

**FITTING EFFECTIVE DIFFUSION MODELS TO DATA
ASSOCIATED WITH A “GLASSY POTENTIAL”: *ESTIMATION,
CLASSICAL INFERENCE PROCEDURES AND SOME HEURISTICS***

*

CHRISTOPHER P. CALDERON [†]

Abstract. A variety of researchers [15, 22, 27, 29, 31, 33] have successfully obtained the parameters of low-dimensional diffusion models using the data that comes out of atomistic simulations. This naturally raises a variety of questions about efficient estimation, goodness-of-fit tests, and confidence interval estimation. The first part of this article uses maximum likelihood estimation (MLE) to obtain the parameters of a diffusion model from a scalar time series. I address numerical issues associated with attempting to realize asymptotic statistics results with moderate sample sizes in the presence of exact and approximated transition densities. Approximate transition densities are used because the analytic solution of a transition density associated with a parametric diffusion model is often unknown. I am primarily interested in how well the deterministic transition density expansions of Ait-Sahalia capture the curvature of the transition density in (idealized) situations that occur when one carries out simulations in the presence of a “glassy” interaction potential. Accurate approximation of the curvature of the transition density is desirable because it can be used to quantify the goodness-of-fit of the model and to calculate asymptotic confidence intervals of the estimated parameters. The second part of this paper contributes a heuristic estimation technique for approximating a nonlinear diffusion model. A “global” nonlinear model is obtained by taking a batch of time series and applying simple local models to portions of the data. I demonstrate the technique on a diffusion model with a known transition density and on data generated by the Stochastic Simulation Algorithm [21].

Key words. Effective diffusion model, multiscale approximation, log likelihood ratio expansion, piecewise polynomial SDE, quasi-maximum likelihood, stochastic process approximation

AMS subject classifications. 62M10,82-08,82B31,82C31,82C43

1. Introduction. Complicated systems are often approximated by overly simplified models. A significant research effort has gone into attempting to efficiently summarize the information contained in a complicated atomistic simulation with a low-dimensional effective model [15, 22, 27, 29, 31, 33]. Atomistic simulations contain many observables, an effective diffusion model aims at representing the salient features of the data in the drift component of a stochastic differential equation (SDE) and lumping the effects of the neglected details into the noise term. The appeal of effective models stems from the fact that information contained in the effective models can easily and quickly be extracted by analytical methods or well-established numerical procedures. The idea being that the “truth” is contained in the atomistic simulation, but the computational load required to get the information is so large that the researcher has a difficult time exploring all of the process parameters under study.

Before one attempts to wrap effective models around the output of an atomistic simulation, a variety of assumptions need to be made about the data and the parametric model. In this paper, I obtain parameter estimates from the classical parametric framework via maximum likelihood estimation (MLE). Here the phrase “classical parametric framework” refers to the fact that one uses a single family of functions (of specified functional form) for the drift and diffusion coefficient functions

*This work was partially supported by a Ford Foundation/NRC Fellowship

[†]Department of Chemical Engineering, Princeton University, Princeton, New Jersey 08544 (ccaldero@princeton.edu).

of the diffusion SDE that depends only on a finite number of parameters. I make the following assumptions about the atomistic system and the parametric diffusion model:

- A small set of order parameters [46, 25] that accurately summarize the full atomistic system are identified and easily measurable. The term “order parameter” is used to refer to the observables modeled in the effective diffusion SDE.
- The parametric model is uniquely identifiable.
- The *exact* transition density of the parametric model has all of the regularity properties that make MLE attractive. Namely it is the *asymptotically* most efficient estimator in the sense that the properly normalized parameter distribution associated with the procedure converges to a normal distribution with the smallest asymptotic variance (in “nonergodic” situations this assumption is relaxed).
- The true parameter value admits a contiguous neighborhood and the process allows for a quadratic expansion of the log likelihood ratio ¹.
- The dynamics of the order parameters can be adequately described by a diffusion SDE.
- The drift and diffusion coefficient of the effective SDE are suitably smooth (to be more specific assume the functions are infinitely differentiable with respect to the parameters and state, but this can be relaxed substantially) and the drift component of the SDE comes from the gradient of an effective potential [15, 31]. Even if the order parameter being considered is governed by a glassy potential (this term is described in section 2), a smooth drift coefficient function can be used to summarize key features of the free energy landscape ².

The first item is extremely important and is an active area of research in our group [46], but will not be addressed in this article. The next two assumptions allow one to trust the parameter estimates and allows one hope for checking if some asymptotic results associated with MLE [28, 48, 50] hold for the sample sizes used. The fifth item is briefly addressed in the final application of the second part of this paper where I estimate the parameters of a diffusion approximation of a jump process (the model of the reduction of nitric oxide on a platinum surface given in [41] is revisited). The final issue is concerned with “model misspecification” [50], dealing with this issue is important if the effective model is to be of any practical use. I first obtain parameter estimates assuming the final item holds, but in section 4.3 classical techniques for testing this assumption *a posteriori* are outlined. Some nonparametric goodness-of-fit tests are better suited for practical implementation [12, 26], but classical tests are of interest because the tests used quantify how well a transition density approximation matches the curvature of the true density.

In order to use MLE one must have an approximation of the transition density of the diffusion process. Unfortunately, for many SDE’s the transition density is not available in closed-form so one must resort to some approximation of the density. The estimation techniques presented are intended for application to interesting atomistic systems (e.g. a time series that comes from a molecular dynamics simulation), but in

¹These terms are briefly introduced in section 4, consult van der Vaart [48] chapters 5-8 for a clear detailed treatment

²For example the global minimum value of the smooth effective free energy surface is the same as that of the more complicated glassy surface

this paper toy models are studied in order to systematically study the consequences of using the transition density expansions of Ait-Sahalia [3, 4] in some large sample statistics applications. Our research group has had success in applying some of these ideas to actual atomistic systems, but this paper’s main concern is in determining exactly how much information can be extracted from the transition density expansions in controlled examples intended to mimic scenarios encountered in some multiscale applications.

The remainder of this paper is organized as follows: Section 2 reviews the multiscale applications that motivated this study. Section 3 lays out the model systems used. Section 4 outlines the techniques and estimation tools used in the presence of “contaminated”³ data and transition densities. Section 5 outlines a simple estimation procedure that can be used to study multiscale systems that satisfy the assumptions stated above. Section 6 contains the numerical results and discussion and the final section gives the conclusion and outlook.

2. Motivation. A situation encountered often in spin glasses [34], protein folding [25], and zeolites [10] is that of a harmonic “glassy” potential energy surface⁴. That is to say, when one looks at a free-energy surface from a distance, the shape of the free energy surface is roughly parabolic. When one looks closely at the details of the surface however, one sees many bumps in the surface (see figure 3.1). In many applications, it is believed that the order parameter of the process “funnels” its way down to the global minimum of the free energy surface [25] in the long time limit. Current computational power does not always allow one to carry out an atomistic simulation long enough to observe such a phenomena because the order parameter can get trapped in a local minima. Sometimes one can reach the global free energy minima by increasing the temperature parameter of the simulation [18, 13]. When this occurs, the force binding the order parameter to the global minimum still has a “bumpy” potential associated with it (but the magnitude of the bumps is smaller because of the new temperature scale). I refer to this hypothetical case as situation I.

Another commonly encountered scenario is one where the order parameter is trapped in a free energy well which is not the global minimum of the surface. Many applications require system information at low temperatures, ruling out the simple technique mentioned in the previous paragraph. If the temperature of the system is so low that on the timescale of the atomistic simulation that the order parameter appears to be approximately bound by a smooth (but not necessarily harmonic) potential in a neighborhood of the local minima, then I will refer to this case as situation II.

Classic statistical mechanics models usually assume that the noise around the local minima is state independent. This assumption does not hold in a variety of interesting systems, [31, 15] so in all models I consider there is state dependent noise in the process. The estimation techniques presented in this paper deal with both situations described in the preceding paragraphs.

³This term is used to convey the fact that one knows that the true data does not follow the exact proposed parametric model

⁴This estimation strategy was developed in order to accurately measure the curvature of complicated free energy surfaces associated with atomistic simulations. If the location(s) of the dominant free energy wells are known by theory or simulation methods, then the estimation methods shown here can be used to measure the curvature at the well minima which can in turn be useful for getting information about transition pathways [14]. If one also has knowledge of where the saddles are and a protocol for starting meaningful simulations around the saddle point then the methods presented can also be used to determine the curvature of the unstable state points.

3. Model Systems. To idealize situation I and II, data is generated from two families of SDE's. This first family is meant to mimic situation I and has the form:

$$(3.1) \quad dX_t = \left(\kappa(\alpha - X_t) + \beta \sin((X_t - \alpha)\omega 2\pi) \right) dt + \sigma \sqrt{X_t} dW_t$$

The second family idealizes situation II and takes the form:

$$(3.2) \quad dX_t = \left(\kappa(\alpha - X_t) + 4\gamma(\alpha - X_t)^3 \right) dt + \sigma \sqrt{X_t} dW_t$$

The parameters are set to $\alpha = 20$, $\kappa = \sigma = 4$, and $\omega = \frac{1}{3}$ throughout. The parameters β and γ take the values $(0, 15, 60, 200)$ and $(0, \frac{1}{400})$. I refer to these cases as situations I A-D and situations II A-B respectively. The situation where γ and β are both zero is known as the Cox-Ingersoll-Ross (CIR) model and is one of the rare situations where an SDE with nonlinear coefficients has an explicit solution [3]. I use this example because it illustrates mean reversion, it demonstrates state dependent noise and most importantly has an exact closed-form transition density which can be used to help determine why a transition density approximation is failing.

Figures 3.1 and 3.2 plot the potential energy surfaces for the cases studied. In all cases, sample paths of the above processes are simulated using the explicit Euler scheme [30] with a step size $\Delta t = 2^{-9}$ given data starting from the invariant density associated with situation I A. The data is observed every 16^{th} step yielding a constant observation window spaced by $\delta t_{obs} = 2^{-5}$ time units (in Situation I A-C the data is also sampled every 64^{th} step giving $\delta t_{obs} = 2^{-3}$). Each plot and graph that are grouped together in this paper used the same Brownian trajectories in order to simulate paths (the only difference between the sample paths is caused by the different drift coefficients) in order to reduce variation due to random number draws.

For the first part of this paper, the data generated by the SDE's above is modeled by the CIR class:

$$(3.3) \quad dX_t = \kappa(\alpha - X_t)dt + \sigma \sqrt{X_t}dW_t$$

For the second part of this paper, the following parametric family is used:

$$(3.4) \quad dX_t = \left(a + b(X_t - X_o) \right) dt + \left(c + d(X_t - X_o) \right) dW_t$$

Where X_o is user specified; the parameter vectors are estimated by techniques associated with maximum likelihood in all cases.

The second terms in the drift coefficient of the data generating processes are used to determine how robust the estimator is against model misspecification [50, 42, 9]. The perturbation terms are not modeled because it assumed that their true (or approximate) functional form are completely unknown and the interest is primarily in the smooth noise and mean reversion parameters (α, κ, σ) . Of course the presence of these extra terms affect the estimation of the parameters, it is shown in section 6 that the effects induced by these perturbation parameters affect things consistently with what a physicist or chemist would intuitively anticipate. The interesting feature demonstrated in the aforementioned section is quantitatively how the MLE procedure carried out with various transition density approximations respond to these perturbation parameters in relation to the exact transition density.

The reason for studying the first two idealized model systems stems from a desire to *carefully* numerically quantify how the MLE procedure with approximated

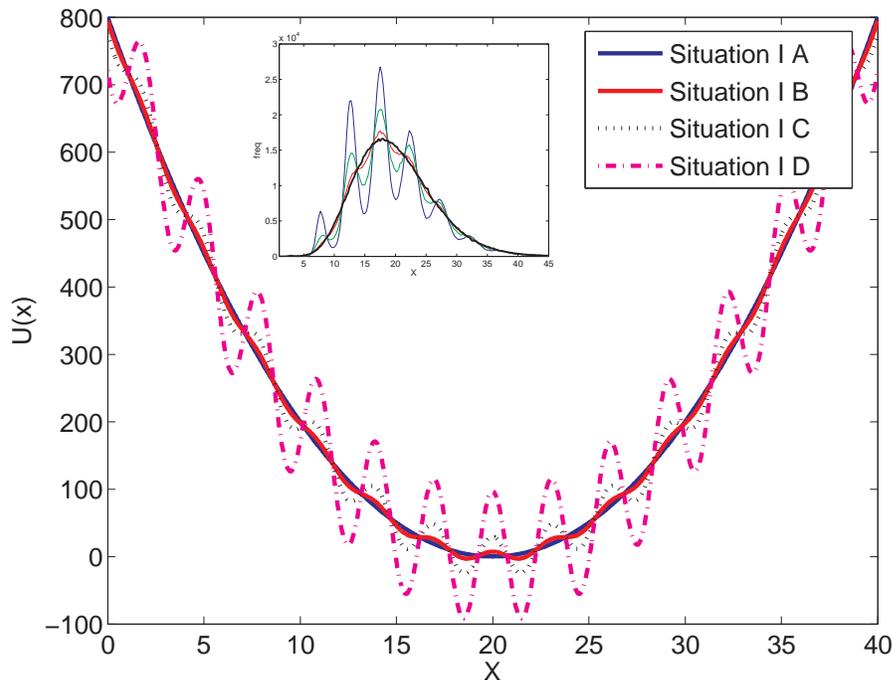
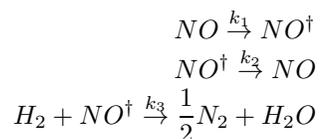


FIG. 3.1. *Situation I Potential* Potential energy function used to determine drift with $\beta = (0, 15, 60, 200)$. The inset shows the empirically measured invariant distribution for the four values of β used (with obvious correspondence between the four cases); the distributions are shown only to give one an idea of how the different parameter values affect the long-term dynamics (MLE parameters only depend on the observation frequency).

transition density performs in tasks beyond point parameter estimation. The results obtained are of course specific to the model and parameter values used, however one can always obtain the parameters of a parametric diffusion model using observations from the particular system being studied and then carry out an idealized set of tests similar to the ones presented here by using established SDE path simulation techniques [30].

