

Estimation of Dynamic Models with Nonparametric Simulated Maximum Likelihood*

Dennis Kristensen[†] Yongseok Shin[‡]

June 2011

Abstract

We propose an easy-to-implement simulated maximum likelihood estimator for dynamic models where no closed-form representation of the likelihood function is available. Our method can handle any simulable model without latent dynamics. Using simulated observations, we nonparametrically estimate the unknown density by kernel methods, and then construct a likelihood function that can be maximized. We prove that this nonparametric simulated maximum likelihood (NPSML) estimator is consistent and asymptotically efficient. The higher-order impact of simulations and kernel smoothing on the resulting estimator is also analyzed; in particular, it is shown that the NPSML does not suffer from the usual curse of dimensionality associated with kernel estimators. A simulation study shows good performance of the method when employed in the estimation of jump-diffusion models.

*We thank the seminar participants at Berkeley, BU, Brown, Columbia, LSE, NYU, Rice, and Stanford for many useful comments. We also thank the referees who offered exceptionally thorough and helpful comments. Kyu-Chul Jung provided excellent research assistance. Kristensen gratefully acknowledges the financial support of the National Science Foundation (SES-0961596) and of the Danish Research Foundation (through a grant to CREATES).

[†]Department of Economics, Columbia University and CREATES, Aarhus University (e-mail: dk2313@columbia.edu).

[‡]Department of Economics, Washington University in St. Louis and Federal Reserve Bank of St. Louis (e-mail: yshin@wustl.edu).

1 Introduction

We propose a simulated maximum likelihood estimator for dynamic models based on nonparametric kernel methods. Our method is designed for models where no closed-form representation of the likelihood function is available. Our method can handle any simulable model without latent dynamics. For any given parameter value, conditioning on available past information, we draw N i.i.d. simulated observations from the model. We then use these simulated observations to nonparametrically estimate the conditional density—unknown in closed form—by kernel methods. The kernel estimate converges to the true conditional density as N goes to infinity, enabling us to approximate the true density arbitrarily well with a sufficiently large N . We then construct the likelihood and search over the parameter space to obtain a maximum likelihood estimator—nonparametric simulated maximum likelihood estimator (NPSMLE).

NPSML was introduced by Fermanian and Salanié (2004), who obtained theoretical results only for static models. In this paper, we generalize their method to dynamic models, including nonstationary and time-inhomogeneous processes. We give general conditions for the NPSMLE to be consistent and have the same asymptotic distribution as the infeasible maximum likelihood estimator (MLE). For the stationary case, we also analyze the impact of simulations on the bias and variance of the NPSMLE. In particular, we show that the estimator does not suffer from the curse of dimensionality despite the use of kernel smoothers. Finally, we show that the theoretical results remain valid even if only simulations from an approximate model are available.

NPSML can be used for estimating general classes of models, such as structural Markov decision processes and discretely-sampled diffusions. As for Markov decision processes, the transition density of endogenous state variables embodies an optimal policy function of a dynamic programming problem, and hence does not typically have a closed-form representation (Rust, 1994; Doraszelski and Pakes, 2007). However, we can closely approximate the optimal policy function numerically, and simulate observations from the model for NPSML. Similarly, as for the estimation of continuous-time stochastic models with discretely-sampled data, the transition densities are well-defined, but only in few special cases can we derive closed-form expressions for them. Again, a large class of continuous-time processes, including jump-diffusions, can be approximated with various discretization schemes to a given level of precision, and we can simulate observations from the model which are then used for NPSML. Indeed, we investigate the performance of NPSML when applied to jump-diffusion models with particular attention to the impact of number of simulations and bandwidth. We find that NPSML performs well even for a moderate number of simulations and that it is quite robust to the choice of bandwidth.

