability integral transform method is more general.

The diffusion process has become quite popular in cognitive psychology during the last two decades (Hanes & Schall, 1996; Luce, 1986; Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff, Van Zandt, & McKoon, 1999). In the study of complicated stochastic models such as the diffusion process, simulations of the model play an important part for three reasons; they enable one to (1) understand the basic characteristics of the model, (2) generate data in order to evaluate estimation techniques for the model parameters, and (3) derive predictions to test the performance of the model.

The problem with simulating the diffusion process is that it is a stochastic process with a continuous state and time space. Roughly, we can distinguish two classes of simulation methods on the basis of accuracy of approximation of the continuous diffusion process. First, there are methods in which a discretization of the continuous state and/or time space of the diffusion model is needed. In a second class of methods, no discretization is made, and the accuracy of the approximation of the continuous state and time space is limited only by the precision of the floating-point implementation. A method from the second class is expected to give always more exact results than any method from the first class. Therefore, it can be used as a benchmark to evaluate the accuracy of methods from the first class.

Apart from differences in accuracy, simulation methods may also differ with respect to their simulation speed. If a huge number of simulations of the model is needed, as, for instance, in Markov chain Monte Carlo applications (Tanner, 1996), speed becomes an important issue.

The aim of this paper is to present an overview and evaluation of techniques to simulate the diffusion process with constant drift rate and variance, and two absorbing boundaries (Cox & Miller, 1970). Such a process is also called a Wiener process with absorbing boundaries, or a Brownian motion with drift rate and absorbing boundaries. Hence, if we speak of a diffusion process, we mean the Wiener process. In a separate section, we will indicate whether or not, and if so, how the proposed simulation methods can be adapted to simulate other types of diffusion processes, but our major focus is on the Wiener process.

The first author was a research assistant of the Fund for Scientific Research (Flanders). The research in this paper was also supported by BOF Grant GOA/02/2000 from the University of Leuven. We thank the editor and two anonymous reviewers for their comments and suggestions. Address all correspondence to F. Tuerlinckx, Department of Statistics, Columbia University, 2990 Broadway MC 4403, New York, NY 10027 (e-mail: tuerlinckx@neyman.stat.columbia.edu).

 Copyright 2001 Psychonomic Society, Inc.

Four methods for simulating the diffusion process are studied here. The first method, used most commonly, is based on the approximation of the diffusion process by a random walk. The second method is built on a rejection algorithm for drawing directly from the first-passage time densities of the diffusion process. The third method is based on simulating a stochastic differential equation characterizing the diffusion process. The fourth method uses the probability integral transform.

These methods have been described in other places before. The theory for the first method is described in almost any textbook on stochastic models (e.g., Feller, 1968). The reference to the second method (Lichters, Fricke, & Schnakenberg, 1995), however, is not readily accessible to psychologists. Moreover, the method is described only for the case with zero drift rate, and as a consequence, it is not directly applicable to the most common diffusion process used in psychology. We therefore present an elaborated and adapted version of the original algorithm. The third method can be derived in a straightforward manner from the definition of the diffusion process and is elaborated in, for instance, Bouleau and Lépingle (1994) and Fahrmeir (1976) (these references contain only information about the case without absorbing boundaries). The fourth method can be found in almost any textbook on statistics (see, e.g., Mood, Graybill, & Boes, 1974).

The structure of the remainder of this paper is as follows. First, a short overview is given of the type of diffusion process that is considered here. Second, the four simulation methods will be presented. Third, these methods will be evaluated with respect to speed and accuracy of simulation. Fourth, the generality of the methods will be discussed, and finally, some conclusions will be presented.

OVERVIEW OF THE DIFFUSION PROCESS

In this section, we provide a brief overview of the basic features of the diffusion process and introduce some notation. A more technical and complete description of the model and its mathematical properties can be found in specialized references such as Cox and Miller (1970), Karlin and Taylor (1981), or Ross (1996).

The diffusion process is a stochastic process that develops through time, and it can be represented by a continuous random variable $X(t)$ ($t \geq 0$), denoting the position of the process in the state space at time t($t \geq 0$; t is also continuous). In case of a diffusion process with absorbing boundaries 0 and a, the state space is restricted to the interval $[0, a]$. The central part of a stochastic process is the transition probability density function $p(x_0, t_0; x, t)$, which is the density for $X(t) = x$ given that the process was at time point t_0 ($t_0 < t$) at position x_0 and that the process did not reach one of the boundaries by time t.

Assume that the diffusion process starts at z ($0 < z < a$) at time 0. The transition probability density function $p(z, 0; x, t)$ should then satisfy the Kolmogorov equations, also called the forward and backward differential equations (Cox & Miller, 1970). Derivations of the Kolmogorov equations are presented in Cox and Miller and in Ross (1996), and their method of solution is explained in Cox and Miller. The forward equation is

$$\frac{1}{2}\sigma^2 \frac{\partial^2 p(z,0;x,t)}{\partial x^2} - \mu \frac{\partial p(z,0;x,t)}{\partial x} = \frac{\partial p(z,0;x,t)}{\partial t}, \tag{1}$$

which has to be solved for $p(z, 0; x, t)$ subject to the conditions

$$\begin{aligned} p(z,0;x,0) &= \delta(x-z) \\ p(z,0;a,t) &= 0 \\ p(z,0;0,t) &= 0. \end{aligned} \tag{2}$$

The first condition is an initial value condition, which means that, at time 0, all mass of the transition probability density function is concentrated at the starting point z. The function $\delta(x - z)$ is the Dirac delta function, a degenerate probability density with all its mass concentrated at 0. The last two conditions are boundary conditions, which impose the restriction that the density equals zero at the boundaries, so that terminated processes disappear from the transition density. The stochastic process described by Equations 1 and 2 is called a diffusion process with constant drift rate, constant variance, and absorbing boundaries, or a Wiener process with absorbing boundaries. Since the exact form of the transition probability density function is not needed in this paper, it will not be presented here.

The parameters μ and σ^2 are the infinitesimal mean and variance of the process: μ is the mean displacement and σ^2 the variance of the displacement of $X(t)$ given that $X(t - \tau) = x$, where τ approaches to 0. The parameter μ is called the drift rate of the diffusion process.

Two important events and corresponding random variables originating from the diffusion process $\{X(t); t \geq 0\}$ are of particular interest for simulating the diffusion process in the context of psychological research. First, there is the event that absorption occurs at the upper boundary. A random variable Y takes the value 1 if absorption occurs at the upper boundary and value 0 if absorption occurs at the lower boundary. The second important event is $[T \leq t \,|\, Y = 1]$, which is the event that the absorption time T is smaller than some value t, given that absorption happens in the upper boundary. We will consider also the event $[T \leq t \,|\, Y = 0]$.

In psychological applications, it is impossible to observe the full sample path $\{X(t); t \geq 0\}$ from start until absorption. Only the boundary of absorption (i.e., the choice response) and the time that it takes until absorption (i.e., the choice response time) are observed. Therefore, it is not required that a simulation algorithm return a complete sample path. Thus, for the simulation, we focus on the random variables Y and $[T \leq t \,|\, Y = 1]$.

Next, the density functions for the random variables defined above will be given. First, the probability of hitting the upper boundary equals

$$\Pr(Y=1)=\frac{\exp\left(-\frac{2z\mu}{\sigma^2}\right)-1}{\exp\left(-\frac{2a\mu}{\sigma^2}\right)-1} \tag{3}$$

From the fact that the probability of absorption is a sure event, it follows that $\Pr(Y=0)$ equals $1-\Pr(Y=1)$. The probability of hitting the lower boundary can also be obtained from Equation 3 by replacing μ by $-\mu$ and z by $(a-z)$. This operation interchanges the role of the two absorbing boundaries but yields an otherwise identical diffusion process.