The final model presented is one for the reduction of nitric oxide (NO) by hydrogen gas on a platinum (Pt) surface. The mechanism is as follows [47]:



(3.5)

NO^\dagger represents NO absorbed onto the Pt surface; this mechanism is used with Gillespie's Stochastic Simulation Algorithm (SSA) [21] technique in order to construct stochastic evolution rules for the amount of NO in the system at any given time. This model is used because it exhibits nonlinear mean reversion with state dependent noise. The jump process is known to converge weakly to a diffusion with a cubic drift term

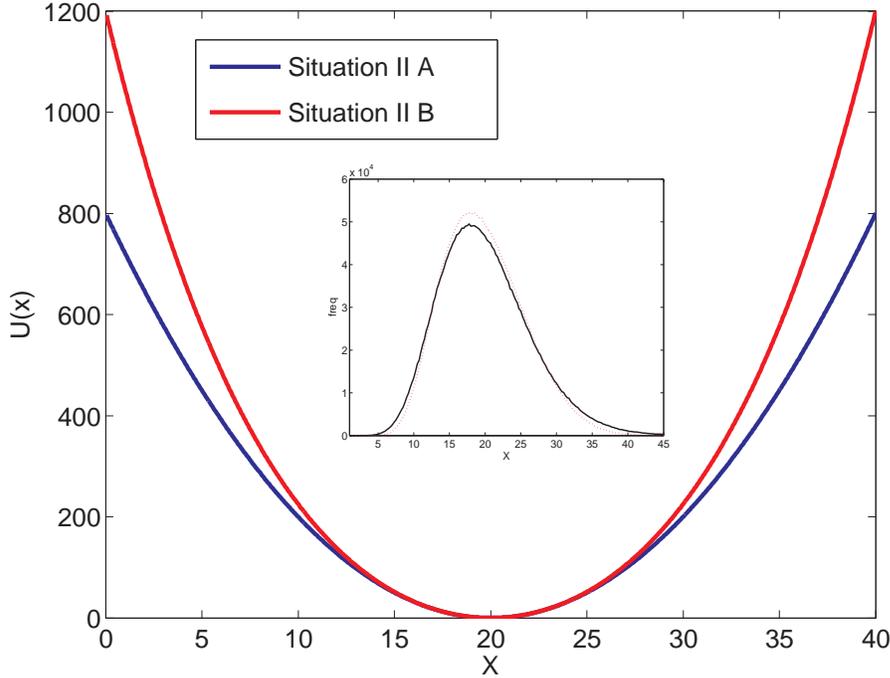


FIG. 3.2. *Situation II Potential* Potential energy function used to determine drift with ($\gamma = 0, \frac{1}{400}$) along with inset of corresponding invariants distributions (the narrower distribution corresponds to the nonlinear perturbation case)

as the system size parameter (denoted by $N_{molecules}$) increases [21, 1]. This yields another “situation II” type scenario, but now the data generating process is not a genuine diffusion⁵. The parameter k_3 is set to a numerical value of 4, the other model parameters used are given in [41]. In the final part of this paper, the parameters of the assumed model are extracted in the “small molecule” case ($N_{molecules} = 3600$). This value is chosen because simple visual inspection indicates that the process has not yet converged to a diffusion ($\delta t_{obs} = 2^{-5}$) with this molecular population size, so the exact functional form of the diffusion model is unknown. A “global”⁶ estimate of the diffusion approximation of the actual process is obtained using techniques presented in the second part of this paper and the invariant density of the obtained nonlinear diffusion model is compared to that of the actual SSA process in section 6.

⁵It should be pointed out that some approximations of the above process by a diffusion model [21] use as many Brownian driving terms as there are elementary reaction steps; our black-box approach only uses one Brownian term for each state component (in this paper the noise contribution of the individual reaction events are lumped into a single noise term).

⁶Actually one can only estimate the function for state points visited, making global a slight misnomer (hence the quotes). One is always free to extrapolate the coefficient functions and get a genuine global diffusion approximation if one somehow knows beforehand that the state points in the time series are adequately representative of the entire portion of phase space having significant probability mass in the infinite time limit.

4. Statistical Tools. In the beginning of this section, some classical statistical tools relevant to this study are outlined. The tools are only briefly defined, references are given throughout which comprehensively describe the details of the theory applied. The tools below are applied to the estimation of the parameters associated with both stationary and “nonergodic”⁷ time series.

MLE’s importance⁸ stems from the fact that one needs some kind of generic metric by which to judge a wide class of parametric models by. In situations where the underlying transition density satisfies a set of regularity assumptions [48], it is very appealing because it provides a consistent estimator with the minimum asymptotic variance. If one has a reliable estimate of the underlying transition density, then one can sometimes (in the stationary ergodic distribution case [28]) determine the asymptotic parameter distribution with a simple deterministic integral [24, 38]. One can also generate test statistics based on output of the MLE procedure which can be used to assess the goodness-of-fit of the parametric model [50].

Next, an optimal simple hypothesis test is introduced. Specifically, the transition density expansions are used in order to create the Neyman-Pearson test statistic [8]. A simple hypothesis test is useful when one wants to test the statistical significance of the magnitude of the changes in effective model parameters when one adds more features to the underlying atomistic simulation (e.g. one would like to determine if the changes in the effective model are significant when the parameters are estimated from the output of atomistic simulations that use a potential with and without electrostatic interactions).

It is already known that the simple models that are wrapped around the data do not faithfully represent the *exact* system dynamics. In section 4.3, methods that can be used to quantify how closely the proposed parametric models represent the data are reviewed. In section 4.4, a heuristic method that can be used in the nonergodic case for obtaining parameter uncertainty estimates is presented.

4.1. Maximum Likelihood Basics. In order to avoid technical complications, it is assumed throughout that the *exact* distribution associated with the parametric model admits a density whose logarithm is well defined almost everywhere and the logarithm of the density is continuously twice differentiable. The principal of maximum likelihood is based on maximizing the following integral with respect to the parameter θ :

$$(4.1) \quad \int_{\Omega} f(\mathbf{x}; \theta) \log \left(f(\mathbf{x}; \theta) \right) d\mathbb{Q}$$

In the above equation, \mathbb{Q} corresponds to the measure of the underlying probability space, $f(\mathbf{x}; \theta)$ corresponds to the Radon-Nikodym derivative (consult [32, 50]) of the law of the random variable \mathbf{x} with respect to the underlying probability space and \mathbf{x} corresponds to a discretely sampled time series (of finite length = M). Assume that if the measure of a set under \mathbb{Q} is zero it implies that the measure of the set under \mathbb{P}_{θ} is also zero (the phrase “ \mathbb{P}_{θ} is absolutely continuous with respect to \mathbb{Q} ” is used

⁷This term is intended to describe situations where the parameter distribution associated with an estimation scheme is not asymptotically normally distributed (with a deterministic covariance matrix); this can occur if the sample size is itself random or if the time series is nonstationary [6].

⁸In practice one usually deals with a quasi-maximum likelihood estimate. The difference between the two is described in section 4.3

to describe this situation). Let $\mathbb{P}_\theta := \int_\Omega f(\mathbf{x}; \theta) d\mathbb{Q}$. For discrete Markovian models, $f(\mathbf{x}; \theta)$ can readily be calculated by the following formula [24]:

$$(4.2) \quad f(\mathbf{x}; \theta) = f(\mathbf{x}_0) \prod_{n=1}^{M-1} f(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta)$$

In the above equation $f(\mathbf{x}_n | \mathbf{x}_{n-1}; \theta)$ represents the conditional probability (transition density) of observing \mathbf{x}_n given the observation \mathbf{x}_{n-1} . In practice, one usually takes a finite sample of data and presents this data to a Monte Carlo scheme that is meant to approximate the integral in equation 4.1 and finds the parameter values that yield the maximum value. In what follows, the function below is referred to as the log likelihood function (assume throughout that \mathbf{x}_0 has a Dirac distribution) :

$$(4.3) \quad \mathcal{L}_\theta := \sum_{i=1}^M \log \left(f(\mathbf{x}_i | \mathbf{x}_{i-1}; \theta) \right)$$

Under our assumptions one has ⁹ the following:

$$\sqrt{M}(\theta - \hat{\theta}) \xrightarrow{\mathbb{P}_{\hat{\theta}}} N(0, \mathcal{F}^{-1})$$

Where in the above $\hat{\theta}$ is the “true” parameter of the model; θ represents the parameter estimated with a finite time series of length M ; $\xrightarrow{\mathbb{P}_{\hat{\theta}}}$ denotes convergence in distribution [48, 24] under $\mathbb{P}_{\hat{\theta}}$; $N(0, \mathcal{F}^{-1})$ denotes a normal distribution with mean zero and covariance matrix \mathcal{F}^{-1} . For a correctly specified model, \mathcal{F} ¹⁰ can be estimated in a variety of ways [50, 38]. Various conditions can be tested to see if asymptotic results are relevant for the finite sample sizes used [48, 38]. In this article sample sizes are moderate, but good agreement with some classical asymptotic predictions [50] is observed. When a closed-form transition density is in hand one can deterministically calculate \mathcal{F} in the stationary ergodic case ¹¹.

In practice, maximum likelihood does fail spectacularly for some simple models because of singularities that can be observed with the log likelihood function ¹². The classical example is the following: one assumes that a distribution is a mixture two Gaussians whose variance $\in (0, \infty)$. Then the finite sample MLE fails to exist in this simple case (see [8, 48] for a more in depth discussion) ¹³.

⁹This actually holds under less stringent regularity assumptions [48]

¹⁰I will adhere to common convention and call this the Fisher information matrix

¹¹In the stationary ergodic case, if one has a time series $(\mathbf{x}_1, \dots, \mathbf{x}_M)$ then one can ignore the initial distribution in the “infinite M ” limit. If the state is n -dimensional, then one can approximate the Fisher information by a $2 \times n$ dimensional deterministic integral $\left(\mathcal{F} \approx \int \frac{\partial \log(f(x|x_o;\hat{\theta}))}{\partial \theta} \frac{\partial \log(f(x|x_o;\hat{\theta}))}{\partial \theta}^T f(x|x_o;\hat{\theta}) dx d\pi(x_o) \right)$, where $d\pi(x_o)$ is the invariant distribution of the process (which in the scalar case can usually be readily calculated in closed-form from the coefficients of the parametric SDE [26]). The difficulties encountered in the nonstationary time series situation is analogous to the situation of using the Metropolis algorithm to sample phase space [18]; in principle a deterministic integral could be evaluated, but with current computational power quadrature is not possible due to the high dimensionality of the problem

¹²Recall this term implies a finite sample size approximation to equation 4.1

¹³In practice, one can partially remedy this situation by finding a local minima using a variant of the technique outlined in section 4.4, but the new comer to MLE should be aware that some failures of MLE [37] are not as easy to remedy, especially when one has data that does not come from the assumed parametric model class

The aforementioned point brings us to an important observation contained in this paper. It is well known that the transition density associated with a diffusion SDE can be solved through a corresponding PDE (via the Kolmogorov equations [4]). One can often prove that the density of the process has the regularity properties needed to make *exact* MLE successful (*a priori* regularity bounds associated with the log likelihood function is trickier) via techniques of harmonic analysis [43], but many useful properties determined by analytic techniques are only applicable to the *exact* transition density. Approximations to the transition density are needed in cases where the transition density is not available in closed-form. The approximations provided by Ait-Sahalia’s Hermite expansion [3, 4] are useful for obtaining point parameter estimates, can capture the curvature of the transition density enough to approximate parameter distributions, and are sometimes accurate enough to create test statistics needed for some hypothesis tests. However, when one uses derivatives of the expansion for creating Wald or Rao ¹⁴ test statistics [8] it can introduce spurious singularities which complicates squeezing all of the information that is theoretically possible from MLE. I use the “Euler” estimator as a crude transition density approximation to demonstrate some points in this paper ¹⁵. It is well known that the Euler estimator is significantly biased [40], this fact helped to initiate a flood of transition density approximations in recent years [3, 4, 5, 7, 19, 44] (just to mention a few).