For the classes of models that NPSML addresses, there are two categories of existing approaches. The first is based on moment matching, and includes simulated methods of moments (Lee and Ingram, 1991; Duffie and Singleton, 1993; Creel and Kristensen, 2009), indirect inference (Gouriéroux et al., 1993; Smith, 1993; Creel and Kristensen, 2011), and efficient methods of moments (Gallant and Tauchen, 1996). These are all general-purpose methods, but cannot attain asymptotic

efficiency—even for models that are Markov in observables—unless the true score is encompassed by the target moments (Tauchen, 1997). More recently, Carrasco et al. (2007) and Altissimo and Mele (2009) developed general-purpose estimators based on matching a continuum of moments that are asymptotically as efficient as MLEs for fully observed systems. One attractive feature of NPSML—which it shares with Carrasco et al. (2007) and Altissimo and Mele (2009)—is that asymptotic efficiency is attained without having to judiciously choose an auxiliary model. For NPSML, the researcher has to choose a kernel and a bandwidth for the nonparametric estimation of transition densities. However, there exist many data-driven methods that guide the researcher in this regard such that our method can be fully automated to yield full efficiency (Jones et al., 1996). Another advantage is that, unlike most of the above methods, NPSML can handle nonstationary and time-inhomogeneous dynamics.

The approaches in the second category approximate the likelihood function itself, and hence is more closely related to NPSML. Examples of this approach include the simulated likelihood method (Lee, 1995) and the method of simulated scores (Hajivassiliou and McFadden, 1998), both of which are designed for limited dependent variable models. Another set of examples are various maximum likelihood methods for discretely sampled diffusions (Pedersen, 1995a,b; Sandmann and Koopman, 1998; Elerian et al., 2001; Aït-Sahalia, 2002, 2008; Brandt and Santa-Clara, 2002). While all these methods result in asymptotically efficient estimators, they are designed only for specific classes of models—i.e. limited dependent variable models or diffusions, and cannot be adapted easily to other classes of models. NPSML is for general purposes in both theoretical and practical senses. Theoretically, we establish its asymptotic properties under fairly weak regularity conditions allowing for a wide range of different models. At the practical level, when the model specification changes, only the part of the computer code that simulates observations needs to be modified, leaving other parts (e.g. kernel estimation of conditional density or numerical maximization of likelihood) unchanged.

The basic implementation of our method requires that it is possible to simulate the current variables of the model conditioning on finitely-many past observations. This excludes models with latent dynamics since these cannot be simulated one step at a time. However, our method can be modified to handle latent dynamics, but this modified version will however not obtain full efficiency. Extensions of our method that obtain full efficiency are worked out in a companion paper (Brownlees et al., 2011) building on the main results obtained here.

The rest of the paper is organized as follows. In the next section, we set up our framework to present the simulated conditional density and the associated NPSMLE. In Section 3, we derive the asymptotic properties of the NPSMLE under regularity conditions. Section 4 provides a detailed description on implementing NPSML with numerical examples, and Section 5 concludes.

2 Nonparametric Simulated Maximum Likelihood

2.1 Construction of NPSMLE

Suppose that we have T observations, $\{(y_t, x_t)\}_{t=1}^T$, $y_t \in \mathbb{R}^k$ and $x_t \in \mathcal{X}_t$. The space \mathcal{X}_t can be time-varying. We assume that the data is generated by a fully parametric model:

$$y_t = g_t(x_t, \varepsilon_t; \theta), \quad t = 1, \dots, T, \quad (1)$$

where $\theta \in \Theta \subseteq \mathbb{R}^d$ is an unknown parameter vector, and ε_t is an i.i.d. sequence with known distribution F_ε and independent of x_t and ε_t . Without loss of generality, assume that F_ε is known and does not depend on t and θ . Our setting accommodates Markov models where $x_t \equiv y_{t-1}$, such that $\{y_t\}$ is a (possibly time-inhomogeneous) Markov process. In this case (1) is a fully-specified model. However, we allow x_t to contain other (exogenous) variables than lagged y_t , in which case (1) is only a partially-specified model. Also, we allow the processes (y_t, x_t) to be nonstationary, for example due to unit-root-type behavior or deterministic time trends.

The model is assumed to have an associated conditional density $p_t(y|x; \theta)$. That is,

$$P(y_t \in A | x_t = x) = \int_A p_t(y|x; \theta) dy, \quad t = 1, \dots, T,$$

for any Borel set $A \subseteq \mathbb{R}^k$. A natural estimator of θ is then the maximizer of the conditional log-likelihood:

$$\tilde{\theta} = \arg \max_{\theta \in \Theta} L_T(\theta), \quad L_T(\theta) = \sum_{t=1}^T \log p_t(y_t | x_t; \theta).$$

If the model (1) is fully specified, i.e. x_t only contains lagged y_t , then this is the full likelihood of the model conditional on the starting value. If on the other hand, x_t contains other variables than lagged y_t , $L_T(\theta)$ is a partial likelihood.