For the random variable T, we will not consider the probability $\Pr(T \le t \mid Y=1)$, but only present the conditional density of first-passage times at the upper boundary:

$$f_{T|Y}(t \mid Y=1) = \frac{\pi\sigma^2}{a^2}\exp\left(\frac{(a-z)\mu}{\sigma^2}-\frac{\mu^2}{2\sigma^2}t\right)\frac{1}{\Pr(Y=1)} \times\sum_{m=1}^{\infty} m\sin\left(\frac{\pi m(a-z)}{a}\right)\exp\left(-\frac{1}{2}\frac{\pi^2\sigma^2m^2}{a^2}t\right), \tag{4}$$

where $f(t \mid Y=1)$ is defined for all $t>0$. The other conditional density of first-passage times, $f(t \mid Y=0)$, can be obtained from Equation 4 by replacing μ by $-\mu$ and z by $(a-z)$.

Typically, the diffusion process is applied to data from a two-alternative forced choice (2AFC) paradigm. However, the response time that is observed consists of more than just the decision time. It is often assumed that two components contribute to the response time: one that involves the decision process, which can be modeled by the diffusion process, and another that involves all the other processes (encoding, response preparation, and execution). A complete model for response times in a 2AFC paradigm should also take into account the second component. For simplicity, it is often assumed that all processes that are unrelated to the decision process consume a constant time T_{er}. If this constant time is added to the decision times, the total response time RT becomes

$$RT = T + T_{er}. \tag{5}$$

Under the assumption of a constant T_{er}, Equation 4 can be easily adapted: In particular, t should be replaced by the difference $(t-T_{er})$. Equation 3 does not change, of course.

If we would make the more realistic approach that T_{er} is a random variable (see, e.g., Ratcliff & Tuerlinckx, 2001), then there are no explicit formulas for the conditional densities, not even for simple choices of densities for T_{er} (e.g., uniform, normal, or exponential). However, once we have a simulation method for the diffusion process, the simulation of a model with a random T_{er} poses no additional problems. For simplicity, in the remainder of the paper it will be assumed that T_{er} is 0.

In some applications of the diffusion process (Ratcliff, 1978; Ratcliff et al., 1999), it is assumed that the drift rate μ and the starting point z show some trial-to-trial variability. The trial-to-trial variability is needed to limit accuracy (otherwise it grows toward 1 if boundary separation increases) and to fit the error response times (Ratcliff & Rouder, 1998). The simulation of such a process can be accomplished by drawing the drift rate and starting point from the appropriate distributions and then simulating the diffusion process with the drawn values. In the following, we will assume that μ and z are both constant, unless specified otherwise.

FOUR METHODS FOR SIMULATING THE DIFFUSION PROCESS

In this section, the four methods are explained and an algorithm for each of them is outlined.

A Random Walk Approximation

In many texts on stochastic processes, the diffusion process with constant drift rate and absorbing boundaries is considered as the continuous version of the random walk process with absorbing states (see, e.g., Cox & Miller, 1970; Feller, 1968). The diffusion process can be derived mathematically by constructing a random walk with very small displacements and small time intervals and by letting the length of both steps and time intervals approach zero. This limiting property is useful for the simulation of the diffusion process. In particular, it shows that a simulation from the diffusion process can be obtained by simulating a random walk with small time intervals and displacements.

Assume a random walk with state space $\{0, \Delta, 2\Delta, \ldots, z-\Delta, z, z+\Delta, \ldots, a\}$, in which 0 and a are the absorbing states and z is the starting point. At every time interval τ, a change occurs: a displacement Δ with probability p and a displacement $-\Delta$ with probability $q=1-p$. Assume furthermore that the following equalities hold:

$$p=\frac{1}{2}\left(1+\frac{\mu\sqrt{\tau}}{\sigma}\right)$$
$$q=\frac{1}{2}\left(1-\frac{\mu\sqrt{\tau}}{\sigma}\right)$$
$$\Delta=\sigma\sqrt{\tau}. \tag{6}$$

Under these conditions, if τ converges to zero, the random walk converges to a diffusion process with drift rate μ and variance σ^2. From a given triplet (μ, σ, τ), a triplet (p, q, Δ) can be computed from Equation 6, and a random walk approximation is easily obtained.

The accuracy of the algorithm depends on the value of τ, which will be denoted as precision in the following. If τ is large, the approximation of the diffusion model will not be very accurate, but the simulation will be fast. On the other hand, if τ is small, the discrete approximation of the

concrete process will be accurate, but the simulation time may take too long to be of practical interest. A good value for τ should return an accurate approximation and the simulation should not take too much time.

To investigate how the simulation time depends on the value of τ, results from the theory of the random walk can be used. For a random walk with time interval τ and state space $\{0, \Delta, 2\Delta, \ldots, z-\Delta, z, z+\Delta, \ldots, a\}$, the expected number of steps before absorption, $E(N)$, in either of the boundaries is equal to (Feller, 1968)

$$E(N) = \begin{cases} \dfrac{\frac{z}{\Delta}}{q-p} - \dfrac{\frac{a}{\Delta}}{q-p} \times \dfrac{1-\left(\frac{q}{p}\right)^{\frac{z}{\Delta}}}{1-\left(\frac{q}{p}\right)^{\frac{a}{\Delta}}} & \text{if } \mu \neq 0 \\ \dfrac{z}{\Delta}\left(\dfrac{a}{\Delta} - \dfrac{z}{\Delta}\right) & \text{if } \mu = 0. \end{cases} \tag{7}$$

With the use of the equality $\Delta = \sigma\sqrt{\tau}$, it is easy to show that for the case $\mu = 0$, if τ is divided by a factor K, $E(N)$ is multiplied by the same factor K. Stated differently, $E(N)$ is of order $O(\tau^{-1})$. If $\mu \neq 0$, this is more difficult to prove, but it can easily be checked numerically. These results show that if the simulation is more accurate, it is also expected to take more time.

To summarize: In the random walk approximation, the continuous time line is divided into small time intervals, and the position of the process at each time point is simulated. The magnitude of the displacement Δ of the process is a monotonic increasing function of magnitude of the time intervals ($\Delta = \sigma\sqrt{\tau}$). Thus, smaller time intervals lead to smaller displacements. A drawback of the random walk approximation is that the expected simulation time increases with the accuracy of the approximation.

A Rejection-Based Algorithm

The second method differs from the first in that there is no attempt to simulate the sample path $\{X(t);\ t \geq 0\}$. Instead, the second method directly simulates the time that it would take for the process to travel a fixed distance R from its starting point. This idea was developed by Lichters et al. (1995). They described an algorithm for the zero drift rate case, and this algorithm will be extended here to the case in which μ can take any real value.

The algorithm of Lichters et al. (1995) makes use of a rejection algorithm for simulating a symmetric diffusion process (i.e., $z = a/2$) with absorbing boundaries. With this algorithm, it is possible to obtain an exact draw from the first-passage density, given absorption at one of the boundaries. This rejection algorithm is then used as a component of a more general algorithm for the simulation from an asymmetric diffusion process ($z \neq a/2$).

The algorithm to simulate a diffusion process will be described in two parts. First, we will discuss the rejection algorithm for the symmetric diffusion process. Second, we will give an outline of a simulation algorithm for the more general asymmetric case.

Part 1: Rejection Method for the Symmetric Wiener Process

In the special case of a symmetric Wiener process, the conditional response time density in Equation 4 simplifies remarkably. See Equation 8 below.

Thus, for a symmetric Wiener process, the conditional first-passage time distributions are the same. This equation can be simplified further by noting that

$$\sin\left(\frac{\pi m}{2}\right) = \begin{cases} 1 & \text{if } m = 1+4d \\ 0 & \text{if } m = 2d \\ -1 & \text{if } m = 3+4d, \end{cases} \tag{9}$$

where d is a positive integer. Equation 8 can now be rewritten as seen in Equation 10 below.