4.2. Optimal Binary Alternative Hypothesis Testing: The Neyman-Pearson Lemma. In this section it is assumed that one has a data set and two parameter vectors. Assume that one parameter vector is the “null hypothesis” and the other parameter vector is the “alternative”. The type I error probability is defined as the probability of rejecting the null hypothesis when in fact the null is true (classically denoted by α). The power of a test against the alternative is the probability of rejecting the null hypothesis when in fact the alternative is true. An optimal test statistic can be found which for a specified α maximizes the power [8]. To create the test statistic define the likelihood ratio by $L := \frac{f(\mathbf{x};\theta_{Alt})}{f(\mathbf{x};\theta_{Null})}$;

one rejects the null if $L > \beta^{NP}$ where β^{NP} is a scalar value that allows the equality $\int_{L > \beta^{NP}} f(\mathbf{x}, \theta_{Null}) dx = \alpha$

In the context of multiscale systems computations, this test is really only practical for stationary ergodic distributions (or for order parameters that are trapped for the duration of a simulation in a local free energy minima), but it provides us with insight as to how well the transition density captures the likelihood ratio of nearby parameter points. Applications shown later require a highly accurate approximation of the likelihood ratio.

4.3. Goodness-of-fit and Model Misspecification. It is possible to test if the log likelihood function is consistent with the proposed model structure by testing

¹⁴These test are concerned with using the Fisher information matrix as a normalizing matrix to create a χ^2 statistic. Both tests can be used to construct confidence ellipsoids around parameter estimates [48, 24, 8]

¹⁵This estimator is motivated by the Euler SDE simulation path technique [30]. One assumes that for a given observation pair $(\mathbf{x}_n, \mathbf{x}_{n+1})$ that a normal distribution can be used whose standard deviation is given by the diffusion coefficient evaluated at \mathbf{x}_n times the square root of the time between successive observations ($\sqrt{\delta t}$) and the conditional mean is given by applying the deterministic explicit Euler scheme to the drift coefficient given \mathbf{x}_n

the following condition [50]:

$$\mathcal{F}_{Hessian} := \frac{\partial^2 \mathcal{L}}{\partial \theta^2} = -\mathcal{F}_{OP}$$

Where the Hessian is evaluated at $\hat{\theta}$ and \mathcal{F}_{OP} is defined by:

$$\mathcal{F}_{OP} := \frac{1}{M} \sum_{i=1}^M \frac{\partial \log(f(x_i|x_{i-1}; \hat{\theta}))}{\partial \theta} \frac{\partial \log(f(x_i|x_{i-1}; \hat{\theta}))}{\partial \theta}^T$$

In the above, the superscript T denotes the transposition operation, \mathcal{F}_{OP} is the ‘‘outer product’’ matrix. For correctly specified models, both \mathcal{F}_{OP} and $-\mathcal{F}_{Hessian}$ [24] are valid estimates of \mathcal{F} . When a closed-form expansion is in hand, these quantities are easily computable after the optimal parameter is located.

For real data it is overly optimistic to expect to be able to exactly parameterize the density of the process with a Euclidean parameter vector. However, in some situations it is meaningful to attempt to project the data onto the proposed model structure [9, 50] yielding a Quasi-Maximum Likelihood Estimator (QMLE). When the true density does not lie in the proposed parametric model class, one can still maximize the integral in equation 4.1 yielding the estimator that minimizes the Kullbeck-Leibler distance [32, 42]. Under assumptions laid out in [50], this QMLE converges to a normal distribution in the infinite sample limit:

$$\sqrt{M}(\theta - \hat{\theta}) \xrightarrow{\mathbb{P}_{\hat{\theta}}} N(0, \mathcal{C})$$

The matrix $\mathcal{C} (:= \mathcal{F}_{Hessian}^{-1} \mathcal{F}_{OP} \mathcal{F}_{Hessian}^{-1})$ replaces \mathcal{F}^{-1} . This fact can be used to test the goodness-of-fit given the data and the optimal parameter vector via the Rao or Wald test statistic [50, 8]. The methodology laid out in [50] is comprehensive and powerful, but in order to develop techniques which can be accessed by the diverse audience involved in multiscale modeling, algorithms available in standard packages/libraries (MATLAB, IMSL) are employed .

4.4. Le Cam’s Method and Likelihood Ratio Expansions (LAQ). Lucien Le Cam was a major contributor to a variety of important asymptotic statistics results [35, 36, 38]; one of his major contributions was concerned with Locally Asymptotic Quadratic (LAQ) expansions of the log likelihood ratio (llr). Denote the llr symbolically by $\Lambda_{h,M}(\theta)$. It is given (for discrete Markovian models) by:

$$\Lambda_{h,M}(\theta) = \mathcal{L}_{\theta + \delta_M h} - \mathcal{L}_{\theta}$$

Here M is again the length of the times series ; h is a vector perturbation with the same dimensionality as θ ; δ_M is a matrix that scales the perturbation and \mathcal{L}_{θ} denotes the llr evaluated at θ ; for later use let $h_M := \delta_M h$. When the assumptions behind LAQ hold [28, 38], one can asymptotically approximate the above statistical experiment by a normal limit experiment ¹⁶. If one is in a neighborhood of the true parameter ($\hat{\theta}$) the LAQ conditions imply

$$\Lambda_{h,M}(\hat{\theta}) - (h_M^T S_M(\hat{\theta}) - \frac{1}{2} h_M^T W_M(\hat{\theta}) h_M)$$

¹⁶See [48, 38] for a clear introduction to this topic, consult [28] for an application to time series analysis

converges (in $P_{\hat{\theta}}$ probability) to 0. In the above expressions, when $\Lambda_{h,M}(\theta)$ is twice differentiable with respect to θ , $S_M(\theta)$ and $-W_M(\theta)$ play the role of the first and second derivative (respectively) of the llr [28]¹⁷ with respect to the parameter vector.

The contiguity condition is a generalization of the concept of absolute continuity [23]; it allows one to determine the asymptotic distribution of “nearby” experiments which is sometimes useful for constructing hypothesis tests. Denote the true parameter by $\hat{\theta}$ and let $\tilde{\theta}$ be contained in a δ_M - neighborhood of $\hat{\theta}$. If contiguity exists in the experiment, it [48] implies the following limit distribution:

$$(4.4) \quad \Lambda_{h,M}(\tilde{\theta}) \xrightarrow{\mathbb{P}_{\tilde{\theta}}} N\left(-\frac{1}{2}h_M^T W_M(\tilde{\theta})h_M, h_M^T W_M(\tilde{\theta})h_M\right)$$

The llr is of interest to us because it provides a method for taking a classical parametric model and producing a scalar random variable which has the above limit distribution in the LAQ case (the practical utility of this theory is realized if the above normal distribution can be used to reliably approximate the more complicated distribution of the llr associated with the proposed parametric model for moderate sample sizes [36]).

There are many other important implications of LAQ, the method is only used in this paper to quantify the uncertainty associated with a nonergodic time series. In the stationary ergodic case the matrix $W_M(\hat{\theta})$ coincides with deterministic quantity \mathcal{F} shown earlier, for nonergodic models the matrix is itself a random variable. I repeat the construction in Le Cam chapter 6 [38] here in order to show how one uses the llr in order to estimate $W_M(\hat{\theta})$ (which can be used to roughly approximate the variance of the limit parameter distribution). To simplify the situation, assume that one is already within a δ_M neighborhood of $\hat{\theta}$ (this neighborhood is centered at $\tilde{\theta}$) and set the matrix $\delta_M = \frac{1}{\sqrt{M}}Id$ where Id is the identity matrix. Suppose $\theta \in \mathbb{R}^k$, then one takes a basis set of \mathbb{R}^k (denote this set by $\{b_1, \dots, b_k\}$) and evaluates:

$$\Lambda_{h,M}(\tilde{\theta} + \delta_M(b_i + b_j))$$

for $i, j = 1, \dots, k$; the components of $W_M(\tilde{\theta})$ are estimated by

$$(4.5) \quad W_M(\tilde{\theta})_{ij} \approx \frac{-[\Lambda_{h,M}(\tilde{\theta} + \delta_M(b_i + b_j)) - \Lambda_{h,M}(\tilde{\theta} + \delta_M(b_i)) - \Lambda_{h,M}(\tilde{\theta} + \delta_M(b_j))]}{b_i b_j}$$

Under the LAQ assumptions, one can claim that $W_M(\tilde{\theta}) \xrightarrow{\mathbb{P}_{\tilde{\theta}}} W_M(\hat{\theta})$ as M tends to infinity. The theory shown here (developed by Le Cam) is purely an asymptotic one, however if one observes that the relation in equation 4.4 holds, it gives one more confidence when equation 4.5 is employed to approximate the quadratic expansion of the llr. In the situation where $W_M(\tilde{\theta})$ is random (as it will be in some of the applications presented), one needs to repeat this procedure for a variety of paths¹⁸ and then take expectations in order to get a rough estimate of the asymptotic

¹⁷This view is convenient for giving an intuitive interpretation, but a classical Taylor expansion of $\Lambda_{h,M}(\tilde{\theta})$ may fail to exist; however the LAQ “derivative” can still be defined

¹⁸In practice one may not have enough paths from an atomistic simulation to realize asymptotic statistics results. If a diffusion model is truly valid, one is always free to use the parameters obtained from a small sample in order to simulate additional paths [30], then use this information to create asymptotic error bounds associated with “perfect” data.

parameter distribution associated with the collection of sample paths. The estimation methods that use the LAQ expansions are purely heuristic, the theory developed above is proved in great detail for independent random variables; asymptotic statements about the llr associated with general Markov chains are harder to make [36, 45]. The potential applications of these types of ideas in multiscale applications are merely shown via numerical results in this paper.

5. Local Polynomial Diffusion Models. Thus far, the main concern has been approximating the transition density and quantifying the goodness-of-fit of our simple models to the data. For the remainder of this section assume that the data can adequately be described by some arbitrary nonlinear effective SDE. Recall, I am working under the assumption that even if the true underlying potential energy surface is rugged, a smooth model of the drift and diffusion coefficients can be used to reliably approximate certain features of the free energy surface. In statistical mechanics, sometimes a limit theory exists for what the large sample system’s effective dynamics converge to [20]. Unfortunately, for many interesting systems, the functional form of the drift and diffusion coefficients are completely unknown. Our research group has used the label “equation free methods” [29] to describe a set of numerical procedures that have been designed to deal with a situation where one has a simulation protocol that is believed to contain useful information and the information in the simulation is believed to be describable by an effective equation (with smooth coefficients), but the equation is unavailable in closed form. The early versions of this procedure used least squares approximation in order to get derivative information. Within the last couple of years, our group has extended the estimation procedure to match local linear SDE models (both vector and scalar) to the output of simulation data. The parametric models proposed are of the type given in equation 3.4. MLE and QMLE are greatly facilitated by the deterministic likelihood expansions developed by Aït-Sahalia [2, 3, 5]

¹⁹

Conceptually, I am just taking advantage of the fact that the accurate Aït-Sahalia Hermite expansion allows us to generalize the piecewise-polynomial (pp) concept to diffusion SDE models. Since I posit the existence of a smooth underlying effective diffusion model, the drift and diffusion coefficient functions are fit locally to linear models. One is free to use a fairly broad class of polynomials in the scalar case with a variant of Aït-Sahalia’s expansion [5], but our future work is concerned with applying the techniques in this paper to the vector case. In that situation, a higher order polynomial of unknown vector functions is not practical from an estimation standpoint. Our experience has shown that if one wants to accurately capture the mean reversion portion of a nonlinear effective model that it usually becomes necessary to model the state dependence of the noise (this point is illustrated in section 6.5). In statistical terms, I am simply just finding the QMLE estimate of the parametric model shown in equation 3.4. However a variety of practical questions arise:

- How does one determine if the estimated model is meaningful?
- In the case where an approximate transition density is used to estimate pa-

¹⁹Deterministic expansions are useful because one can quickly carry out parameter optimization and can easily evaluate the functionals needed to carry out variations of MLE in situations where MLE fails [38, 28]. Another obvious advantage is that in the case of stationary ergodic distributions one can easily carry out a deterministic quadrature and determine the asymptotic parameter distributions in situations where the parameter distributions converge to a normal random variable

rameters, how sensitive is the QMLE projection to the quality of the approximation?

- How does one choose the size of the neighborhood where a local linear model is valid?
- How does one piece together the local models in order to create a global nonlinear effective diffusion model?

The first issue is readily handled by the theory discussed earlier. The second issue is concerned with semiparametric estimation and robust statistics. These are important and active area of research in statistics [9, 48] (and in the future could make great contributions to multiscale modeling), but results are currently difficult (for this author) to translate into practical and reliable numerical methods associated with the ideas laid out here. In section 4.1 it is shown that the true model is not too sensitive to “contaminated data” and this quality is retained by the expansions of Ait-Sahalia, however is not by the overly crude Euler estimator.

A simple estimation procedure is well-equipped to handle the third issue. The method requires the user to obtain a batch of trajectories from a simulation. In order to carry out a classical estimation procedure, one must create a partition of state space where a local linear diffusion model is an accurate representation of the underlying nonlinear diffusion (see figure 5.1). Once the data set is collected, one is free to make as many partitions as desired. The simple idea is to only present observation pairs $(\mathbf{x}_n, \mathbf{x}_{n+1})$ to the log likelihood function that have “ \mathbf{x}_n ” within the selected neighborhood. The neighborhood size must be chosen to be small enough that the QMLE parameters estimates of the pp SDE model are statistically meaningful²⁰ and large enough to contain enough samples to obtain the desired parameter accuracy (and have asymptotic limits remain useful for the sample size). The main difficulty that the method faces is determining a neighborhood yielding a satisfactory compromise between the two aforementioned factors. This procedure is greatly facilitated by the equations of a limit process if one is known. Otherwise one must use numerical experimentation to attempt to determine how smooth the underlying coefficient functions are (recall the assumption that a simple low-dimensional model exists). In multiscale modeling we have the convenience of controlling the initial conditions of the (assumed meaningful) atomistic simulation making this type of pp SDE modeling more practical (otherwise one can only estimate parameters around state points visited frequently by the sample path). Note that this method makes the time series length itself a random variable. In addition, many applications of this idea lead to estimation of a nonstationary time series. It has already been noted earlier that it is computationally difficult to calculate the Fisher information matrix in these circumstances. These facts indicate that the LAQ likelihood ratio approximation might be useful. It should be noted that optimal tests rarely exist in this type of situation, however in the estimation literature some heuristic techniques have been recommended [6, 17]. The technical details of these methods require some fairly specialized arguments and due to this author’s ignorance of both recent developments and the more important aspects of the theory, goodness-of-fit tests are avoided in this situation²¹.