Suppose now that $p_t(y|x; \theta)$ does not have a closed-form representation, and thus the maximum likelihood estimation of θ is not feasible. In terms of the model (1), this occurs when either the inverse of $g_t(x_t, \varepsilon_t; \theta)$ w.r.t. ε_t does not exist, or when the inverse does not have a closed-form expression.¹ Such a situation may arise, for example, when the function g involves a solution to a dynamic programming problem, or when we are dealing with discretely-sampled jump-diffusions. In such cases, although $p_t(y|x; \theta)$ is not available in closed form, we are still able to generate simulated observations from the model: A solution to a dynamic programming problem can be represented numerically, and a jump-diffusion can be approximated by various discretization schemes up to a given level of precision.

We here propose a general method to obtain a simulated conditional density, which in turn will be used to obtain a simulated version of the MLE. For any given $1 \leq t \leq T$, $y_t \in \mathbb{R}^k$, $x_t \in \mathcal{X}_t$, and

¹If the inverse has a closed-form expression, we have $p_t(y|x; \theta) = p_\varepsilon(g_t^{-1}(y, x; \theta)) \left| \frac{\partial g_t^{-1}(y, x; \theta)}{\partial y} \right|$, and the likelihood is easily evaluated.

$\theta \in \Theta$, we wish to compute a simulated version of $p_t(y_t|x_t; \theta)$. To this end, we first generate N i.i.d. draws from F_ε , $\{\varepsilon_i\}_{i=1}^N$, through a random number generator, and use these to compute

$$Y_{t,i}^\theta = g_t(x_t, \varepsilon_i; \theta), \quad i = 1, \dots, N.$$

By construction, the N simulated i.i.d. random variables, $\{Y_{t,i}^\theta\}_{i=1}^N$, follow the target distribution: $Y_{t,i}^\theta \sim p_t(\cdot|x_t; \theta)$, $i = 1, \dots, N$. They can therefore be used to estimate $p_t(y|x; \theta)$ with kernel methods. Define:

$$\hat{p}_t(y_t|x_t; \theta) = \frac{1}{N} \sum_{i=1}^N K_h(Y_{t,i}^\theta - y_t), \quad (2)$$

where $K_h(v) = K(v/h)/h^k$, $K : \mathbb{R}^k \mapsto \mathbb{R}$ is a kernel, and $h > 0$ a bandwidth.² Under regularity conditions on p_t and K , we obtain:

$$\hat{p}_t(y_t|x_t; \theta) = p_t(y_t|x_t; \theta) + O_P(1/\sqrt{Nh^k}) + O_P(h^2), \quad N \rightarrow \infty,$$

where the remainder terms are $o_P(1)$ if $h \rightarrow 0$ and $Nh^k \rightarrow \infty$.

Once (2) has been used to obtain the simulated conditional density, we can now construct the following simulated MLE of θ_0 :

$$\hat{\theta} = \arg \max_{\theta \in \Theta} \hat{L}_T(\theta), \quad \hat{L}_T(\theta) = \sum_{t=1}^T \log \hat{p}_t(y_t|x_t; \theta).$$

When searching for $\hat{\theta}$ through numerical optimization, we use the same draws for all values of θ . We may also use the same batch of draws from $F_\varepsilon(\cdot)$, $\{\varepsilon_i\}_{i=1}^N$, across different values of t and x . Numerical optimization is facilitated if $\hat{L}_T(\theta)$ is continuous and differentiable in θ . With (2), if K and $\theta \mapsto g_t(x, \varepsilon; \theta)$ are $r \geq 0$ times continuously differentiable, then $\hat{L}_T(\theta)$ has the same property. This follows from the chain rule and the fact that we use the same random draws $\{\varepsilon_i\}_{i=1}^N$ for all values of θ .