For Equation 10, it is shown in Appendix A that the following inequality holds:

$$f_{T|Y}(t \mid Y=1) \leq M\lambda \exp(-\lambda t) \tag{11}$$

$$f_{T|Y}(t \mid Y=1) = f_{T|Y}(t \mid Y=0) = \frac{\pi\sigma^2}{a^2}\left[\exp\left(\frac{a\mu}{2\sigma^2}\right) + \exp\left(-\frac{a\mu}{2\sigma^2}\right)\right] \times \exp\left(-\frac{\mu^2}{2\sigma^2}t\right)\sum_{m=1}^{\infty} m \sin\left(\frac{\pi m}{2}\right)\exp\left(-\frac{1}{2}\frac{\pi^2\sigma^2 m^2}{a^2}t\right). \tag{8}$$

$$f_{T|Y}(t \mid Y=1) = \frac{\pi\sigma^2}{a^2}\left[\exp\left(\frac{a\mu}{2\sigma^2}\right) + \exp\left(-\frac{a\mu}{2\sigma^2}\right)\right]\exp\left(-\frac{\mu^2}{2\sigma^2}t\right) \times \sum_{n=0}^{\infty}(2n+1)(-1)^n \exp\left(-\frac{1}{2}\frac{\pi^2\sigma^2(2n+1)^2}{a^2}t\right). \tag{10}$$

where

$$M = \frac{\frac{\pi\sigma^2}{a^2}\left[\exp\left(\frac{a\mu}{2\sigma^2}\right)+\exp\left(-\frac{a\mu}{2\sigma^2}\right)\right]}{\left(\frac{\mu^2}{2\sigma^2}+\frac{\pi^2\sigma^2}{2a^2}\right)} \quad (12)$$

and

$$\lambda = \left(\frac{\mu^2}{2\sigma^2}+\frac{\pi^2\sigma^2}{2a^2}\right). \quad (13)$$

Since the inequality in Equation 11 holds for any t, it may be used for a rejection algorithm for simulating from $f_{T|Y}(t|Y=0)$ or $f_{T|Y}(t|Y=1)$. The rejection algorithm is explained in Appendix B (see also Press, Flannery, Teukolsky, & Vetterling, 1986; Tanner, 1996).

In the sampling algorithm for a symmetric diffusion process, it is first decided whether a realization from $f_{T|Y}(t|Y=1)$ or $f_{T|Y}(t|Y=0)$ has to be simulated. Next, a realization is drawn from $f_{T|Y}(t|Y=1)$ or $f_{T|Y}(t|Y=0)$, dependent on the value for Y. Assume for simplicity that the outcome of the first step of the algorithm is 1 (i.e., $Y=1$). However, notice that the algorithm is the same for the case $Y=0$, since $f_{T|Y}(t|Y=1)=f_{T|Y}(t|Y=0)$.

A rejection algorithm is initiated by sampling from a candidate-generating distribution. In our case, this is the exponential distribution with parameter λ as defined in Equation 13. Sampling a realization t^* from an exponential distribution involves the following computation:

$$t^* = -\frac{1}{\lambda}|\ln(1-u)| \quad (14)$$

where u is a uniform random number from $U(0,1)$ and λ is defined as in Equation 13. The draw t^* from the exponential distribution is a candidate draw from $f_{T|Y}(t|Y=1)$ and is accepted as a draw from that distribution if

$$v \le \frac{f_{T|Y}(t^*|Y=1)}{M\lambda\exp(-\lambda t^*)} \quad (15)$$

where v is another random number from $U(0, 1)$, and M is defined as in Equation 12. After some rewriting, the condition in Equation 15 becomes

$$v \le 1+(1-u)^{-F}\sum_{n=1}^{\infty}(2n+1)(-1)^n(1-u)^{F(2n+1)^2}, \quad (16)$$

where

$$F = \frac{\pi^2\sigma^4}{\pi^2\sigma^4+\mu^2a^2}.$$

Convergence of the infinite sum in Equation 16 is fast. The number of terms will typically be lower than 15. Convergence is worst if u is close to zero.

Since the presented algorithm is a rejection algorithm, the acceptance rate is equal to $1/M$. The acceptance rate is

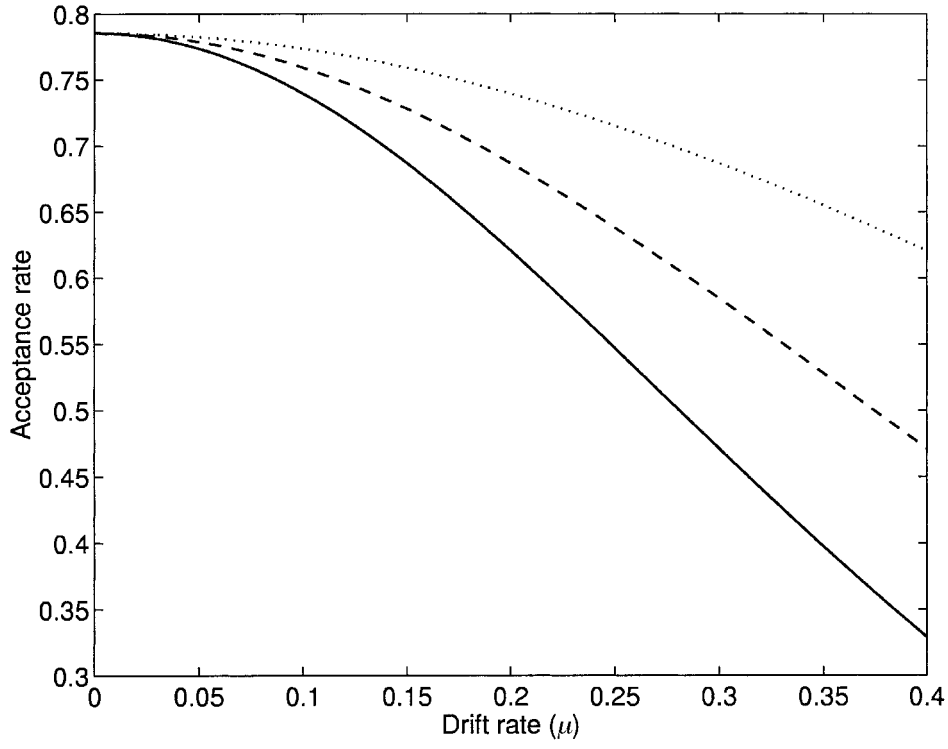

Figure 1. The acceptance rate as a function of drift rate for three values of a. The solid line represents the case $a = 0.16$; the dashed line, $a = 0.12$; and the dotted line, $a = 0.08$ ($\sigma = 0.1$).

not constant, since M is a function of μ, a, z, and σ^2 (see Equation 12). If $\mu = 0$, then the acceptance rate equals $\pi/4 = 0.785$, and for $\mu \neq 0$, it will be always lower than $\pi/4$. In Figure 1, the acceptance rate as a function of drift rate μ (on the abscissa) is shown for three different values of a (σ is always equal to 0.1). Both the range of μ and the different values of a are realistic values, in the sense that one may encounter them in real applications (given that $\sigma =$ 0.1); they produce decision times around 1 sec. It can be seen that in the most extreme case in Figure 1 ($a = 0.16$ and $\mu = 0.4$), the overall acceptance rate is still about 0.3.

In conclusion, a symmetric diffusion process can be simulated very easily in two steps: (1) Determine the boundary of absorption, based on the probabilities of reaching each of the boundaries; and (2) draw a response time from the conditional response time density, given the boundary of absorption.

Part 2: Rejection-Based Method for the General Wiener Process

We now proceed with the general Wiener process, without the requirement of a starting point symmetrically between the absorbing boundaries. The rejection method described in the previous subsection will be applied one or more times in order to obtain a realization from the more general process. In particular, given the starting point, a new diffusion process is defined with symmetric absorbing boundaries that lie between the original boundaries, and the rejection method is applied. If the simulated value is located at one of the original absorbing boundaries, the process stops, and if not, again a new symmetric diffusion process is defined, and this goes on until absorption occurs at one of the original boundaries. Each time a fixed distance is traveled (the distance in the symmetric diffusion process from the starting point to the absorbing boundary) and how long it takes to travel the fixed distance is simulated.