²⁰The goodness-of-fit techniques presented in the previous section can check this, however in practice this procedure greatly benefits from modern nonparametric techniques [26, 12]

²¹Before proceeding, I would like to explicitly point out that if one has an initial distribution of points clustered near each other initially and as time proceeds the ensemble mean slowly “funnels” its way down to global or local minima, then the problem is slightly easier because one only needs to determine the boundary of the “right hand end points” of the collection of time series. The point stressed is that data only needs to be collected *once* and a statistically analyzed on or off-line (the

In regards to the final item, I merely present two simplistic methods of piecing together our local models in section 5. I also demonstrate that the LAQ expansion can give a useful quantification of parameter uncertainty for our toy models. The information contained in an LAQ approximation can be used in more sophisticated “matching” schemes. Again this area is well outside of this author’s research area; if the problem at hand is of any real significance one should consult [11, 16, 49] for modern matching techniques, but at some stage one will almost certainly have to appeal to some form of heuristics (a.k.a., “the art of numerical computation”).

5.1. Relevance to Multiscale Numerical Methods. The details of the tools used may obscure their utility in multiscale modeling, so the connection is summarized here. The computational load to carry out a full simulation is too great, so a diffusion approximation is made which hopes to capture the relevant information in the underlying model. Goodness-of-fit tests are needed to quantify the quality of the approximation. If the model is found to be statistically acceptable, then one would be interested in the confidence intervals associated with the estimate. Different state points have different levels of noise; having a reliable method for theoretically predicting parameter uncertainty as a function of sample size at different state points can assist one in designing “efficient experiments” because one can allocate computational resources in an intelligent manor. QMLE is a nice tool for achieving all these tasks, however it can fail if it is followed “verbatim”, especially when one uses an approximation of the transition density. Estimators based on llr methods are appealing in situations where standard QMLE fails because they are applicable to a wider class of problems and the distribution of the llr asymptotically converges to a manageable distribution for certain model classes [36].

6. Results and Discussion. In order to get MLE (QMLE) parameter estimates in all of the section to follow, the IMSL Nelder-Mead search algorithm is employed using the termination criterion of 5×10^{-8} with an initial parameter distribution dictated by a uniform distribution around the slightly biased (+10%) true parameter values in the cases of known or approximately known models and from an assumed limit process in the final application.

6.1. Maximum Likelihood Estimation Results. Tables 6.1 and 6.2 report the empirically measured mean and standard deviation of the parameter distributions as well as the asymptotic predictions for the standard deviation for situation I A-D and situation I A-C (using data spaced $\delta t = 2^{-5}, 2^{-3}$ units respectively). We see that as β increases the estimated mean reversion parameter decreases in magnitude. This makes sense because the parametric model in equation 3.3 assumes a single free energy minimum; as the magnitude of the sine wave added to the drift increases, one experiences a situation where the “well-depths” associated with the local minima of the glassy potential increase, retarding the rate of mean reversion (centered around the global minimum). The exact magnitude of this effect depends heavily on several factors, some of which are: the magnitude of the noise, the frequency of the sine wave’s perturbation, the amplitude of the sine wave, and the sampling frequency. The last two effects are quantified in tables 6.1 and 6.2. Inspection of these tables shows that the when the true density is used, the effective model estimated from the data is relatively independent of the sampling frequency. Note that as the observation frequency decreases the quality of the expansion (in time) naturally decreases, but

deterministic expansions of Aït-Sahalia make the former a possibility)

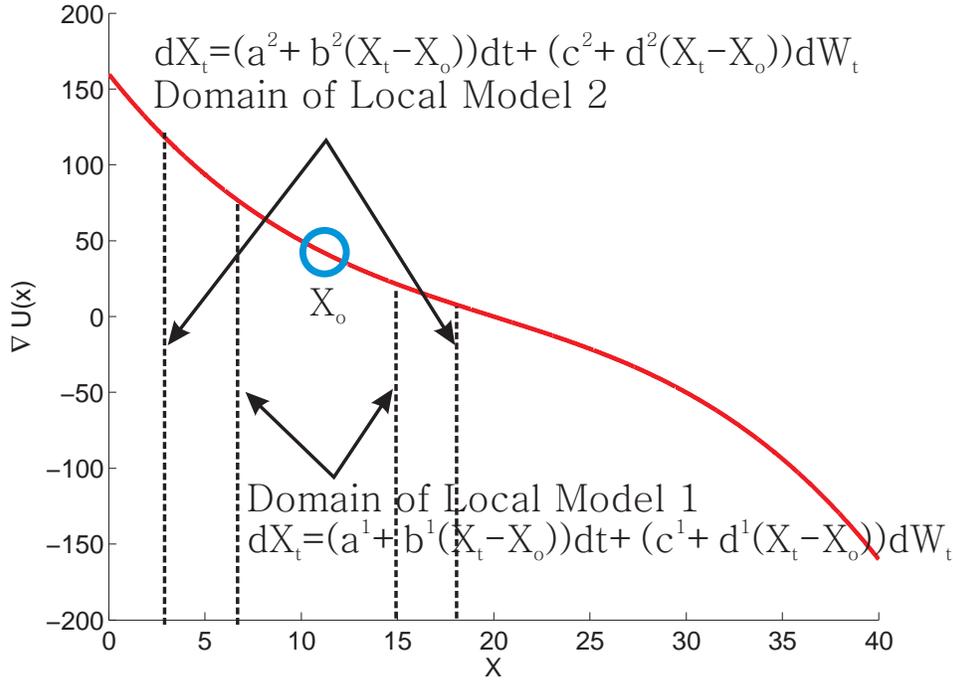


FIG. 5.1. LAQ Screening Method Illustration Applied to Situation II B Data The circle corresponds to a state point where a parametric model was obtained using that particular “ X_0 ” in the parametric model given in equation 3.4. For smooth models the quality of the linear approximation depends on the neighborhood size as well as the curvature of the drift and diffusion coefficients (latter not shown). Once the data is collected, one is free to vary X_0 and/or the neighborhood size where the local SDE is assumed to model the data.

even for relatively “large” times between samples Aït-Sahalia’s expansion remains accurate. The quality of the Euler expansion is very sensitive to the time between samples and it introduces a significant bias even when the “perfect data” is presented to the estimator. I continue to show the results for the Euler estimator despite these shortcomings because in section 6.4 the estimator demonstrates a redeeming quality which can be used in conjunction with the transition density estimator of Aït-Sahalia in situations where the latter fails. The tables also show that the sample size is large enough to partially realize asymptotic results in most cases. Table 6.3 shows similar results; the cubic perturbation introduced into the drift results in a higher mean reversion rate. In section 6.4 it is demonstrated that if one presents “screened” data (those observation pairs that fall within a small neighborhood centered around the estimated α parameter) that the bias introduced by the nonlinear drift perturbation steadily decreases in magnitude as the neighborhood size decreases which intuitively makes sense given our assumptions. Unfortunately, this simple procedure creates a more complicated statistical problem in regards to theoretical parameter distribution predictions because a deterministic approximation of the parameter distribution is much harder to get using this technique.

TABLE 6.1

Situation I Parameter Distributions The data used to obtain the parameter distributions was $N = 2000$ sample paths of an SDE sampled over $M = 4000$ time intervals evenly spaced at $\delta t = 2^{-5}$ taken from the invariant distribution of an Euler path simulation. The empirical mean and standard deviation of the parameter distributions are reported as well as the asymptotic predictions of the standard deviation. For correctly specified models (case A) the asymptotic standard deviation is given by \mathcal{F}_{OP} (calculable by a deterministic integral) and for misspecified models the standard deviation predicted by \mathcal{C} is reported

Scenario/ Expansion Method	$\langle \alpha \rangle$	$\langle \kappa \rangle$	$\langle \sigma \rangle$	$\sigma_{\alpha}^{\text{Asymp}}$	$\sigma_{\kappa}^{\text{Asymp}}$	$\sigma_{\sigma}^{\text{Asymp}}$
				$\sigma_{\alpha}^{\text{Emp}}$	$\sigma_{\kappa}^{\text{Emp}}$	$\sigma_{\sigma}^{\text{Emp}}$
Sit I A True	20.0052	4.0428	4.0168	0.40026	0.26975	0.047621
				0.40152	0.27501	0.047197
Sit I A Ait-Sahalia	20.0454	3.9628	4.0159	0.39938	0.25999	0.047575
				0.40249	0.26681	0.047197
Sit I A Euler	20.0052	3.7961	3.7869	0.40000	0.25298	0.044722
				0.40146	0.24439	0.042119
Sit I B True	19.9985	4.0148	4.0102	0.40089	0.27011	0.047769
				0.39973	0.26831	0.045601
Sit I B Ait-Sahalia	20.0374	3.9362	4.0093	0.3956	0.24097	0.047761
				0.40102	0.26296	0.045592
Sit I B Euler	19.9985	3.7707	3.7822	0.40097	0.23932	0.042992
				0.40052	0.23995	0.041049
Sit I C True	19.8675	3.7844	3.9277	0.41654	0.26111	0.047758
				0.41787	0.27209	0.047093
Sit I C Ait-Sahalia	19.9016	3.7085	3.9267	0.41061	0.23487	0.047743
				0.4196	0.27554	0.047109
Sit I C Euler	19.8674	3.5374	3.7195	0.42007	0.2335	0.043175
				0.41835	0.24537	0.041998
Sit I D True	19.3406	2.9368	3.581	0.49515	0.22916	0.047533
				0.48422	0.27015	0.058569
Sit I D Ait-Sahalia	19.3614	2.8953	3.5787	0.48505	0.21071	0.047373
				0.48679	0.26281	0.050485
Sit I D Euler	19.3387	2.6775	3.4387	0.51899	0.21057	0.0439
				0.48306	0.24238	0.044822

6.2. Optimal Binary Alternative Hypothesis Testing Results. In figure 6.1, the top/right axis corresponds to the Neyman-Pearson critical value versus the theoretical type I error probability for all three transition density expansions using the null as the true (known) parameters of situation II A and the alternative as the parameters obtained using situation II B data with QMLE (using the exact CIR density). The left/bottom axis shows the analytically calculated cumulative distribution of rejecting the null under the alternative using the three different transition densities (plugging in the same alternative parameters estimated by QMLE) as well as the empirically measured distribution of the likelihood ratio (see caption for additional details). We see that the sample size is large enough to realize agreement between the measured and limit distributions. The expansion of Ait-Sahalia provides a very good approximation of this distribution, whereas the distribution predicted by the Euler approximation deviates substantially.

6.3. Goodness-of-fit and Model Misspecification Results. Here it is demonstrated how the various transition densities perform when one tries to use them to evaluate the derivatives needed to estimate some goodness-of-fit statistics shown in [50]. Before proceeding, a mildly problematic aspect of the expansion of Ait-Sahalia is pointed out. To illustrate the problem, the histograms of $\frac{\partial^2 \mathcal{L}_a}{\partial \alpha^2}$ are plotted for 500

TABLE 6.2

Situation I Parameter Distributions Same information as table 6.1 except the time intervals are evenly spaced at $\delta t = 2^{-3}$. Note how the parameter distribution quality degrades (as compared to table 6.1), this is due in part to the failure of the transition density expansion (as evident from situation I A). Another thing to note is that when the true CIR density is used, the parameters estimated are relatively independent of the sampling frequency (which is a very desirable quality). This is not the case for the two transition density expansions, however the magnitude of the discrepancy between this table and the previous is much smaller for the Ait-Sahalia expansion.

Scenario/ Expansion Method	$\langle \alpha \rangle$	$\langle \kappa \rangle$	$\langle \sigma \rangle$	$\sigma_{\alpha}^{\text{Asymp}}$	$\sigma_{\kappa}^{\text{Asymp}}$	$\sigma_{\sigma}^{\text{Asymp}}$
				$\sigma_{\alpha}^{\text{Emp}}$	$\sigma_{\kappa}^{\text{Emp}}$	$\sigma_{\sigma}^{\text{Emp}}$
Sit I A True	19.9996	4.0296	4.019	0.20206	0.16652	0.056847
				0.20085	0.16673	0.057301
Sit I A Ait-Sahalia	20.1162	3.6394	3.9971	0.18208	0.10863	0.055179
				0.22812	0.23505	0.06125
Sit I A Euler	19.9996	3.1652	3.2323	0.20088	0.12692	0.045736
				0.20088	0.10329	0.038508
Sit I B True	19.9915	4.0095	4.0123	0.7049	0.4502	0.03758
				0.20131	0.1659	0.056841
Sit I B Ait-Sahalia	20.118	3.6523	3.9941	0.76493	0.41996	0.037335
				0.22589	0.21442	0.059281
Sit I B Euler	19.9915	3.1517	3.2306	0.30418	0.26104	0.049705
				0.20138	0.10267	0.038656
Sit I C True	19.8718	3.7377	3.9093	0.32825	0.29782	0.048524
				0.20573	0.16222	0.055943
Sit I C Ait-Sahalia	19.9835	3.4578	3.8975	0.33702	0.26354	0.041954
				0.22621	0.21599	0.05993
Sit I C Euler	19.8717	2.9659	3.1967	0.38195	0.24293	0.04877
				0.20574	0.10375	0.03815

TABLE 6.3

Situation II Parameter Distributions Same information as table 6.1 except $M = 4500$.