Since $\hat{p}_t(y_t|x_t; \theta) \xrightarrow{P} p_t(y_t|x_t; \theta)$, $\hat{L}_T(\theta) \xrightarrow{P} L_T(\theta)$ as $N \rightarrow \infty$ for a given $T \geq 1$ under regularity conditions. The main theoretical results of this paper demonstrate that $\hat{\theta}$ inherits the properties of the infeasible MLE, $\tilde{\theta}$, as $T, N \rightarrow \infty$, under suitable conditions.

The precision of $\hat{\theta}$ relative to $\tilde{\theta}$ clearly depends on the quality of the approximation of $p_t(y|x; \theta)$ by $\hat{p}_t(y|x; \theta)$. Let us note the following important points concerning the impact of the simulated density. Firstly, because we use i.i.d. draws, the density estimator is not affected by the dependence structure in the observed data. In particular, our estimator works whether the observed data are i.i.d. or nonstationary. Secondly, the simulated density, $\hat{p}_t(y|x; \theta)$, suffers from the usual curse of dimensionality for kernel density estimators with its variance being of order $1/(Nh^k)$. The curse of dimensionality only depends on $k \equiv \dim(y_t)$ here since we do not smooth over x_t , and so the dimension of x_t is irrelevant in itself. Still one could be concerned that for high-dimensional

²Here and in the following, we will use K to denote a generic kernel.

models (i.e. large k), a much larger number of simulations must be used to control the variance component of the resulting estimator relative to, say, standard simulated method of moments (where the variance is of order $1/N$). However, as we demonstrate in Section 3, this is fortunately not the case: The summation of the log-densities in the computation of $\hat{L}_T(\theta)$ functions as an additional smoothing device such that the additional variance of $\hat{L}_T(\theta)$ due to simulations recovers the standard parametric rate $1/N$. This is a well-known phenomenon in the literature on two-step semiparametric estimators: They can obtain parametric rates despite the fact that they depend on first-step nonparametric estimators—e.g. Kristensen (2010).

On the other hand, a disadvantage of our estimator is that for a fixed N and $h > 0$, the simulated log-likelihood function is a biased estimate of the actual one. First, to obtain consistency, we will have to let $N \rightarrow \infty$ which is a feature shared by most non-linear simulation-based likelihood methods.³ In addition, we have to let $h \rightarrow 0$ which is an additional complication relative to other simulation-based estimators where no such nuisance parameter appears. In particular, one has to choose h for a given sample and simulation size. However, if one is willing to make a stronger assumption about the identification of the model, the choice of the bandwidth may be less important. For example, in the stationary case, the standard identification assumption states that $\mathbb{E}[\log p(y_t|x_t; \theta)] < \mathbb{E}[\log p(y_t|x_t; \theta_0)]$ for $\theta \neq \theta_0$. A stronger identification condition implying the former is

$$\mathbb{E} \left[\log \left(\int K(v)p(y_t + hv|x_t; \theta)dv \right) \right] < \mathbb{E} \left[\log \left(\int K(v)p(y_t + hv|x_t; \theta_0)dv \right) \right], \quad \theta \neq \theta_0,$$

for all $0 \leq h \leq \bar{h}$ for some $\bar{h} > 0$.⁴ Under the latter identification condition, one can show consistency of our estimator for any fixed $0 < h \leq \bar{h}$ as $N \rightarrow \infty$. A similar identification condition can be found in Altissimo and Mele (2009). Still, for a fixed $h > 0$, the resulting estimator will no longer enjoy full efficiency. To obtain this, one has to let $h \rightarrow 0$. Moreover, the above argument assumes knowledge of the threshold $\bar{h} > 0$ for a given model and so in practice the bandwidth selection problem still remains. However, it suggests that one can still identify parameters in large samples when a given $h > 0$ is chosen and that the NPSMLE will be fairly robust to the choice of h . This is supported by our simulation study which shows that the NPSML performs well within a fairly broad range of bandwidth choices. Still a careful choice of the bandwidth will in general lead to better performance of the estimator.

While we here focus on the kernel estimator, one can use other nonparametric density estimators as well. Examples are the semi-nonparametric estimators of Wahba (1981), Phillips (1983), Gallant and Nychka (1987), and Fenton and Gallant (1996); the log-spline estimator of Stone (1990); and

³See Lee and Song (2009) for an exception.