The central property of the diffusion process on which the algorithm is based is the Markov property: Once it is determined that the process is at a certain state, $x(t)$ at time t, the future behavior of the process does not depend on how state $x(t)$ was reached.

The algorithm is explained further in detail on basis of a flow chart (Figure 2) and an example (Figure 3). First, the variables of the algorithm are initialized. The integer k is a counter, $t^{(0)}$ is the simulated time at step 0, and $z^{(0)}$ is the starting point of the first symmetric diffusion process. At the start of the algorithm, $z^{(0)}$ equals z, the starting point of the original diffusion process. In the example in Figure 3, the absorbing boundaries of the general diffusion process are 0 and 0.12, and the starting point is 0.04. Therefore, the starting point of the first symmetric diffusion process is also 0.04. The other two variables, w and r, are position and time variables. They will be assigned values in the course of the algorithm.

Next, a new diffusion process with the starting point exactly in between the two absorbing boundaries is defined. The starting point has already been initialized, and the lower and upper boundary ($l^{(k)}$ and $u^{(k)}$) are now chosen so that one of them equals an absorbing boundary of the original diffusion process, in particular the absorbing boundary closest to the new starting point. The drift rate and variance remain unchanged. In the example, the absorbing boundaries of the original diffusion process are 0 and 0.12. Therefore, with starting point at 0.04, the new absorbing boundaries for step 0 are 0 (the most nearby original absorbing boundary) and 0.08. Notice that the starting point 0.04 lies exactly halfway between these boundaries. New absorbing boundaries are represented by a dotted line in Figure 3.

In a third step, the rejection algorithm as described in the previous subsection is applied to the new symmetrical diffusion process. In the flow chart, the rejection algorithm is denoted by REJ, and the outcome of the simulation assigns values to w and r. The position w must equal $l^{(k)}$ or $u^{(k)}$. After this, the total simulation time is updated. In the example, the first application of the rejection algorithm results in $w = 0.08$ and $r = 0.20$ and the observation is denoted by "x" in Figure 3.

If the simulation of the symmetrical diffusion process leads to an absorption at one of the boundaries of the original diffusion process, then the simulation stops. In that case, the final boundary of absorption is known, as is the total simulated first-passage time. However, if absorption occurs at the other boundary, the rejection algorithm will be applied another time but with the new starting point w, the current position in the state space. In the example, the rejection algorithm has to be applied another time, since w is not equal to 0.0. The new starting point, $z^{(1)}$, is now equal to 0.08, and the new absorbing boundaries are 0.04 (dotted line) and 0.12. The second application of the rejection algorithm ends up in the upper boundary of the original diffusion process, so that the algorithm stops. The simulated time in this step was 0.22, so that the total simulated diffusion time is equal to 0.42.

The advantage of the rejection-based method is that the simulation time is not proportional any more to the physical time to run through the process. Another advantage over the random walk method is that it does not involve a discrete approximation to the continuous diffusion process. Therefore, the accuracy of the simulation is bounded only by the numerical accuracy of the computer system on which the simulation takes place.

It is difficult to obtain theoretical results about the simulation time for the rejection-based algorithm in Figure 2. The simulation time depends on the acceptance rate for the input parameter values and the number of times that the absorbing boundaries of the original diffusion process do not coincide with the chosen boundaries from the rejection algorithm. However, both the acceptance rate and the expected number of times that a subdivision of the original diffusion process in a symmetric subprocess has to be made depend on so many things that it is difficult to make general assertions.

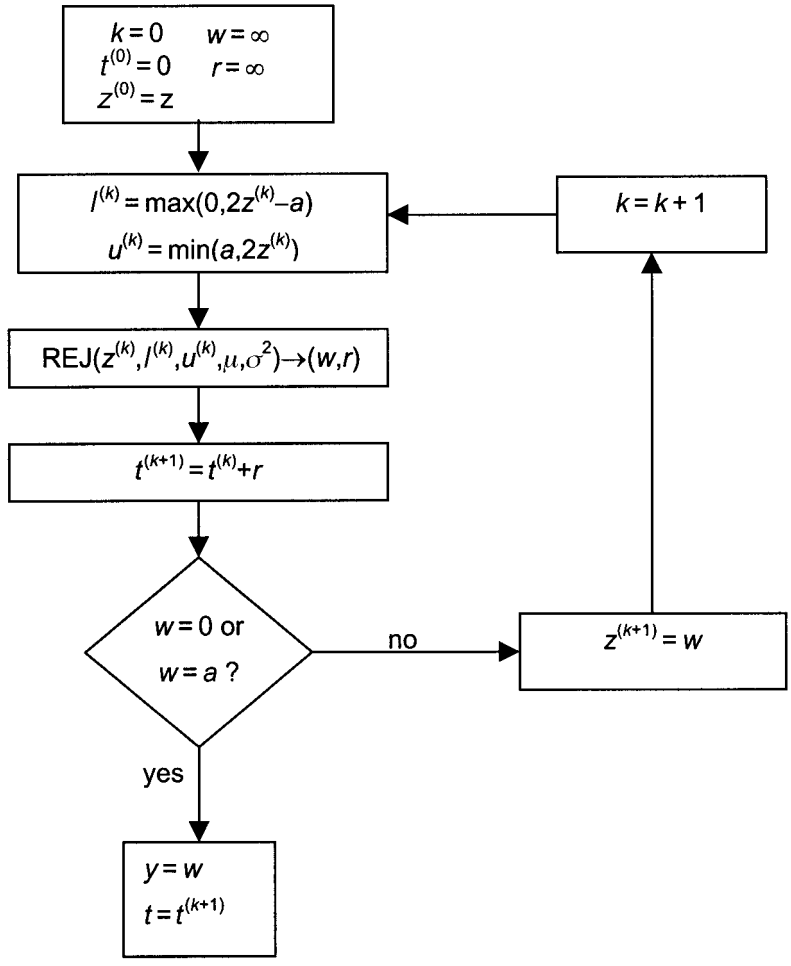

Figure 2. Flow chart for the rejection-based algorithm to simulate from a diffusion process.

Simulation of the Stochastic Differential Equation

Another way of defining the diffusion process is via a stochastic differential equation (SDE; Karlin & Taylor, 1981; Smith, 2000) defined for $X(t)$,

$$dX(t) = \mu dt + \sigma dB(t), \tag{17}$$

under the condition that $X(0) = z$ and where $B(t)$ is a standard Brownian motion process. For each $t > 0$, $B(t)$ has a standard normal distribution with $E[B(t)] = 0$ and $\mathrm{Var}[B(t)] = t$. Moreover, $\mathrm{Cov}[B(t),B(s)] = \min(t,s)$.

It can be proved that the sample path $\{X(t);\ t \geq 0\}$ of a diffusion process is (almost surely) continuous but the derivative $dX(t)/dt$ (for any $t \geq 0$) does not exist (see Karlin & Taylor, 1981, p. 341, for an illustration). Therefore, the change in $X(t)$ in a small time interval should be represented by a stochastic differential equation instead of by the usual deterministic differential equation. Intuitively, Equation 17 states that the change in $X(t)$ in a small time interval dt is normally distributed with mean μdt and variance $\sigma^2 dt$. The latter follows from inserting $\mathrm{Var}[B(t)] = t$ and $\mathrm{Cov}[B(t),B(s)] = \min(t,s)$ in the usual formula for the variance of the difference $dB(t) = B(t + dt) - B(t)$.

A discrete version of the continuous SDE in Equation 17 can be defined as follows:

$$x(t_{k+1}) - x(t_k) = \mu\left(t_{k+1} - t_k\right) + \sigma\left[b(t_{k+1}) - b(t_k)\right]. \tag{18}$$

Assume that the time intervals are constant, or, in other words, $\Delta t_k = \tau$ (for all $k = 1, 2, \ldots$). It follows from the definition of the Brownian motion that the difference $\Delta b_k = b(t_{k+1}) - b(t_k)$ is normally distributed with mean 0 and variance τ.