Scenario/ Expansion Method	$\langle \alpha \rangle$	$\langle \kappa \rangle$	$\langle \sigma \rangle$	$\sigma_{\alpha}^{\text{Asymp}}$	$\sigma_{\kappa}^{\text{Asymp}}$	$\sigma_{\sigma}^{\text{Asymp}}$
				$\sigma_{\alpha}^{\text{Emp}}$	$\sigma_{\kappa}^{\text{Emp}}$	$\sigma_{\sigma}^{\text{Emp}}$
Sit II A True	19.9877	4.0512	4.0169	0.37737	0.25432	0.044898
				0.38691	0.26327	0.04653
Sit II A Ait-Sahalia	20.0272	3.9695	4.0159	0.37654	0.24512	0.044854
				0.38803	0.25535	0.046481
Sit II A Euler	19.988	3.8045	3.7862	0.37712	0.23852	0.042164
				0.38732	0.23417	0.041944
Sit II B True	19.826	4.9256	4.018	0.30542	0.28151	0.045772
				0.31247	0.26417	0.047077
Sit II B Ait-Sahalia	19.8544	4.8679	4.0186	0.29819	0.25567	0.045828
				0.30856	0.25292	0.047127
Sit II B Euler	19.826	4.5645	3.7382	0.30545	0.24273	0.039965
				0.31238	0.22756	0.040756

sample paths using the three expansions. The particular histogram shown using the Ait-Sahalia expansion had 25 observations that were disregarded because of unusually high values in the calculated quantities. The expansion of Ait-Sahalia is very accurate, but its functional derivatives can unfortunately introduce spurious singularities into the transition density approximation. To show a specific example of this, take the logarithm of the order one Ait-Sahalia CIR expansion²² and calculate the second derivative with respect to α and plug in the observation pair (obtained via an Euler

²²Mathematica code available from <http://www.princeton.edu/yacine/research.htm>

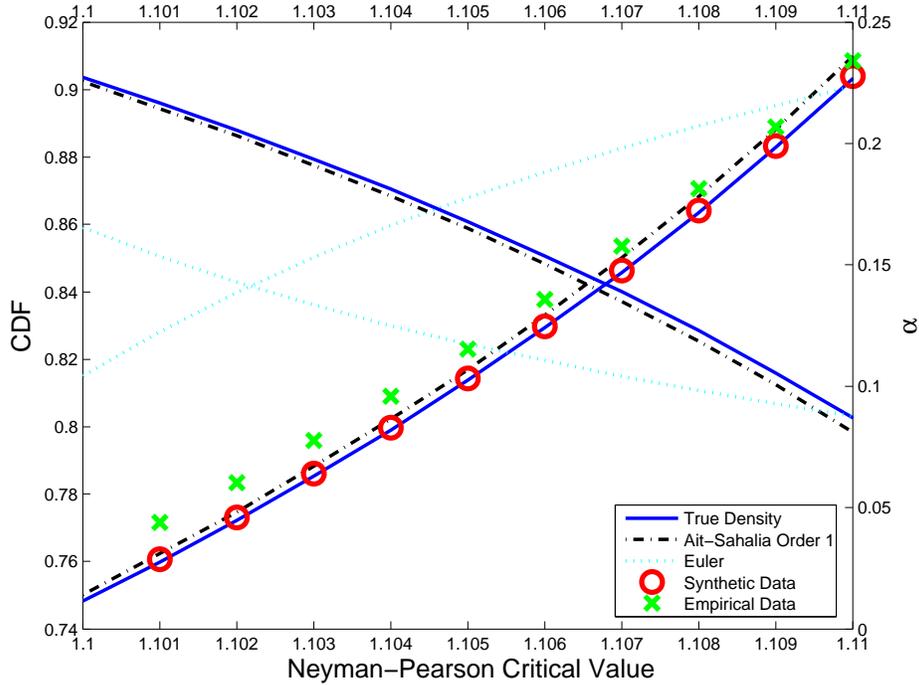


FIG. 6.1. *Neyman-Pearson Results* The Neyman-Pearson test carried out on situation II B data using the QMLE parameter estimate as the alternative and the exact (known) parameters associated with situation II A as the null. The curves represent the deterministic calculation of the type I error probability as a function of the critical value (right axis) and the theoretical CDF of the likelihood ratio obtained by using the various transition density expansions assuming ergodic sampling of the invariant distribution (left axis). The “x’s” correspond to using the actual situation II B data (nonlinear potential) and the “o’s” correspond to the distribution of the likelihood ratio obtained using the QMLE parameters estimated to simulate sample paths of the parametric model proposed (plugging in the average of the QMLE parameter distribution) and then using the exact CIR density to evaluate the likelihood ratio.

path simulation) $[x_n, x_{n-1}] = [3.1826960, 3.305275]$ and set $\sigma = 4, \kappa = 4, \delta t = 2^{-5}$. The bottom panel in figure 6.2 plots the resulting function of α . For the particular model and parameter values shown, this type of situation is only encountered when calculating functional derivatives. In the transition density associated with the model given in equation 3.4, sample values where singularities are hit when evaluating the pure log likelihood function are encountered. A standard method for dealing with this is to apply a one-step MLE estimator²³. A minor modification of Le Cam’s method shown in section 4.4 provides one possible construction of a one-step MLE estimator. Many one-step methods require a parameter guess that is within a “reasonable neighborhood” of the true parameter value (the exact size of the neighborhood can be

²³The basic ideas of these procedures is to use a simple restricted parameter set that is “rich” in the parameter space [9]. An example of this would be to take a discrete mesh of parameter space and optimize the log likelihood over this finite set. The main idea is to keep the parameter values away from singularities associated with the finite sample log likelihood. Refer to the literature on asymptotically centering estimators in [38, 35] for another example of a one-step method.

TABLE 6.4

Le Cam LAQ One-Step Test The data used was situation II A with $\delta t = 2^{-5}$. The LAQ motivated one-step expansion was carried out at the point $\theta = (20.1, 4.5, 4.1)$, consult Le Cam [38] chapter 6 for details. The one-step estimate of the parameter results from adding the quantity below to θ . This type of procedure is necessary when standard MLE misbehaves. The table shows that the LAQ update gets one close to the true parameter when an extremely accurate approximation of the transition density is in hand. The Ait-Sahalia order one and Euler updates are not as good, but still usable. The major problem with these estimators is that the contiguity condition becomes difficult to verify (a condition that needs to be met before the LAQ expansion can be used with confidence). The last two columns display -2 times the mean and the variance of the llr measured by evaluating it at $\hat{\theta}$ and $\hat{\theta} + h_M$ where $h_M = (1.2, 0, 0.12)$. In the infinite sample limit these two quantities will be identical (assuming h_M is continually scaled properly). For the exact and order four Ait-Sahalia expansion one observes close results. For the other two expansions this is not the case (see text for further discussion).

Expansion	$\Delta\alpha$	$\Delta\kappa$	$\Delta\gamma$	$-2\mu_{\Lambda_{\theta\text{true}}}$	$\sigma_{\Lambda_{\theta\text{true}}}$
CIR Exact	-0.9507	-0.4791	-0.0931	12.9402	13.8679
Ait-Sahalia Order 4	-0.9507	-0.4791	-0.0931	12.9405	13.8679
Ait-Sahalia Order 1	-1.0065	-0.7499	-0.1208	13.6013	39.0793
Euler	-0.9621	-0.7344	-0.3156	41.8144	16.0470

chosen using an approximation of the Fisher information matrix). Table 6.4 illustrates that if one starts with a slightly perturbed guess of the true parameter that Le Cam’s method can get fairly close to the true parameter vector if the transition density approximation is very accurate (the llr expansion was obtained from situation II A data and the one-step used is outlined in Le Cam [38] chapter 6). The asymptotic likelihood expansions are valuable in parametric estimation; unfortunately this application requires an extremely accurate transition density approximation²⁴. Throughout this paper the first order (in time) Ait-Sahalia expansions have been used; only in table 6.4 do I report results from a higher order expansion (order four). Recall in section 6.1 that the order one Ait-Sahalia expansion provided an accurate approximation of the likelihood ratio for a simple hypothesis test. The test statistic generated for the Neyman-Pearson test only required the ratio of one observation pair at two nearby parameter vectors. The transition density expansion does introduce a systematic bias into the approximation, and the nature the bias does exhibit a smooth dependence on the parameter values making the ratio of two nearby probability densities also exhibit a significant bias and these small errors accumulate as the time series length grows (see equation 4.3) complicating matters. Some techniques and analysis have already been developed for approximating the LAQ expansion of random variables in the case where the density is not known explicitly [28]; most techniques developed require one to empirically measure the transition density. For our applications, the empirical distribution techniques are mildly inconvenient (from a computational standpoint) because one needs to determine an empirical density approximation for each observation pair.

Now let us return to the goodness-of-fit issue. First, the condition $\mathcal{F}_{Hessian} =$

²⁴Instead of applying one-step methods (a full discussion would slightly overload this paper and distract from the simpler points), in the simple parametric models studied I remedy the singularity issue by using the Euler approximation to determine when a singularity is hit. It is possible to distinguish between spurious singularities from log likelihood function singularities in the CIR case because the true transition density is known. In all of the CIR applications the singularities hit were in fact spurious. I simply threw out sample paths where the absolute value of the logarithm of the transition density of the Ait-Sahalia expansion differed from that of the Euler approximation by a factor of three anywhere along the discretely observed path (this occurred less than 4% in the CIR studies, in the SSA studies this number increased to roughly 25%).

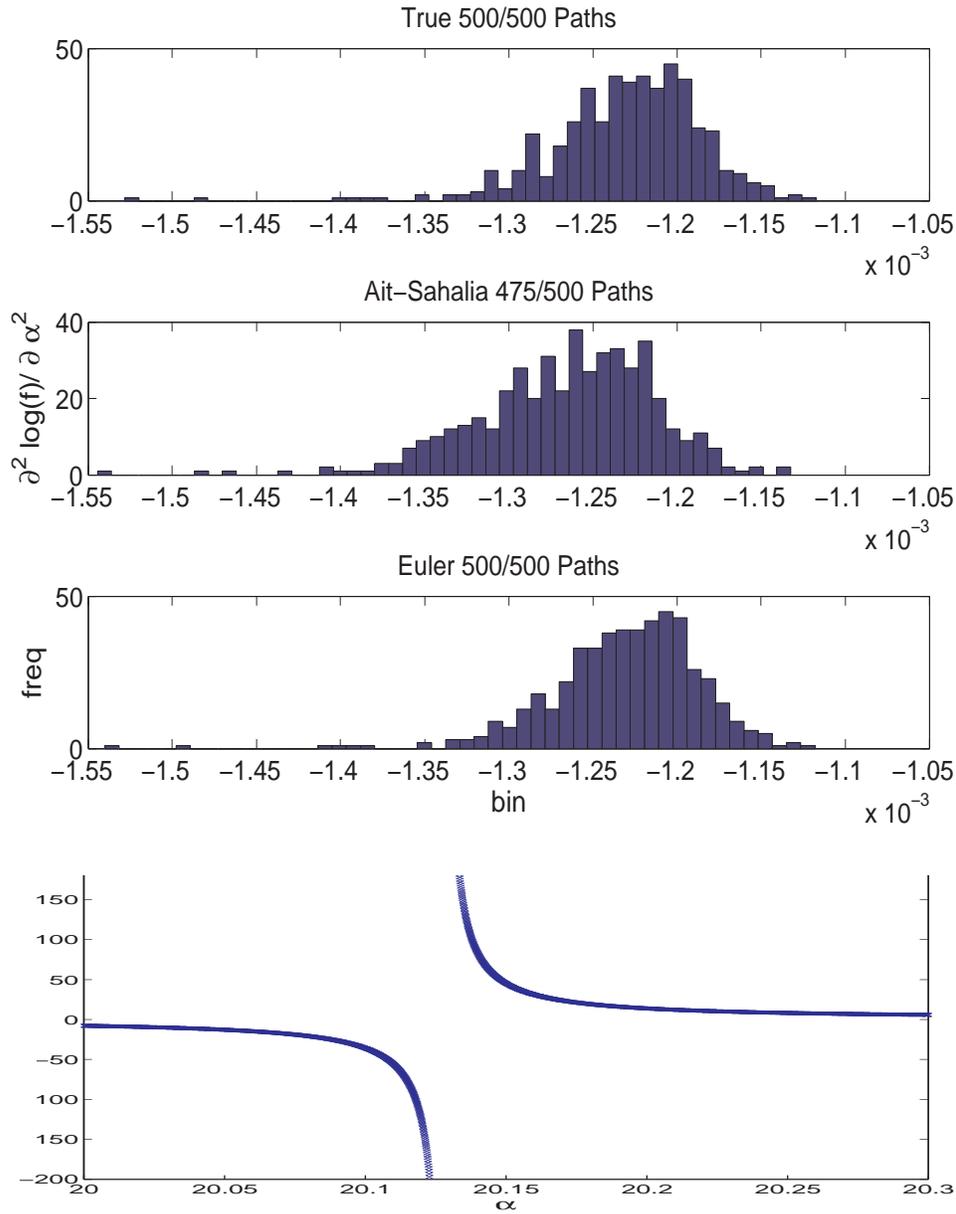


FIG. 6.2. *Histogram of Diagonal Component of $\mathcal{F}_{Hessian}$ Corresponding to α for the Three Transition Density Expansions* The histograms all have slightly different shapes indicating that the curvatures of the three transition densities differ (data taken from situation I C with $\delta t = 2^{-5}$). In the second plot 25 observations were thrown out because there were outliers caused by spurious singularities in the transition density expansion. The bottom panel gives one demonstration of how the Ait-Sahalia expansion can introduce spurious singularities into the log-likelihood function. The data shown is the second derivative of transition density expansion with respect to α as a function of α for a particular observation (see text for discussion). The fat tails of the log likelihood function make a simple screening of outliers (caused by spurious singularities) very difficult if one does not have the true transition density to reference.