⁴This follows from the following inequality:

$$\begin{aligned} \mathbb{E}[\log p(y_t|x_t; \theta)] &= \lim_{h \rightarrow 0} \mathbb{E} \left[\log \left(\int K(v)p(y_t + hv|x_t; \theta)dv \right) \right] < \lim_{h \rightarrow 0} \mathbb{E} \left[\log \left(\int K(v)p(y_t + hv|x_t; \theta_0)dv \right) \right] \\ &= \mathbb{E}[\log p(y_t|x_t; \theta_0)]. \end{aligned}$$

the wavelet estimator of Donoho et al. (1996).

Example: Discretely-Observed Jump Diffusion Consider an \mathbb{R}^k -dimensional continuous-time stochastic process $\{y_t : t \geq 0\}$ that solves the following stochastic differential equation:

$$dy_t = \mu(t, y_t; \theta) dt + \Sigma(t, y_t; \theta) dW_t + J_t dQ_t. \quad (3)$$

The model contains both continuous and jump components. $W_t \in \mathbb{R}^l$ is a standard Brownian motion, while Q_t is an independent pure jump process with stochastic intensity $\lambda(t, y_t; \theta)$ and jump size 1. The functions $\mu : [0, \infty) \times \mathbb{R}^k \times \Theta \mapsto \mathbb{R}^k$ and $\Sigma : [0, \infty) \times \mathbb{R}^k \times \Theta \mapsto \mathbb{R}^{k \times l}$ is the drift and the diffusion term respectively, while J_t measures the jump sizes and has density $v(t, y_t; \theta)$.

Such jump diffusions are widely used in finance to model the dynamics of stock prices, interest rates, exchange rates and so on (Sundaresan, 2000). Suppose we have a sample y_1, \dots, y_T —without loss of generality, we normalize the time interval between observations to 1—and wish to estimate θ by maximum likelihood. Although under regularity conditions (Lo, 1988) the transition density $p_t(y|x; \theta)$ satisfying $P(y_{t+1} \in A | y_t = x) = \int_A p_t(y|x; \theta) dy$ is well-defined, it cannot in general be written in closed form which in turn complicates estimation.⁵ However, discretization schemes (Kloeden and Platen, 1992; Bruti-Liberati and Platen, 2007) can be used to simulate observations from the model for any given level of accuracy, enabling NPSML. We re-visit this example in Section 4 where we provide a detailed description of implementing NPSML in practice.

2.2 Extensions and Alternative Schemes

Latent Dynamics Our method can be modified to handle dynamic latent variables: Suppose y_t is generated from

$$[y_t, w_t] = g(y_{t-1}, w_{t-1}, \varepsilon_t; \theta),$$

where w_t is unobserved/latent and $\varepsilon_t \sim \text{i.i.d. } F_\varepsilon$. The full likelihood function will require computation of conditional densities on the form $p(y_t | y_{t-1}, y_{t-2}, \dots, y_0; \theta)$ which in general is complicated due to the expanding information set; see e.g. Brownlees et al. (2011). We can however construct a simulated version of the following "limited information" likelihood (LIL) given by $L_T(\theta) = \sum_{t=1}^T \log p(y_t | x_t; \theta)$ where x_t is a set of conditioning variables chosen by the econometricians, say, $x_t = (y_{t-1}, \dots, y_{t-m})$ for some $m \geq 1$. There will be an efficiency loss from estimating θ using this LIL relative to the full likelihood, but the LIL is a lot easier to implement: First simulate a (long) trajectory $\{Y_t^\theta\}_{t=1}^{\tilde{N}}$ by

$$\left[Y_t^\theta, W_t^\theta \right] = g\left(Y_{t-1}^\theta, W_{t-1}^\theta, \varepsilon_t; \theta \right), \quad t = 1, \dots, \tilde{N},$$

⁵Schaumburg (2001) and Yu (2007), building on the approach of Ait-Sahalia (2002), use analytic expansions to approximate the transition density for univariate and multivariate jump diffusions, respectively. Their asymptotic result requires that the sampling interval shrink to zero. The simulated MLE of Pedersen (1995a,b) or Brandt and Santa-Clara (2002) need to be substantially modified before they can be applied to Lévy processes.