Like the random walk approximation, the discrete version of the SDE is an approximation to the diffusion process. The stochastic process that is defined in Equation 18 can also be seen as a random walk for time intervals $0, \tau, 2\tau, \ldots$ with normally distributed displacements. The simulation process stops if the $x(t_k) \geq a$ or $x(t_k) \leq 0$ for a certain k. Smith (1990) has shown that in case that τ approaches 0, this random walk process converges to a diffusion process with drift rate μ and variance σ^2.

$$E(N) \approx \begin{cases} \dfrac{1}{\mu\tau} \times \dfrac{a-(a-z)\exp\left(\dfrac{2\mu z}{\sigma^2}\right)-z\exp\left(-\dfrac{2\mu(a-z)}{\sigma^2}\right)}{\exp\left(-\dfrac{2\mu(a-z)}{\sigma^2}\right)-\exp\left(\dfrac{2\mu z}{\sigma^2}\right)} & \text{if } \mu \neq 0 \\ \dfrac{(a-z)z}{\sigma^2\tau} & \text{if } \mu = 0. \end{cases} \tag{19}$$

Wald (1947) gives the approximate expected number of time intervals needed to reach a boundary, as shown in Equation 19.

As was the case for the ordinary random walk, the expected simulation time is $O(\tau^{-1})$.

A problem with this method and the stopping rule is that, in the final simulation step, the state variable always exceeds the boundary, and this leads to inaccuracies in the approximation of the theoretical diffusion process, which can be quite substantial as Luce (1986) and Smith (1990) have noted. The reason why Equation 19 gives only the approximate number of steps is that in its derivation the boundary excess is ignored (Cox & Miller, 1970).

Some ad hoc methods have been presented to correct for this excess over a boundary. The correction method of Heath and Kelly (1988) yields the best results. The method corrects for an excess by widening the boundary separation:

$$\begin{aligned} a &\rightarrow a + \zeta\sigma\sqrt{\tau} \\ 0 &\rightarrow -\zeta\sigma\sqrt{\tau}, \end{aligned} \tag{20}$$

where ζ is a constant whose optimal value is 0.586 (Smith, 1990).

Probability Integral Transform Method

To explain the probability integral transform method, assume that for a random variable X, $g_X(x)$ is a probability density and that $G_X(x)$ is its associated cumulative distribution function. Suppose a new random variable $U = G_X(X)$ is defined. It is easy to prove that U is uniformly distributed over [0, 1] (see Appendix C). Conversely, if a realization u from a uniform density is obtained, its transform, $x = G_x^{-1}(u)$, is a draw from the density $g_X(x)$ (see also Appendix C).

To apply this theory to the case of the diffusion process, let us define $F_y(t) = F_{T|Y}(t \mid Y = y)$ as the conditional cumulative distribution function for response y ($y = 0, 1$). The formula for $F_1(t)$ is as shown in Equation 21.

Obtaining a realization from the diffusion process proceeds, then, as follows. First, decide on the boundary of absorption by using the probabilities $\Pr(Y = 0)$ and $\Pr(Y = 1)$; this renders an observed value y. Second, draw a uniform random number u and compute $t = F_y^{-1}(u)$. The pair (y, t) is then a realization from the diffusion process.

Computing the inverse of the cumulative distribution function has to be done numerically. For this purpose, the

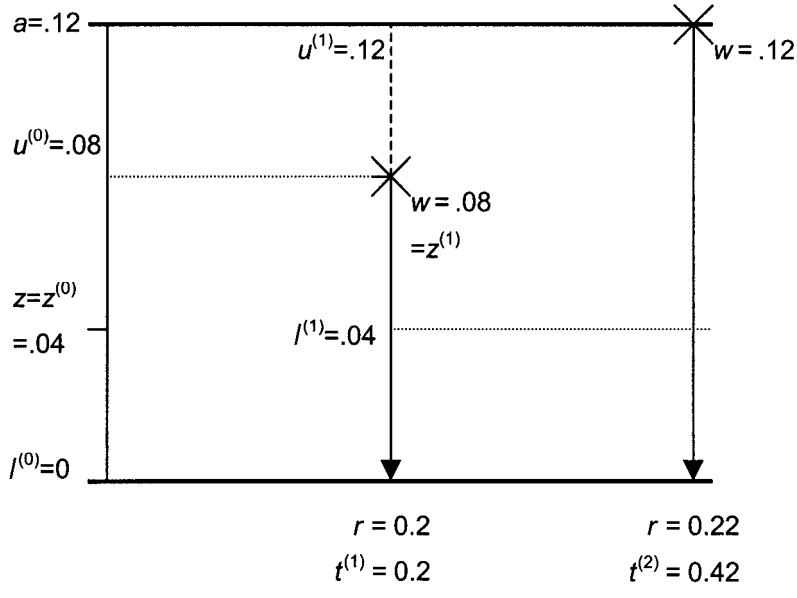

Figure 3. An example of the rejection-based method for the simulation of the diffusion process.

$$F_1(t) = 1 - \frac{1}{\Pr(Y=1)} \frac{\pi\sigma^2}{a^2} \exp\left(\frac{(a-z)\mu}{\sigma^2}\right) \times \sum_{m=1}^{\infty} \frac{2m \sin\left(\frac{\pi m(a-z)}{a}\right) \exp\left[-\frac{1}{2}\left(\frac{\mu^2}{\sigma^2} + \frac{\pi^2 m^2 \sigma^2}{a^2}\right)t\right]}{\left(\frac{\mu^2}{\sigma^2} + \frac{\pi^2 m^2 \sigma^2}{a^2}\right)}. \quad (21)$$

function $F_y(t)$ can be approximated by using a grid of points (nodes). Denote the sequence of N grid points by $t_1,\ldots,t_N$. For each grid point t_i, the corresponding function value $h_i = F_y(t_i)$ is calculated. A uniformly distributed random variable u is drawn and the associated realization t is determined as follows:

$$t = \begin{cases} t_1 & \text{if } u \le F_y(t_1) \\ t_{i+1} & \text{if } F_y(t_i) < u \le F_y(t_{i+1}), i = 1,\ldots,N-1 \\ t_N & \text{if } u > F_y(t_N). \end{cases} \quad (22)$$

If N is taken large enough, the approximation will be highly accurate. More complex approximations are also possible, such as linear interpolation between two grid points.

EVALUATION OF THE SIMULATION METHODS

In this section, a Monte Carlo study is presented to evaluate and compare the three simulation methods on two criteria: speed and accuracy. The speed of the random walk and SDE approximations depends on the precision parameter τ. If a high precision is used (a small value of τ), a lot of small simulation steps may be needed, and that will slow down the methods. Likewise, the speed of the probability integral transform method depends on the number of tabulated grid points, N. For the probability integral transform method, setting up the table with function values can be a large cost, but once that table is available, a value is easily determined. On the other hand, the speed of the rejection-based method depends on the amount of time that it takes to evaluate Equation 16. These considerations show that it is difficult to make predictions about which method will be faster.

Concerning accuracy, some firm statements can be made, since the rejection-based method is an exact simulation method and hence is by definition better than every approximation method; therefore it functions as a benchmark (it is needed to evaluate the accuracy of the three other methods). It is interesting, however, to compare the performance of the three other methods with the performance of the rejection-based algorithm. In particular, we are interested in the precision that is needed for the other methods to be almost as accurate as the rejection-based method.

To avoid superfluous notation, we will abbreviate the names of the three simulation methods. The random walk approximation will be denoted as the RW method, the rejection-based method as the RB method, the stochastic differential equation approximation as the SDE method, and the probability integral transform method as the PIT method.

Design and Procedure

Three independent variables are manipulated in the Monte Carlo experiment: drift rate (μ), combination of boundary separation and starting point ($a|z$), and, if used by the algorithm, the precision τ. First, there are five levels for the drift rate μ: 0.0, 0.1, 0.2, 0.3, and 0.4. Second, six levels for the combination of absorbing boundary and starting point are chosen (denoted as $a|z$): 0.08|0.04, 0.12|0.06, 0.16|0.08, 0.08|0.02, 0.12|0.04, and 0.16|0.04. Third, three levels of precision τ are chosen for the RW and SDE methods: 0.01, 0.001, and 0.0001 (ordered from low to high precision). For the PIT method, three levels of precision are chosen (operationalized in number of grid points): 500, 5,000, and 50,000.