$-\mathcal{F}_{OP}$ it tested. If the condition holds $T := \sum_{i,j=1}^k \mathcal{F}_{Hessian}^{ij} + \mathcal{F}_{OP}^{ij}$ should be a mean zero random variable (the superscripts denote matrix components and k is the number of parameters estimated). If it is assumed that the classical central limit theorem (CLT) holds and that the sample sizes used are large enough to appeal to the CLT and approximate the sample mean by a normal distribution of unknown variance, then the classical t-test can be employed [8] (more sophisticated tests are proposed in [50]). Figure 6.3 plots the t-test results for testing if $F_{Hessian} = -\mathcal{F}_{OP}$ for various sample sizes (the Ait-Sahalia data is screened by the technique mentioned in footnote 24). The figure shows that as the number of paths grows (the equality tested is valid in the infinite time series length limit; here the time series length is held fixed), that it becomes easier to detect discrepancies between the data and the assumed model. Even when the true transition density is used with the correct model, one observes that as sample size increases (in paths) that it becomes easier to reject the null. This is because the equality $\mathcal{F}_{Hessian} = -\mathcal{F}_{OP}$ is an asymptotic result. A finite size time series can *never* fully realize asymptotic results; as more paths are analyzed the departure from the limit becomes easier to detect. One should also note that as the number of paths increases it also becomes easier to detect the errors in the transition density expansion. Before proceeding to check for model misspeciation in the presence of an approximated transition density, one should determine if asymptotic results are valid for “perfect data” using the same time series length, then proceed to check the accuracy of the expansion by some MC tests (like the ones presented in the previous section).

For extreme cases (situation I C and D) the t-test safely rejects the null for small sizes (see figure 6.3 and the inset). This indicates that the data is well outside the particular parametric model class under study, but one can still test if some asymptotic QMLE results hold. Figure 6.4 plots the exact probability density of the $\chi^2(3)$ random variable as well a test statistic given in [50] (which will be denoted by \mathcal{HW} ²⁵) We see that both the approximation and the exact transition density appear close to the predicted limit distribution indicating that the asymptotic results approximately hold for the sample sizes used. I quantitatively compare how well the empirical distributions match the limit distribution by using the Kolmogorov-Smirnov (KS) test with various sample sizes in the inset of figure 6.4²⁶. The average p-value is obtained by applying the one sample KS test [8] using 500 draws (with replacement) of size $N_{samples}$ from the pool of \mathcal{HW} random variables associated with the $N_{samples}$ paths using the $\chi^2(3)$ density as the null. The result of this procedure is shown in the inset of figure 6.4. We see that the test statistic created using the Ait-Sahalia expansion is rejected before that associated with the true density indicating the possibility that the errors in the expansion cause a *mildly* inflated rejection rate

²⁵ $\mathcal{HW} := Mg(\theta) \left[\nabla g(\theta) C_{(\theta)} \nabla g(\theta)^T \right]^{-1} g(\theta)^T$ where $g(\cdot)$ is the gradient (written as a row vector)

of the log likelihood function. If $\theta \in \mathbb{R}^k$ then under conditions stated in [50] this random variable converges in distribution to a $\chi^2(k)$ random variable. This statistic was originally proposed by Halbert White [50].

²⁶In this particular application, we are still significantly away from the limit distribution. The sample sizes used for estimation are large enough that a test as powerful as the KS would reject equality of the two distributions with very high certainty if one used all of the data at hand. For this reason I only present portions of the data to the KS test. In many interesting atomistic simulation studies, one can only afford to generate a couple hundred sample paths making this type of goodness-of-fit test useful in practice. An alternative would be to partition the CDF into bins and use the χ^2 square goodness-of-fit test [39].

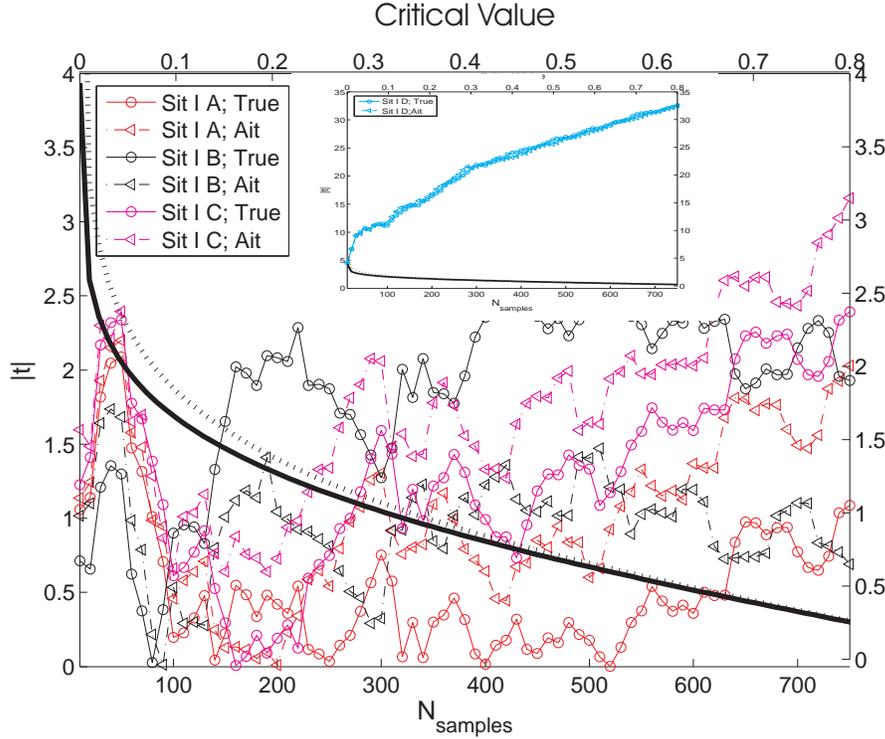


FIG. 6.3. *Model Misspecification t-test Results* The left/bottom axis plots the t-statistic obtained by summing the components of \mathcal{F}_{OP} and $\mathcal{F}_{Hessian}$ for various sample sizes (time series length is fixed, different $N_{samples}$ values correspond to using different sample path numbers to create the random variables required to generate the sample mean and standard deviation needed for the t-test). The right/top axis is to be used with the lines without symbols; the top axis is the critical value (cv) of the t-test (lowest $|t|$ needed to reject the null for a given α). Samples were drawn at random, but once a T r.v. was drawn it contributed to the cumulative average plotted. The solid line corresponds to the theoretical cv of a sample size =375 and the dashed line corresponds to a sample size =10. The inset plots the situation I D test which can be rejected quickly with high confidence.

in situation I C.

6.4. Le Cam’s Method and Likelihood Ratio Expansions Results. Table 6.5 presents parameter estimates of the screened data (the neighborhood size used is given in the captions) from situation II B. We see that as the neighborhood size decreases the estimator mean moves closer to that of the “uncontaminated” model but the variance of the estimator increases. I use Le Cam’s LAQ expansion with the Ait-Sahalia transition density expansions in order to quantify the uncertainty in the measurement. One should observe that the LAQ matrix is closely related to the variance of the parameter distributions (for a more precise description of the connection consult [28, 6]). The agreement between the estimated and measured parameter variance is not as sharp as it was in the case of a stationary distribution, but there one had the convenience of evaluating a deterministic integral (in situation A). Furthermore the asymptotic results are harder to realize because some observation pairs are not used to obtain parameter estimates (due to the screening). If one has

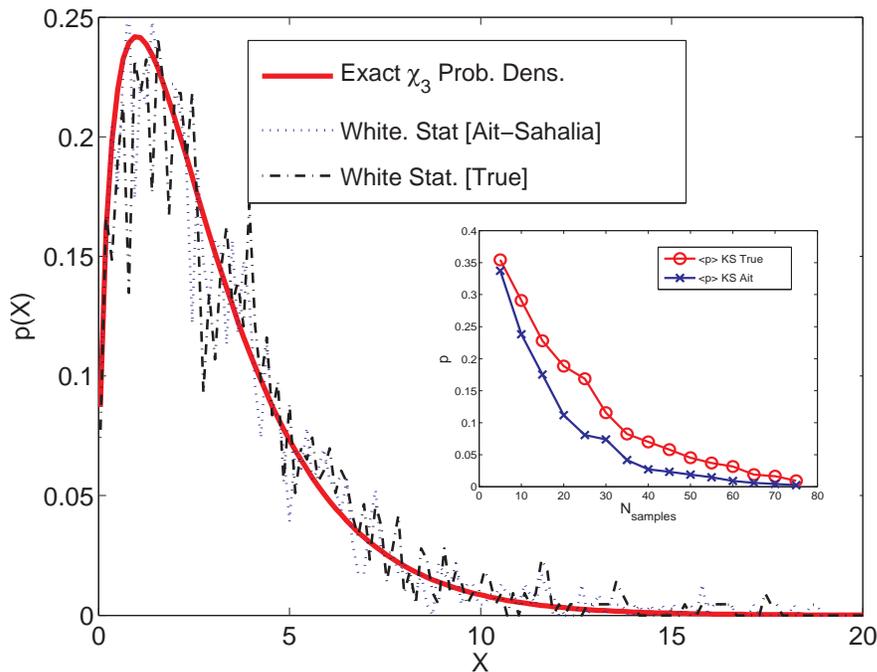


FIG. 6.4. *Goodness-of-Fit with Kolmogorov-Smirnov Test* The $\chi^2(3)$ probability density is plotted along with the empirical distribution of the test-statistic proposed in [50] for situation I C data. The inset shows a plot of the p-value (minimum α needed to reject the null given the data) obtained using the one-sample Kolmogorov-Smirnov test versus the number of paths used to create the test statistic. If the full set of samples were used, the test would have no problem in rejecting the null. This is due to the fact that asymptotic results are never truly realized with finite sample sizes. The plot illustrates that the test statistic generated by the Ait-Sahalia expansion does faithfully capture the average curvature of the log likelihood function in a misspecified model.

a situation where the neighborhood size is small enough to give “meaningful” QMLE parameter estimates and the number of observation pairs is large enough to appeal to asymptotic methods, then the LAQ screening method can be used to define a cost function which can be used to intelligently choose the screening neighborhood size.

6.5. Local Polynomial Diffusion Models Results. In the first application, the pp SDE method is applied to data generated by the unperturbed CIR model ($\gamma = \beta = 0$); I pretend that the functional form of the drift or diffusion coefficient of the data generating process is unknown and instead use the model in equation 3.4. Five arbitrary state points denoted by $\{E_i\}_{i=1,5}$ (shown as circles in figure 6.5) are chosen and the parameters of the affine SDE are obtained there. The LAQ expansion is used in order to obtain parameter uncertainty estimates and the results are compared to the empirical parameter distribution measured by carrying out QMLE on the screened data in table 6.6 .

At this point we have in hand, estimates of the constant and linear sensitivities of the coefficients of the SDE. The “global” drift and diffusion coefficient functions

TABLE 6.5

Situation II B Parameter Distributions with LAQ Screening The data used to obtain the parameter distributions was the same as that used in table 6.3 except here only results of using the Ait-Sahalia expansion are presented using screened data. The base point (X_o) of all of the local SDE models shown below is 20 which corresponds to the true α parameter and the interval used to filter the data is given in the first left column. The empirical mean and standard deviation of the parameters are given along with the LAQ prediction of the parameter uncertainty.

Neighborhood Size	$\langle \alpha \rangle$	$\langle \kappa \rangle$	$\langle \sigma \rangle$	$\sigma_\alpha^{\text{Asymp}}$	$\sigma_\alpha^{\text{Asymp}}$	$\sigma_\kappa^{\text{Asymp}}$
				$\sigma_\alpha^{\text{Emp}}$	$\sigma_\kappa^{\text{Emp}}$	$\sigma_\sigma^{\text{Emp}}$
(16, 24)	19.993	4.080	4.006	0.577	1.079	0.094
				0.560	1.008	0.085
(15, 25)	20.047	4.228	4.003	0.502	0.780	0.078
				0.530	0.802	0.078
(14, 26)	20.021	4.365	4.006	0.450	0.629	0.068
				0.474	0.627	0.066

TABLE 6.6

Piecewise Polynomial SDE Parameter Distributions of Situation II B The data used to obtain the parameter distributions was $N = 2000$ sample paths of an SDE sampled over $M = 4000$ time intervals evenly spaced at $\delta t = 2^{-5}$ using expansion point X_o with the neighborhood size given in the second column. The empirical mean and standard deviation of the parameter distributions are reported as well as the LAQ predictions of the standard deviation.