where $\{\varepsilon_t\}_{t=1}^{\tilde{N}}$ are i.i.d. draws from F_ε . We can then use these simulations to construct a simulated version of $p(y_t|x_t; \theta)$ by the following kernel estimator of the conditional density,

$$\check{p}(y|x; \theta) = \frac{\sum_{t=1}^{\tilde{N}} K_h(Y_t^\theta - y)K_h(X_t^\theta - x)}{\sum_{t=1}^{\tilde{N}} K_h(X_t^\theta - x)}. \quad (4)$$

where $X_t^\theta = (Y_{t-1}^\theta, \dots, Y_{t-m}^\theta)$. Similar ideas were utilized in Altissimo and Mele (2009) and Creel and Kristensen (2009).

A disadvantage of the above method is that the convergence of \check{p} relative to \hat{p} will be slower due to (i) the dimension of (Y_t^θ, X_t^θ) can potentially be quite large and (ii) the simulated variables are now dependent. So one will have to choose a larger \tilde{N} for the simulated conditional density in (4) relative to the one in (2).

To handle (ii), one will typically have to assume a stationary solution to the dynamic system under consideration, and either have to start the simulation from the stationary distribution, or assume that the simulated process converges towards the stationary distribution at a suitable rate. For the latter to hold, one will need to impose some form of mixing condition on the process, as in Altissimo and Mele (2009) and Creel and Kristensen (2009). Then a large value of \tilde{N} is needed to ensure that the simulated process is sufficiently close to its stationary distribution—that is, one has to allow for a burn-in.

The estimator in (4) may work under nonstationarity as well. Recently, a number of papers have considered kernel estimation of nonstationary Markov processes. The kernel estimator proves to be consistent and asymptotically mixed-normally distributed when the Markov process is recurrent (Karlsen and Tjøstheim, 2001; Bandi and Phillips, 2003). However, the convergence rate will be path-dependent and relatively slow.

In the remainder of this paper we focus on (2). The properties of (4) can be obtained by following the same strategy of proof as the one we employ for (2). The only difference is that, to obtain $\check{p} \xrightarrow{P} p$ in the sup-norm one has to take into account the dependence of the simulated values. This can be done along the lines of Creel and Kristensen (2009) where kernel regressions and simulations are combined to compute GMM estimators for dynamic latent variable models.

Discrete Random Variables Discrete random variables can be accommodated within our framework. Suppose y_t contains both continuous and discrete random variables. For example, $y_t = (y_{1t}, y_{2t}) \in \mathbb{R}^{k+l}$ where $y_{1t} \in \mathbb{R}^k$ is a continuous random variable while $y_{2t} \in \mathcal{Y}_2 \subset \mathbb{R}^l$ is a random variable with (potentially infinite number of) discrete outcomes, $\mathcal{Y}_2 = \{y_{2,1}, y_{2,2}, \dots\}$. We could then use a mixed kernel to estimate $p_t(y|x)$. For given simulated observations $Y_{t,i}^\theta = (Y_{1t,i}^\theta, Y_{2t,i}^\theta)$, $i = 1, \dots, N$:

$$\hat{p}_t(y_1, y_2|x; \theta) = \frac{1}{N} \sum_{i=1}^N K_h(Y_{1t,i}^\theta - y_1) \mathbb{I}\{Y_{2t,i}^\theta = y_2\}, \quad (y_{1t}, y_{2t}) \in \mathbb{R}^{k+l}, \quad (5)$$

where $\mathbb{I}\{\cdot\}$ is the indicator function and $K : \mathbb{R}^k \mapsto \mathbb{R}$ is the kernel from before. However, the resulting simulated log-likelihood will be discontinuous and optimization may be difficult. One

could replace the indicator function used for the discrete component with a smoother. Examples of smoothers can be found in Cai et al. (2001) and Li and Racine (2007, ch. 2). These will increase bias but reduce variance of the estimator, and at the same time lead to a continuous function. However, in general, $\theta \mapsto Y_{2t,i}^\theta$ itself will not be continuous so either way, with a discrete component, $\hat{L}_T(\theta)$ based on (5) is no longer continuous w.r.t. θ .