The three independent variables are completely crossed. For each cell of the design, 20 samples of 3,000 observations are simulated. All four methods are programmed in FORTRAN, and the Monte Carlo experiment is run on a personal computer with a Pentium III 600-MHz processor.

For each criterion, speed and accuracy, we examine two dependent variables. First, with respect to speed, we compute the time that it takes to draw 3,000 observations. Second, the mean number of steps taken for the RW and SDE methods to reach an absorbing barrier is computed, and, for the RB method, the mean number of applications of the rejection algorithm before an absorbing barrier is reached. For the PIT method, no steps have to be taken in the simulation.

With respect to the accuracy of the simulation, again two dependent variables are examined. Both response probability and the conditional response time density are simulated, and hence for both aspects the accuracy of approximation is evaluated. First, we examine the proportion of times that the process hits the upper boundary. Second, for the processes that hit the upper boundary, a measure of maximal deviation between the observed and theoretical conditional cumulative distribution function is computed:

$$H = \sup_{0<t<\infty} |\hat{F}_1(t) - F_1(t)|, \quad (23)$$

where $\hat{F}_1(t)$ is the empirical distribution function for the simulated values from the diffusion process, and $F_1(t)$ is the

theoretical conditional distribution function for the diffusion process (Equation 21). It is expected that H is smallest for the rejection-based algorithm, and it will be investigated how close the other methods approximate this value.

If Equation 23 is multiplied by the square root of the number of observations $\sqrt{n}$), it becomes a Kolmogorov–Smirnov test statistic (Mood et al., 1974). Under the condition that $\hat{F}_1(t)$ converges to $F_1(t)$ if the sample size goes to infinity, the distribution of the test statistic $D_n = \sqrt{n}H$ known. Given that n is large enough, D_n can be used to test the hypothesis that a sample of simulated values with empirical distribution function $\hat{F}_1(t)$ does come from the target distribution $F(t)$. Under the null hypothesis that this is true, the proportion of rejections of the null hypothesis should not exceed 5% (1%) if the test is performed at the nominal significance level of $\alpha = .05$ ($\alpha = .01$).

In this paper, emphasis is put on the raw value H instead of on the Kolmogorov–Smirnov statistic, because the outcome of the Kolmogorov–Smirnov test depends on the sample size. Any method that is not exact will be rejected as inaccurate if the number of simulations is large enough, and that is not needed here. Therefore, the Kolmogorov–Smirnov test will not be considered any more, and only the raw values H are used.

Results for the Speed of Simulation

In Table 1, one can find the mean time (in seconds) the algorithms needed to draw 3,000 realizations from the diffusion process. The RW method with the lowest precision is the fastest, followed by the RB method. For equal precision levels, the SDE method is the slowest. The results suggest that for the RW and SDE methods, the simulation time is $O(\tau^{-1})$. The simulation time for the PIT method is comparable to that of the RW method. However, if the number of simulations was larger than 3,000, the PIT method would perform faster than the RW method because, beyond the fixed cost of computing the table with function values, it takes no time to select values from it.

Table 2 shows the mean number of steps for the RW and SDE methods, again as a function of the precision level τ.

Table 1
Mean Time (in Seconds) to Draw 3,000 Realizations from the Diffusion Process for the Four Simulation Methods

Method	Precision Level	Mean Time
Random walk	0.01	0.0252
	0.001	0.2669
	0.0001	2.5717
Rejection-based		0.1663
Stochastic differential	0.01	0.1688
equation	0.001	1.3905
	0.0001	12.9905
Probability integrated	500	0.0264
transform	5,000	0.2256
	50,000	2.3164

Note—For the RW and SDE method, precision level is τ. For the PIT method, precision level is N, the number of grid values used to approximate $F_1(t)$.

For the RB method, the mean number of times that the rejection algorithm has to be applied is shown in the same table. The SDE method takes more time, because it takes more steps than the RW method to reach the boundary and it requires more complex calculations (one has to draw normally distributed numbers instead of uniformly distributed ones as in the RW method). The PIT method is not mentioned in this table, since the selection of a value is done in one step.

The third column of Table 2 contains the expected number of steps before absorption. (The expected number of steps is computed for each combination of drift rate, a, and z, and then the mean is taken.) The expected number of times the rejection algorithm has to be applied for the RB method is not available, since there is no formula to compute it. For the RW method, the observed and theoretical results are very similar, but for the SDE method, the observed and theoretical results differ, especially for the lowest precision. The reason is that Equation 20 is only an approximating formula and not the exact one. This discrepancy between the observed and expected number of steps was also noted by Luce (1986) and Smith (1990).

The effects of the different parameter values (μ, a, and z) on the speed of simulation are not shown in detail here; we will only briefly mention some interesting findings. The mean simulation time for the RW and SDE methods decreases if the (absolute) value of the drift rate becomes larger. Moreover, if the boundary separation a is large, it will take longer to simulate the process. For the RB method, a drift rate closer to zero leads to a shorter simulation time. This is expected, because the acceptance rate of the rejection algorithm attains its maximum at $\mu = 0$. Also for the RB method, the simulation time is larger the more asymmetric the diffusion process is. If the starting point is close to the lower boundary and the drift rate large, then more applications of the rejection algorithm are needed than when the starting point lies symmetrically. In the latter case, by definition, only one application of the rejection algorithm is necessary. For the PIT method, a larger boundary separation a and a smaller drift lead to a larger simulation time.

Results for the Accuracy of the Simulation

First, we checked whether the simulation methods yield good estimates of the probability of absorption at the upper boundary. To do so, the sum of the absolute deviancies between the simulated proportions and theoretical probabilities (Equation 3) is calculated. Because the RB and PIT method do not use approximations at this stage, there cannot be accuracy problems. However, this does not hold for the two other methods (RW and SDE). Only the results for the highest precision levels of the RW and SDE methods are computed. The approximation for the RW method is as good as for the RB and PIT method. In particular, the sums of absolute deviancies are 0.0368 for the RW method, 0.0279 for the RB method, and 0.0352 for the PIT method. The SDE method produces inaccurate results, especially

Table 2
Mean Number of Steps Needed to Reach a Boundary for the Random Walk (RW) Method, the Rejection-Based (RB) Method, and the Stochastic Differential Equation (SDE) Method

Method	Precision Level	Observed Mean Number		Expected Mean Number
		M	*SD*	
Random walk	0.01	26.655	14.406	26.738
	0.001	274.469	148.547	265.181
	0.0001	2,651.922	1,440.373	2,650.635
Rejection-based		1.423*	0.449	
Stochastic differential equation	0.01	35.727	18.245	26.505
	0.001	293.127	154.795	265.051
	0.0001	2,739.748	1,470.129	2,650.505

Note—For the RW and SDE method, precision level is τ. For the PIT method, precision level is N, the number of grid values used to approximate $F_1(t)$. *This mean number of steps is the mean number of times a new rejection algorithm has to be initiated. As such, its meaning differs from the other mean number of steps for the RW and SDE methods.

for the asymmetric cases. The sum of the absolute deviancies for the SDE method was 0.1862, which is clearly larger than for the other two. For lower precision, the results are even worse and therefore they are not presented.

Next, it is checked whether the obtained sample of absorption times at the upper boundary can be considered as a sample from the conditional density in Equation 4. We consider only the conditional density $f_{T|Y}(t \mid Y = 1)$, but the conclusions also hold for $f_{T|Y}(t \mid Y = 0)$.