Xo	Neighborhood	$\langle a \rangle$	$\langle b \rangle$	$\langle c \rangle$	$\langle d \rangle$	σ_a^{Asymp}	σ_b^{Asymp}	σ_c^{Asymp}	σ_d^{Asymp}
						σ_a^{Emp}	σ_b^{Emp}	σ_c^{Emp}	σ_d^{Emp}
16	(11, 22)	15.909	-4.227	15.600	0.534	1.939	0.696	0.295	0.067
						1.999	0.713	0.278	0.061
18	(13, 25)	8.226	-4.118	16.987	0.462	1.898	0.6344	0.288	0.061
						1.984	0.667	0.293	0.062
20	(16, 24)	-0.021	-3.879	17.862	0.437	2.408	1.131	0.425	0.085
						2.419	1.157	0.43774	0.088
22.5	(15.5, 29.5)	-9.994	-4.064	18.971	0.435	2.284	0.624	0.334	0.058
						2.490	0.615	0.338	0.0848
25	(21, 29)	-20.602	-4.050	19.982	0.391	3.310	1.092	0.514	0.091
						3.648	1.426	0.598	0.105

can be approximated by the following interpolation procedure:

$$f_j(x) = \frac{w^L f_j^L(x) + w^R f_j^R(x)}{Z}$$

where $j = 1, 2$ corresponds to the drift and diffusion coefficients respectively, w^L and w^R are the weights associated with the left and right expansion point, $f_j^L(x)$ is the affine approximation to the nonlinear function based on the closest expansion point whose value is less than or equal to x (similarly for $f_j^R(x)$, but use the nearest expansion point strictly greater than x). The weights were assigned by the *ad hoc* rule:

$$w^L(x) = \frac{(\sigma_C + \sigma_D(x)(x - E_i))^{-1} (E_{i+1} - x)}{E_{i+1} - E_i}$$

$$w^R(x) = \frac{(\sigma_C + \sigma_D(x)(E_{i+1} - x))^{-1} (x - E_i)}{E_{i+1} - E_i}$$

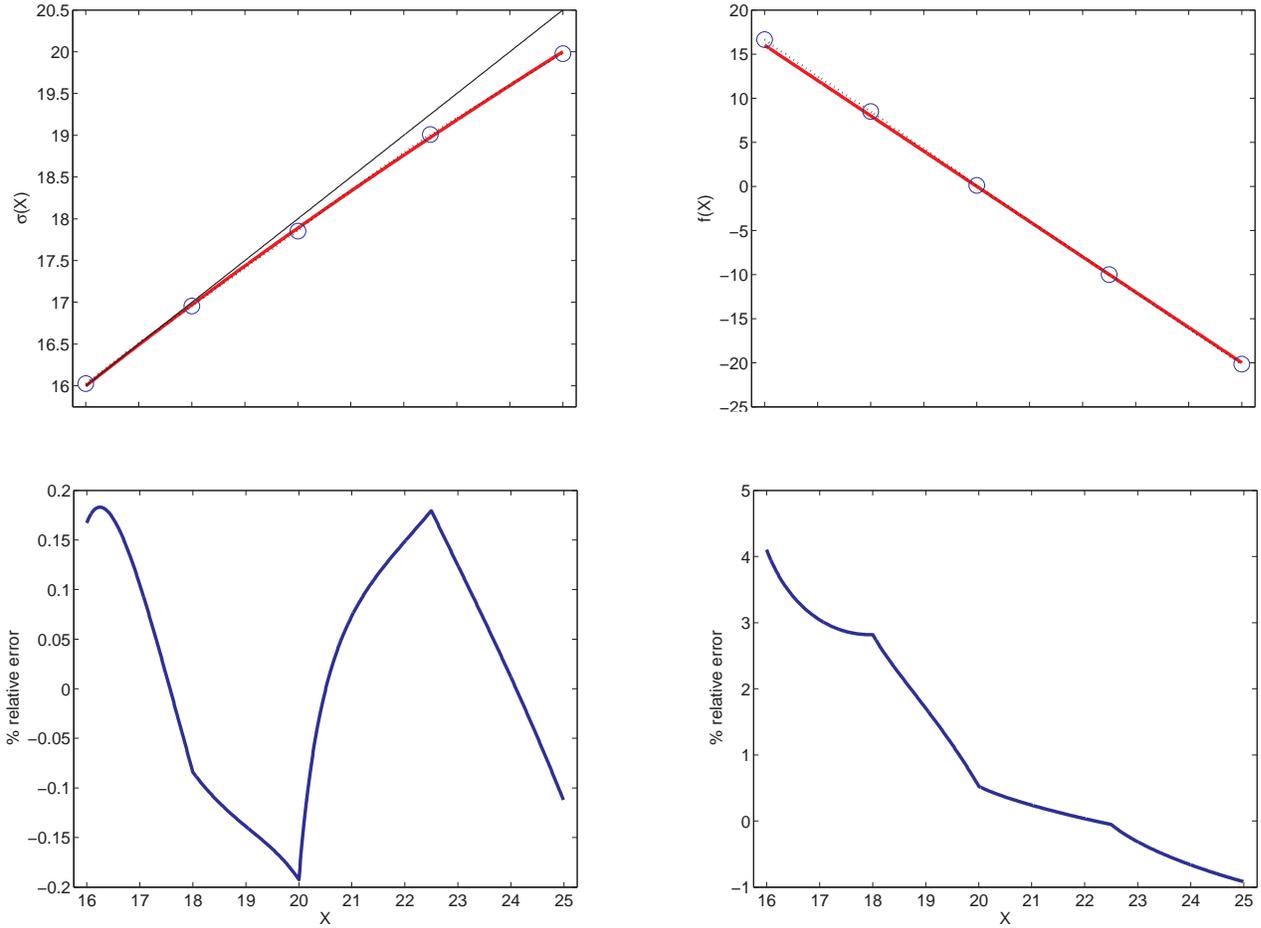


FIG. 6.5. *Piecewise Polynomial Approximation of CIR Model* The top panels plot the diffusion (left) and drift (right) function obtained by the interpolation procedure given in section 6.5. The top left panel plots the first order Taylor expansion (evaluated at $X_o = 16$) of the diffusion coefficient from the known function to show how much the true function deviates from linearity and how well the procedure detects this change. The bottom figures plot the corresponding relative errors (using the exact known SDE coefficient functions). See footnote 27 for a discussion on systematic versus random errors.

$$Z = w^L + w^R$$

The results of this simple interpolation procedure are shown as a dotted lines in figure 6.5 (in the rule above, σ_C and σ_D are the empirically measured parameter standard deviations of the constant and linear term of the diffusion term; the interpolation rule for the drift is analogous). Figure 6.5 plots the relative error of using this procedure using the known SDE coefficient functions as the “truth”. The errors in the diffusion coefficient function do not indicate a systematic bias in contrast to the drift coefficient ²⁷.

²⁷I offer an intuitive explanation for this fact; if one *naively assumes* that the true paths come

In the final application, the parameters of the model are unknown. The parameters were obtained with the LAQ screening technique. The SSA process was run until N sample paths reached approximately “invariant distributions”²⁸. Then $N = 200$ paths were used to obtain initial parameter estimates; for expansion points in the tails of the “invariant distribution” an additional $N = 600$ paths were simulated in order to get sharper estimates of the poorly sampled state points²⁹. Here a B-spline [11] was used (MATLAB’s cubic smoothing spline, “spaps” was used) to piece together the constants (a, c) of the local models in order to smoothly interpolate between state points. Each point was given the same weight when creating the spline, the LAQ technique for measuring parameter uncertainty could be used to give different weights to the measured parameters, but this procedure was not carried out here because the inherent jump nature of the data complicates matters slightly³⁰. The information contained in the linear coefficients (b, d) was not used for interpolation purposes because the variance of the parameter of the linear terms was much larger than those of the constant terms in the case studied. However, if one ignores the linear terms in the local model the constants estimated are significantly affected (see figure 6.6). One observes significant departure of the SSA model parameters from the limiting drift function (plotted as a dashed red line). The number of particles in the SSA system was increased to 2×10^5 and the parameters of the effective model in the drift were estimated and demonstrate that the limiting drift function is measured with the estimators used (see inset in figure 6.6). As the system size increases, the noise decreases, limiting the state points visited (hence the drift function estimated does not cover as large a range as the first SSA case considered). The estimated diffusion function in figure 6.6 demonstrates dependence on the transition density used (Ait-Sahalia order one and Euler). In this application the Ornstein-Uhlenbeck (OU) model was also used to demonstrate the importance of accounting for state dependent noise (note the systematic difference in the constant noise parameter estimated using

from an Euler simulation with Euler step sizes corresponding to the observation frequency then one is in the “Gaussian case”. In between expansion points, the Taylor expansion of the known diffusion function consistently over estimates the diffusion function at nearby points (the diffusion coefficient is concave). The *true* transition density of the SDE has a smaller variance (the data generating process) than the deterministic Taylor series prediction. When one uses the simple Euler density approximation in QMLE, the magnitude of the mean reversion parameter appears to be larger than it really is because of the error in the Taylor series approximation of the diffusion function (larger noise is expected making the deterministic trend appear stronger). The actual situation is much more complicated due to the fact that the distributions are not Gaussian, a complicated nonlinear function is used to obtain parameter estimates (QMLE), finite sampling effects, etc.; however figure 6.5 is consistent with this overly simplified intuitive explanation.

²⁸It is known that the free energy surface of this model contains two stable free energy wells and one saddle; no particles were able to overcome the large free energy barrier for the time series lengths used, hence the quotes on “invariant distributions”

²⁹Parameters were initially optimized over individual paths in order to estimate the parameter distribution variance. When parameters were optimized on a pathwise basis, a significant fraction ($\approx 25\%$) of observation pairs resulted in *assumed* spurious singularities in the Ait-Sahalia expansion. To constrain the parameter space explored in the optimization, I found the QMLE parameters with *all* of the data (over time and paths). This helped prevent the optimization routine from attempting to evaluate the log likelihood function at parameter values that cause spurious singularities because the parameter space explored was reduced because the trial QMLE parameters needed to be good for all of the paths.

³⁰In practice one could overcome this difficulty by obtaining the QMLE, generate SDE sample paths with the model parameters obtained, then find the LAQ parameter variance of a genuine diffusion. In applications where the uncertainty associated with using the local SDE model technique on imperfect data is desired, the problem is much harder. One should consult the specialized literature [9, 45, 48] for guidance; this author can not make any sound general recommendations.

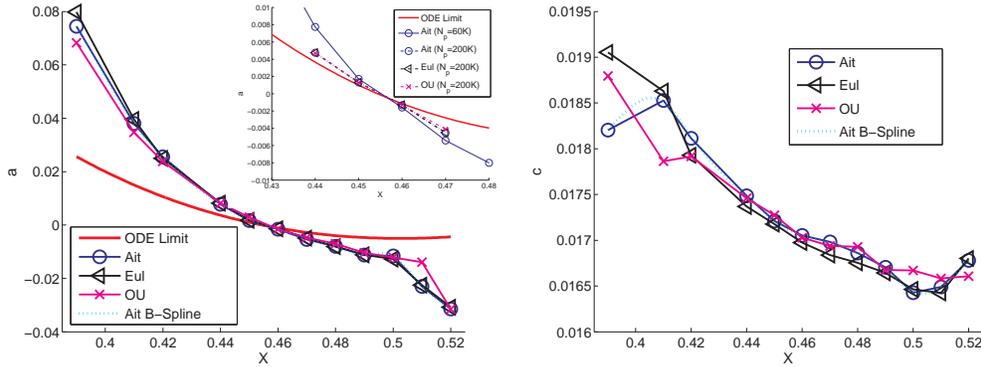


FIG. 6.6. *Estimated Nonlinear Drift (Global)* The LAQ screening method was used to obtain the plots shown. The results of three different transition density estimates are shown (the OU model sets d in equation 3.4 equal to zero). The left panel contains the drift coefficient corresponding to a SSA simulation containing 60^2 particles. The thick solid line corresponds to infinite system size drift function. The dotted line corresponds to the B-spline of the Ait-Sahalia Expansion. The inset shows the drift coefficient corresponding to a SSA simulation containing 2×10^5 particles. Observe how the drift function convergence towards the expected limit equation with “larger” systems, however for small particle systems the results are substantially different. The right panel displays the measured diffusion coefficient (the infinite sample limit has zero noise)

the three different transition density estimators).

Figure 6.7 plots the “invariant” empirical CDF of the SSA data versus that predicted by a sample path simulation of our obtained nonlinear diffusion approximation (constructed from the B-spline interpolation of the estimated local SDE models). The inset plots the difference between the two empirical CDF’s. Our interest in this application was in getting a parametric description of the “invariant” distribution, if the dynamics of the process are of more interest consult [12] for a useful test which is made possible if one has an approximation of the transition density.