Instead, we will here assume that there exists a function $Y_{2t,i}^\theta = g_2(y_{2t}, x_t, \varepsilon_i; \theta)$ that is smooth in θ such that

$$\mathbb{E} \left[Y_{2t,i}^\theta | Y_{1t,i}^\theta = y_{1t} \right] = p_t(y_{2t} | y_{1t}, x_t; \theta). \quad (6)$$

Thus, $Y_{2t,i}^\theta$ now denotes a simulated value of the associated density, and not the outcome of the dependent variable. We then propose to estimate the joint density by

$$\hat{p}_t(y_{1t}, y_{2t} | x_t; \theta) = \frac{1}{N} \sum_{i=1}^N K_h(Y_{1t,i}^\theta - y_{1t}) Y_{2t,i}^\theta. \quad (7)$$

To motivate the above assumption and the resulting estimator, we first note that a discrete random variable can always be represented as $y_{2,t} = D(z_t)$ for some continuous variables $z_t \in \mathbb{R}^m$ and some function $D : \mathbb{R}^m \mapsto \mathcal{Y}_2$ which we, for the sake of the argument, assume does not depend on (t, x, θ) . For example, most limited dependent variables can be written in this form, c.f. Manrique and Shephard (1998) and the references therein. We assume that z_t satisfies $z_t = g_Z(x_t, \varepsilon_t; \theta)$ for some function g_Z that can be written in closed form, and has associated conditional density $p_{z_t|x_t}(z|x)$. Clearly, $p_t(y_2|x) = P(y_{2t} = y_2 | x_t = x)$ satisfies

$$p_t(y_2|x) = P(z_t \in D^{-1}(y_2) | x_t = x) = \int_{D^{-1}(y_2)} p_{z_t|x_t}(z|x) dz.$$

The last integral is equal to $\int_{\mathbb{R}^m} \frac{p_t(z|x)}{p_D(z|y_2)} p_D(z|y_2) dz$ for any density $p_D(z|y_2)$ with support $D^{-1}(y_2)$. If $p_{z_t|x_t}(z|x)$ is known in closed form, this integral can then be simulated by

$$\hat{p}_t(y_{2t}|x_t) = \frac{1}{N} \sum_{i=1}^N \mathcal{K}^{(1)}(\tilde{Z}_{t,i}^\theta | y_{2t}, x_t), \quad \mathcal{K}^{(1)}(\tilde{Z}_{t,i}^\theta | y_{2t}, x_t) = \frac{p_{z_t|x_t}(\tilde{Z}_{t,i}^\theta | x_t)}{p_D(\tilde{Z}_{t,i}^\theta | y_{2t})}, \quad (8)$$

where $\tilde{Z}_{t,i}^\theta \stackrel{iid}{\sim} p_D(z|y_{2t})$, as is standard in the estimation of limited dependent variable models.

If $p_{z_t|x_t}(z|x)$ cannot be written in closed form, we propose to use $\hat{p}_{z_t|x_t}(z_t|x_t) = \frac{1}{N} \sum_{i=1}^N K_b(Z_{t,i}^\theta - z)$, where $Z_{t,i}^\theta = g_Z(x_t, \varepsilon_i; \theta)$ and $b > 0$ is another bandwidth. If $\int_{D^{-1}(y_2)} K_b(Z_{t,i}^\theta - z) dz$ can be written in closed form, we follow Fermanian and Salanié (2004, pp. 709–710 and 724–725) and use:

$$\hat{p}_t(y_{2t}|x_t) = \frac{1}{N} \sum_{i=1}^N \mathcal{K}_b^{(2)}(Z_{t,i}^\theta, y_{2t}), \quad \mathcal{K}_b^{(2)}(Z_{t,i}^\theta, y_{2t}) = \int_{D^{-1}(y_{2t})} K_b(Z_{t,i}^\theta - z) dz. \quad (9)$$

If this is not the case, we can use

$$\hat{p}_t(y_{2t}|x_t) = \sum_{i=1}^N \hat{\mathcal{K}}_b^{(2)}(Z_{t,i}^\theta, y_{2t}), \quad \hat{\mathcal{K}}_b^{(2)}(Z_{t,i}^\theta, y_{2t}) = \frac{1}{N} \sum_{j=1}^N \frac{K_b(Z_{t,i}^\theta - \tilde{Z}_{t,j}^\theta)}{p_D(\tilde{Z}_{t,j}^\theta | y_{2t})}. \quad (10)$$