Table 3 contains the mean deviance value H as defined in Equation 21 for the four simulation methods. For the RW, SDE and PIT methods, we distinguish between the three levels of precision. The results in Table 3 show that a precision lower than 0.0001 does not yield accurate simulations for both the RW and SDE methods. If the precision is 0.0001, the random walk gives results similar to those for the RB method. But even for the smallest time intervals that we studied, the SDE method is not as accurate as the other two methods. For the PIT method, setting up a table of 5,000 grid points already gives an H value that is undistinguishable from the one for the RB method (taking into account the standard deviation). A fortiori, the same holds of course for 50,000 grid points. The accuracy of the PIT method would even be improved if linear interpolation was used.

The lack of accuracy of the SDE method is startling at first sight. To make sure that the method works but needs an excessively large precision, the simulations are repeated with $\tau = .000001$. With this precision, drawing a sample of 3,000 realizations takes more than 20 min. The sum of absolute deviations between the probability of upper boundary absorption and the actual proportion of absorptions drops to .0335, a value close to the reported ones for the three other methods. Moreover, the mean value of H is equal to 0.0198 ($SD = .0078$), and this is also closer to the values of the three other methods (see Table 3).

Table 3
Mean Values of the Deviance Measure H for the Four Simulation Methods

Method	Precision	Mean H	*SD*
Random walk	0.01	0.0759	0.0381
	0.001	0.0320	0.0118
	0.0001	0.0191	0.0072
Rejection-based		0.0190	0.0070
Stochastic differential equation	0.01	0.1926	0.0547
	0.001	0.0665	0.0213
	0.0001	0.0301	0.0115
Probability integrated transform	500	0.0491	0.0210
	5,000	0.0205	0.0070
	50,000	0.0189	0.0075

Finally, we examine the differences in H for different conditions of the design but without showing the results in detail. Because interesting results are only to be expected for the RW, SDE, and PIT methods, those are the only ones that are discussed.[1] First, the systematic effects of the independent variables for the RW simulations with the highest precision ($\tau = .0001$) are studied. In conditions with zero drift rate, the accuracy is always worse than in conditions with nonzero drift rate. Moreover, conditions with symmetric starting points return more accurate results than do conditions with asymmetric starting points.

Second, for the SDE approximation, conditions with a small boundary separation lead to less accurate simulations. This is understandable, because if the boundaries are close together, it will be more likely that an excess over a boundary occurs, despite the correction method. Completely in line with this fact, from the two conditions with a small boundary separation (0.08), the condition with starting point at 0.02 performs worst. It appears that the correction method does not succeed in completely removing the inaccuracy due to an excess over a boundary. The problem of excess over a boundary is the reason why the SDE approximation is the least accurate method.

Third, for the PIT method, we looked at the effects for a table of 5,000 grid values. The most important trend is that simulations with zero drift rate are the least accurate ones. There are no systematic effects from starting point and boundary separation.

GENERALITY OF THE METHODS

The random walk approximation can be generalized to simulate other types of diffusion processes where the drift

rate and variance parameter are not just constant but depend on the position the process occupies. Being functions of the state of the process, these parameters are now denoted by $\mu(x)$ and $\sigma(x)$. In this more general case, not the limit of a simple random walk is taken but the limit of a birth–death process, which is the discrete analogue of these more general kinds of diffusion processes. However, the structure of the basic part of the algorithm (Equation 6) remains unchanged. Details about this generalization are given by Bhattacharya and Waymire (1990), Busemeyer and Townsend (1992), and Diederich (1997). The method can also be applied to diffusion processes with reflecting boundaries. Diffusion models with reflecting boundaries are used in cognitive research by, for instance, Schwarz (1993). The random walk is also easily generalized to nonstationary models where at a fixed time point t, the drift rate changes from one constant value to another while the variance remains constant (Ratcliff, 1980). Such models are, for instance, applied in Ratcliff and Rouder (2000).

The rejection-based method is not applicable to more general diffusion processes such as the Ornstein–Uhlenbeck process, since in that case there is no explicit formula for the conditional first-passage time densities. For other diffusion processes for which these explicit formulas are available, the inequality in Equation 11 may not be true, and so it is difficult to generalize the method. The method can be applied, however, to diffusion processes with constant drift rate and variance and reflecting boundaries. The rejection-based method is not applicable to models with a change in drift rate from one constant to another at a certain time point.

The same technique of approximating the SDE is also applicable if μ and σ depend on the state that the process is in. For that case, more advanced methods than the one described in Equation 18 are available. Fahrmeir (1976) considered one-step methods for solving ordinary differential equations (ODEs), and Equation 18 is the so called stochastic analogue of the Euler method for solving ODEs (see, e.g., Burden & Faires, 1997). More advanced one-step methods such as Heun, modified Euler, and Runge–Kutta are more appropriate for variable drift rate and variance diffusion processes (for the simple diffusion process considered in this paper, these methods give the same expression as Equation 18). Fahrmeir (1976) discusses theoretical convergence results of the numerical ODE solvers and proposes some modifications, but only for the case without absorbing boundaries; and for that case, some difficulties may rise as aforementioned. The method is easily generalized to the case with reflecting boundaries and changes in drift rate from one constant value to another at a certain time point.

It is difficult to apply the probability integral transform method for the Ornstein–Uhlenbeck process, since no closed-form expression is available for the cumulative distribution function. Changes in drift rate as those introduced by Ratcliff (1980) and applied by Ratcliff and Rouder (2000) are, however, easy to handle (from time 0 up to t, one takes the cumulative distribution function of the first process, and from t on, the cumulative distribution function from the second process). Moreover, random drift rate and random starting point models (Ratcliff, 1978; Ratcliff et al., 1999) can be simulated by inverting the cumulative distribution function that is obtained by first integrating over both drift rate and starting point distributions (this may, however, have an effect on the accuracy of the method). The probability integral transform method could be used in a hybrid construction with some of the other methods. For instance, a limited number of simulations from the random walk method could be used to set up the empirical cumulative distribution function of some process, and then, starting from the empirical cumulative distribution function, the probability integral transform method is used to generate the remaining simulations. This technique is particularly useful for simulating diffusion processes for which there is no closed-form expression for the cumulative distribution function available (e.g., the Ornstein–Uhlenbeck process). It combines the advantages of two methods: the generality of the random walk method, and the speed and accuracy of the probability integral transform method.

DISCUSSION

In this paper, four methods for the simulation of a Wiener process are proposed and compared with each other with respect to their simulation speed and the accuracy of the results. With respect to these two criteria, it turns out that the two best methods are the rejection-based method and the probability integral transform method. The advantage of the rejection-based method is that it is exact. However, if a huge sample from the same distribution has to be drawn, the probability integral transform method is to be preferred, since its accuracy is comparable (certainly for the choice probabilities but also for the response times) to that of the rejection-based method and it will be at least as fast or even faster.

If the time intervals are chosen to be small enough (high precision), also the random walk gives accurate results. The third method, based on the approximation of a stochastic differential equation, is the slowest and does not yield accurate results. The fourth method, using the probability integral transform method, performs almost as fast and as accurately as the rejection-based method.

The poor performance of the stochastic differential equation method is not very surprising, since that method suffers from the boundary excess problem and this is not completely corrected for by widening the boundaries. Perhaps the problem can be rectified by using an algorithm with adaptive precision: lower precision in the neighborhood of the starting point, but higher precision when the process approaches one of the boundaries.

Although the rejection-based method is an exact and fast method, a drawback of the rejection-based method is that it is restricted in its application: Only diffusion processes with constant drift rate and variance can be simulated. With only minor adaptations, the random walk method,

stochastic differential equation method, and probability integral transform method can handle more general diffusion processes.

The algorithms are compared with respect to their speed of simulation, but they could also be compared with respect of their "speed of implementing." The gain in time by speeding up the simulation may be lost in some situations by implementing the method. Of course, the speed of implementing depends strongly on the expertise of the researcher, but a few general points are worth mentioning. If the stochastic differential equation method is left aside because of its inaccuracy, the random walk is certainly the easiest to implement—after that, the probability integral transform method, and finally the rejection-based method. Hence, if only a few simulations are needed for a quick inspection of the properties of a model, the random walk may be the most preferable method.