7. Conclusions and Outlook. In this paper I have demonstrated that the expansions of Ait-Sahalia can potentially be a useful tool in the parametric estimation associated with computational studies of multiscale systems. Parameter estimates can accurately be obtained and the curvature of the model is accurately represented by the expansion in a variety of applications. These facts can be exploited to estimate parameter distributions and construct useful inference procedures. The overly simple Euler expansion behaves poorly in accuracy of the estimate and in the curvature of the transition density (as shown early on by Lo [40]), but it has the redeeming feature that it does not introduce any spurious singularities into the transition density (for the CIR model). In large sample statistics applications, point singularities can *usually* be remedied [35, 36, 38] by using likelihood ratio expansions. Unfortunately, many of these techniques require an extremely accurate estimate of the transition density. For imperfect transition densities with systematic errors, this becomes mildly problematic (from a computational standpoint)³¹.

I have also demonstrated a heuristic method for locally approximating a nontriv-

³¹Methods proposed in [5] are applicable to a wider class of models and help the accuracy of the transition density expansion in the scalar case, but the vector case poses a more challenging problem

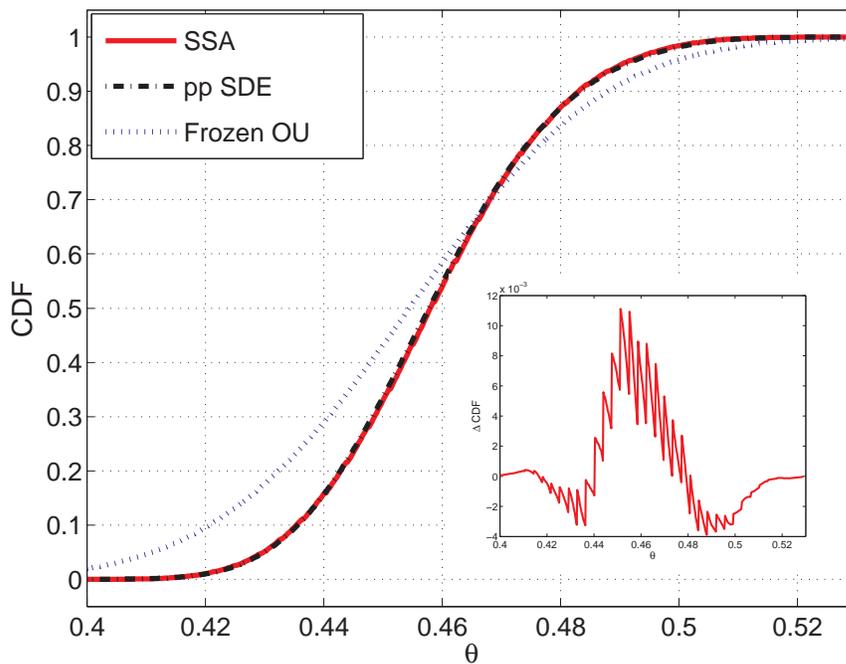


FIG. 6.7. *CDF and KS test* ($N_{molecules} = 3600$) The CDF of the SSA “invariant distribution” (empirically measured) plotted against that of the piecewise-polynomial (pp) local SDE model (simulated by long time Euler integration [30]). In addition the (actual) invariant distribution of the OU model obtained is plotted (the model parameters of the local OU model were measured at the *single* state point where the drift is zero). This was done to show what effect neglecting the nonlinear drift and state dependent diffusion has on the invariant CDF. The inset plots the difference between the “invariant CDF” of the SSA data and that of the pp SDE model.

ial SDE by a collection of simple local models. The application was inspired by the need to accurately measure the parameters, quantify the uncertainty and determine the goodness-of-fit of parametric diffusion models around atomistic data. The pp SDE technique presented is simple in nature, but it raises many deep questions. Insight from the semiparametric, robust and large sample statistics communities would greatly assist in further developing this type of numerical method. The general method is appealing because it can be used to estimate nonlinear effective diffusion models where the effective SDE’s coefficients are smooth, but of unknown functional form. The resulting parametric model structure (which is typically a complicated function due to the “matching” used) can then be passed on to diffusion path simulation methods in order to generate additional data which can be used in order to construct confidence bands or carry out established inference procedures.

From a practical point of view it would be desirable to apply the techniques in this paper to Ait-Sahalia’s method in the vector case because many interesting physical systems depend on a couple of “reaction coordinates” [27, 46, 14]. Unfortunately the vector version of the expansion usually requires one to use an additional Taylor series approximation (which increases the chance of spurious singularities and decreases the

quality of the curvature estimate); this researcher has been able to use the expansions in order to get useful parameter estimates, but has not had as much success in pushing them as far as the scalar expansions. A numerical method which can be used in conjunction with the expansions of Ait-Sahalia in order to get detailed information about the log likelihood ratio associated with smooth SDE's (diffusion and jump models) is currently being explored by the author.

8. Acknowledgements. The author would like to thank Ioannis Kevrekidis and Adam Meadows for comments and advise and Yacine Ait-Sahalia for providing help with the transition density expansions. Any errors are entirely due to the fault of the author.

REFERENCES

- [1] S. ETHIER AND T. KURTZ, *Markov processes : Characterization and convergence*, Wiley, 1986.
- [2] Y. AÏT-SAHALIA, *Transition densities for interest rate and other nonlinear diffusions*, Journal of Finance, 54 (1999), pp. 1361–1395.
- [3] ———, *Closed-form likelihood expansions for multivariate diffusions*, Technical Report, (2001).
- [4] ———, *Maximum-likelihood estimation of discretely-sampled diffusions: A closed-form approximation approach*, Econometrica, 70 (2002), pp. 223–262.
- [5] G. BAKSHI AND N. JU, *A refinement to Ait-Sahalia's (2002) "Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach"*, Journal of Business, 78 (2005), (to appear).
- [6] I. BASAWA AND D.J. SCOTT, *Asymptotic Optimal Inference for Non-Ergodic Models*, Springer-Verlag, 1983.
- [7] B.M. BIBBY AND M. SØRENSEN, *Martingale estimation functions for discretely observed diffusion processes*, Bernoulli, 1 (1995), pp. 17–39.
- [8] P. J. BICKEL AND K.J. DOKSUM, *Mathematical Statistics: Basic Ideas and Selected Topics*, Prentice Hall, Upper Saddle River, NJ, 2001.
- [9] P. J. BICKEL, C. A. J. KLAASSEN, Y. RITOV, AND J. A. WELLNER, *Efficient and Adaptive Estimation for Semiparametric Models*, Johns Hopkins University Press, Baltimore, 1993.
- [10] R.A. CURTIS AND M.W. DEEM, *A statistical mechanics study of ring size, ring shape, and the relation to pores found in zeolites*, J. Phys. Chem. B, 107 (2003), pp. 8612–8620.
- [11] C. DE BOOR, *A practical guide to splines*, Springer, New York, 2001.
- [12] F.X. DIEBOLD, T. GUNTHER, AND A. TAY, *Evaluating density forecasts with applications to financial risk management*, International Economic Review, 39 (1998), pp. 863–883.
- [13] D.J. EARL AND M.W. DEEM, *Optimal allocation of replicas to processors in parallel tempering simulations*, J. Phys. Chem. B, 1078 (2004), pp. 6844–6849.
- [14] B. ENSING, A. LAIO, F.L. GERVASIO, M. PARRINELLO, AND M.L. KLEIN, *A minimum free energy reaction path for the E2 reaction between fluoro ethane and a fluoride ion*, JACS, 126 (2004), pp. 9492–9493.
- [15] R. ERBAN, I.G. KEVREKIDIS, D. ADALSTEINSSON, AND T. ELSTON, *Gene regulatory networks: A coarse-grained, equation-free approach to multiscale computation*, arXiv:physics/0508112, (2005).
- [16] J. FAN, *Nonlinear time series: Nonparametric and parametric methods*, Springer, New York, 2003.
- [17] P. FEIGIN, *Asymptotic theory of conditional inference for stochastic processes*, Stochastic Processes and their Applications, 22 (1986), pp. 89–102.
- [18] D. FRENKEL AND B. SMIT, *Understanding Molecular Simulation: From Algorithms to Applications*, Academic-Press, 2002.
- [19] A.R. GALLANT AND G. TAUCHEN, *Which moments to match?*, Econometric Theory, 12 (1996), pp. 657–681.
- [20] G. GIACOMIN AND J.L. LEBOWITZ, *Phase segregation dynamics in particle system with long range interactions II: Interchange motion*, Siam J. Appl. Math., 58 (1998), pp. 1707–1729.
- [21] D.T. GILLESPIE AND L.R. PETZOLD, *Improved leap-size selection for accelerated stochastic simulation*, J Chem. Phys. , 119 (2003), pp. 8229–8234.
- [22] J. GULLINGSRUD, R. BRAUN, AND K. SCHULTEN, *Reconstructing potentials of mean force through time series analysis of steered molecular dynamics simulations*, J. Comp. Phys., 151 (1999), pp. 190–211.

- [23] P.R. HALMOS AND L.J. SAVAGE, *Application of the Radon-Nikodym theorem to the theory of sufficient statistics*, The Annals of Mathematical Statistics, 20 (1949) , pp. 225–241.
- [24] J.D. HAMILTON, *Time Series Analysis*, Princeton University Press, 1994.
- [25] C. HARDIN, M.P. EASTWOOD, M. PRENTISS, Z. LUTHEY-SCHULTEN, AND P.G. WOLYNES, *Folding funnels: The key to robust protein structure prediction*, J. Comp. Chem., 23 (2002), pp. 138–146.
- [26] Y. HONG AND H. LI, *Nonparametric specification testing for continuous-time models with applications to term structure of interest rates*, The Review of Financial Studies, 18 (2005), pp. 37–84.
- [27] G. HUMMER AND I.G. KEVREKIDIS, *Coarse molecular dynamics of a peptide fragment: Free energy, kinetics, and long-time dynamics computations*, J Chem. Phys., 118 (2003), pp. 10762–10773.
- [28] P. JEGANATHAN, *Some aspects of asymptotic theory with applications to time series models*, Econometric Theory, 11 (1995), pp. 818–887.
- [29] I.G. KEVREKIDIS, C.W. GEAR, AND G. HUMMER, *Equation-free: The computer-aided analysis of complex multiscale systems*, AIChE Journal, 50 (2004), pp. 1346–1355 .
- [30] P. KLOEDEN AND E. PLATEN, *Numerical Solution of Stochastic Differential Equations*, Springer-Verlag, 1992.
- [31] D.I. KOPELEVICH, A.Z. PANAGIOTOPOULOS, AND I.G. KEVREKIDIS, *Coarse-grained kinetic computations for rare events: Application to micelle formation*, J Chem. Phys. , 122 (2005), p. 044908.
- [32] S. KULLBACK AND R.A. LEIBLER, *On information and sufficiency*, The Annals of Mathematical Statistics, 22 (1951), pp. 79–86.
- [33] R. KUPFERMAN AND A. M. STUART, *Fitting SDE models to nonlinear Kac-Zwanzig heat bath models*, Physica D., 199 (2004), pp. 279–316.
- [34] E. LA NAVE, F. SCIORTINO, P. TARTAGLIA, M.S. SHELL, AND P. G. DEBENEDETTI, *Test of non-equilibrium thermodynamics in glassy systems: The soft-sphere case*, Phys Rev. E, 68 (2003), p. 032103.
- [35] L. LE CAM, *Locally asymptotically normal families of distributions*, University of California Publications in Statistics, 3 (1960), pp. 37–98.
- [36] ———, *Asymptotic Methods in Statistical Decision Theory*, Springer-Verlag, New York, 1986.
- [37] ———, *Maximum likelihood: An introduction*, International Statistical Review, 58 (1990), pp. 153–172.
- [38] L. LE CAM AND G. L. YANG, *Asymptotics in Statistics: Some Basic Concepts*, Springer-Verlag, 2000.
- [39] E.L. LEHMANN, *Testing statistical hypotheses*, New York, 1959.
- [40] H.W. LO, *Maximum likelihood estimation of generalized Ito processes with discretely sampled data*, Econometric Theory, 4 (1988), pp. 231–247.
- [41] A. MAKEEV, D. MAROUDAS, AND I. G. KEVREKIDIS, *Coarse stability and bifurcation analysis using stochastic simulators: Kinetic Monte Carlo examples*, J. Chem. Phys., 116 (2002), pp. 10083–10091.
- [42] P.W. MILLAR, *The minimax principle in asymptotic decision theory*, Lecture Notes in Math., (1983), pp. 75–265.
- [43] D. NUALART, *The Malliavin Calculus and Related Topics*, Springer, New York, 1995.
- [44] A.R. PEDERSEN, *A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations*, Scandinavian J. of Statistics, 22 (1995), pp. 55–71.
- [45] H. RIEDER, *Robust asymptotic statistics*, Springer-Verlag, New York, 1994.
- [46] O. RUNBORG, J. KRISHNAN, AND I. G. KEVREKIDIS, *Bifurcation analysis of nonlinear reaction-diffusion problems using wavelet-based reduction techniques*, Comp. Chem. Eng., 28 (2004), pp. 557–574.
- [47] E. SCHUTZ, N. HARTMANN, Y. KEVREKIDIS, AND R. IMBIHL, *Microchemical engineering of catalytic reactions*, Catalysis Letters, 54 (1998), pp. 181–186.
- [48] A. VAN DER VAART, *Asymptotic Statistics*, Cambridge University Press, 1998.
- [49] G. WAHBA, *Spline models for observational data*, SIAM, Philadelphia, 1990.
- [50] H. WHITE, *Maximum likelihood estimation of misspecified models*, Econometrica, 50 (1982), pp. 1–25.