REFERENCES

Bhattacharya, R. N., & Waymire, E. C. (1990). *Stochastic processes with applications.* New York: Wiley.

Bouleau, N., & Lépingle, D. (1994). *Numerical methods for stochastic processes.* New York: Wiley.

Burden, R. L., & Faires, J. D. (1997). *Numerical analysis* (6th ed.). Pacific Grove, CA: Brooks/Cole.

Busemeyer, J. R., & Townsend, J. T. (1992). Fundamental derivations from decision field theory. *Mathematical Social Sciences*, **23**, 255-282.

Cox, D. R., & Miller, H. D. (1970). *The theory of stochastic processes.* London: Methuen.

Diederich, A. (1997). Dynamic stochastic models for decision making under time constraints. *Journal of Mathematical Psychology*, **41**, 260-274.

Fahrmeir, L. (1976). On the simulation of stochastic processes with continuous state and parameter space. In L. Dekker (Ed.), *Simulation of systems* (pp. 67-71). Amsterdam: North-Holland.

Feller, W. (1968). *An introduction to probability theory and its applications* (Vol. 1, 3rd ed.). New York: Wiley.

Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, **274**, 427-430.

Heath, R. A., & Kelly, L. (1988). An application of a discriminability index for the assessment of individual differences. In R. A. Heath (Ed.), *Current issues in cognitive development and mathematical psychology* (pp. 180-196). Newcastle, NSW: University of Newcastle.

Karlin, S., & Taylor, H. M. (1981). *A second course in stochastic processes.* New York: Academic Press.

Lichters, R., Fricke, T., & Schnakenberg, J. (1995). *Stochastic simulation of diffusion with absorbing and reflecting boundary conditions* [On-line]. Available: http://www.physik.rwth-aachen.de/group/thphys/tpd/Werke/absrefl.html

Luce, R. D. (1986). *Response times.* New York: Oxford University Press.

Mood, A. M., Graybill, F. A., & Boes, D. C. (1974). *Introduction to the theory of statistics.* New York: McGraw-Hill.

Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1986). *Numerical recipes: The art of scientific computing.* New York: Cambridge University Press.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, **85**, 59-108.

Ratcliff, R. (1980). A note on modeling accumulation of information when the rate of accumulation changes over time. *Journal of Mathematical Psychology*, **21**, 178-184.

Ratcliff, R., & Rouder, J. (1998). Modeling response times for two-choice decisions. *Psychological Science*, **9**, 347-356.

Ratcliff, R., & Rouder, J. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human Perception & Performance*, **26**, 127-140.

Ratcliff, R., & Tuerlinckx, F. (2001). *Estimating the parameters of the diffusion model.* Manuscript submitted for publication.

Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, **106**, 261-300.

Ross, S. M. (1996). *Stochastic processes* (2nd ed.). New York: Wiley.

Schwarz, W. A. (1993). A diffusion model of early visual search: Theoretical analysis and experimental results. *Psychological Research*, **55**, 200-207.

Smith, P. L. (1990). A note on the distribution of response times for a random walk with Gaussian increments. *Journal of Mathematical Psychology*, **34**, 445-459.

Smith, P. L. (2000). Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of Mathematical Psychology*, **44**, 408-463.

Tanner, M. A. (1996). *Tools for statistical inference: Methods for the exploration of posterior distributions and likelihood functions* (2nd ed.). New York: Springer-Verlag.

Wald, A. (1947). *Sequential analysis.* New York: Wiley.

NOTE

1. We expect to find no systematic effects of drift rate and starting point on the accuracy of the simulation for the RB method. Nevertheless, some systematic effects appear, and these are attributed to the fact that the number of cases on which the deviance measure H is computed differs.

APPENDIX A

To prove inequality 11, first rewrite $f_{T|Y}(t\,|\,Y=1)$:

$$f_{T|Y}(t\,|\,Y=1)=\frac{\frac{\pi\sigma^2}{a^2}\left[\exp\left(\frac{a\mu}{2\sigma^2}\right)+\exp\left(-\frac{a\mu}{2\sigma^2}\right)\right]}{\left(\frac{\mu^2}{2\sigma^2}+\frac{\pi^2\sigma^2}{2a^2}\right)}\left(\frac{\mu^2}{2\sigma^2}+\frac{\pi^2\sigma^2}{2a^2}\right)$$

$$\times\sum_{n=0}^{\infty}(2n+1)(-1)^n\exp\left\{-\frac{1}{2}\left(\frac{\mu^2}{\sigma^2}+\frac{\pi^2\sigma^2(2n+1)^2}{a^2}\right)t\right\}$$

$$=M\lambda\sum_{n=0}^{\infty}(2n+1)(-1)^n\exp\left\{-\frac{1}{2}\left(\frac{\mu^2}{\sigma^2}+\frac{\pi^2\sigma^2(2n+1)^2}{a^2}\right)t\right\},$$

where M and λ are defined in Equations 12 and 13, respectively.

The even terms of the infinite sum ($n = 0, 2, \ldots$) are positive, and the odd terms ($n = 1, 3, \ldots$) are negative. Also, the term for $n + 1$ is always smaller in absolute value than the term for n. It follows that the difference be-

APPENDIX A (Continued)

tween consecutive terms starting from $n = 1$ is negative. Therefore, the total of the infinite sum starting from $n = 1$ is also negative. This leads to the following inequality:

$$f_{T|Y}(t \mid Y = 1) \leq \frac{\frac{\pi\sigma^2}{a^2}\left[\exp\left(\frac{a\mu}{2\sigma^2}\right) + \exp\left(-\frac{a\mu}{2\sigma^2}\right)\right]}{\left(\frac{\mu^2}{2\sigma^2} + \frac{\pi^2\sigma^2}{2a^2}\right)}\left(\frac{\mu^2}{2\sigma^2} + \frac{\pi^2\sigma^2}{2a^2}\right)$$
$$\times \exp\left[-\frac{1}{2}\left(\frac{\mu^2}{\sigma^2} + \frac{\pi^2\sigma^2}{a^2}\right)t\right]$$
$$= M\lambda \exp(-\lambda t).$$

APPENDIX B

Suppose one wants to sample a realization from the density function $g(t)$, but no direct algorithm (such as there exists for sampling from, e.g., a normal distribution) exists for this purpose. Now suppose that another density function $h(t)$ can be found, from which it is easy to sample (e.g., the exponential or normal density), and there also exists a constant M such that the following inequality holds:

$$g(t) \leq Mh(t), \tag{B1}$$

which means that $g(t)$ is dominated by the function $Mh(t)$ for all t. This dominating function is called the majorization or envelope function. The algorithm then proceeds as follows:

1. Sample a t^* from $h(t)$. This t^* is called a candidate value.

2. Sample u from $U(0,1)$, independently from t^*.

3. If

$$u \leq \frac{g(t^*)}{Mh(t^*)},$$

then accept the draw t^*; otherwise go back to step 1.

It is shown in Tanner (1996) that a t^* that is accepted in step 3 is a realization from $g(t)$. Moreover, it can be shown that 1/M is the probability that a candidate t^* will be accepted as a draw from $g(t)$, which is also called the acceptance rate. The acceptance rate determines the efficiency of the simulation method: If the acceptance rate is low, many candidate values need to be generated before one is accepted as coming from $g(t)$.

APPENDIX C

If $U = G_X(X)$ and X has density $g_X(x)$, then $\Pr(U \leq u) = \Pr(G_X(X) \leq u) = \Pr(X \leq G_X^{-1}(u)) = G_X(G_X^{-1}(u)) = u$. Hence, U is uniformly distributed. The reverse is as follows: If $X = G_X^{-1}(U)$ and U is uniformly distributed over $[0,1]$, then $\Pr(X \leq x) = \Pr(G_X^{-1}(U) \leq x) = \Pr(U \leq \Pr(G_X(x)) = G_X(x)$.

(Manuscript received June 2, 2000;
revision accepted for publication January 18, 2001.)