In all three cases, we can write the resulting simulated joint density in equation (7) by choosing $Y_{2t,i}^\theta = \mathcal{K}^{(1)}(\tilde{Z}_{t,i}^\theta, y_{2t}|x_t)$, $Y_{2t,i}^\theta = \mathcal{K}_b^{(2)}(Z_{t,i}^\theta, y_{2t})$ and $Y_{2t,i}^\theta = \hat{\mathcal{K}}_b^{(2)}(Z_{t,i}^\theta, y_{2t})$, respectively. Here, $\theta \mapsto Y_{2t,i}^\theta$ is smooth with a bias that disappears as $b \rightarrow 0$ and variance that is bounded in b . Thus, the order of the variance of $\hat{L}_T(\theta)$ is not affected by any added discrete variables, and the curse of dimensionality remains of order $k = \dim(y_{1t})$.

Quasi Maximum Likelihood Estimation The use of our approximation method is not limited to actual MLEs. In many situations, one can define quasi- or pseudo-likelihood which, even though it is not the true likelihood, identifies the parameters of the true model. One obvious example of this is the standard regression model, where the MLE based on Gaussian errors (i.e. the least-squares estimator) proves to be robust to deviations from the normality assumption. Another example is estimation of (G)ARCH models using quasi-maximum likelihood—e.g. Lee and Hansen (1994). These are cases where the quasi-likelihood can be written explicitly. If one cannot find explicit expressions of the quasi-likelihood, one can instead employ our estimator, simulating from the quasi-model: Suppose for example that data has been generated by the model (1), but the data-generating distribution of the errors is unknown. We could then choose a suitable distribution F_ε , draw $\{\varepsilon_i\}_{i=1}^N$ from F_ε and then proceed as in Section 2.1. The resulting estimator would no longer be a simulated MLE but rather a simulated QMLE. In this setting, the asymptotic distribution should be adjusted to accommodate the fact that we are not using the true likelihood to estimate the parameters. This obviously extends to the case of misspecified models as in White (1984).

The above procedure is one example of how our simulation method can be applied to non- and semiparametric estimation problems where an infinite-dimensional component of the model is unknown. Another example is the situation where data has been generated by the model (1) with known distribution F_ε , but now $\theta = (\alpha, \gamma)$ where α and γ are finite- and infinite-dimensional parameters respectively. An application of our method in this setting can be found in Kristensen (2010) where γ is a density. Again, our asymptotic results have to be adjusted to allow for θ to contain infinite-dimensional parameters.

3 Asymptotic Properties of NPSMLE

Given the convergence of the simulated conditional density towards the true one, we expect that the NPSMLE $\hat{\theta}$ based on the simulated kernel density estimator will have the same asymptotic properties as the infeasible MLE $\tilde{\theta}$ for a suitably chosen sequence $N = N(T)$ and $h = h(N)$. We give two sets of results: The first establishes that $\hat{\theta}$ is first-order asymptotic equivalent to $\tilde{\theta}$ under general conditions, allowing for nonstationarity. Under additional assumptions, including stationarity, we derive expressions of the leading bias and variance components of $\hat{\theta}$ relative to the actual MLE due to simulations and kernel smoothing, and give results for the higher-order asymptotic properties of $\hat{\theta}$.

We allow for a mixed discrete and continuous distribution of the response variable, and

write $y_t = (y_{1t}, y_{2t}) \in \mathcal{Y}_1 \times \mathcal{Y}_2$, where $\mathcal{Y}_1 \subseteq \mathbb{R}^k$ and $\mathcal{Y}_2 = \{y_{2,1}, y_{2,2}, \dots\} \subset \mathbb{R}^l$. Here, y_{1t} has a continuous distribution, while y_{2t} is discrete. The joint distribution can be written as $p_t(y_1, y_2|x; \theta) = p_t(y_2|y_1, x; \theta)p_t(y_1|x; \theta)$ where $p_t(y_{2,i}|y_1, x; \theta)$ are conditional probabilities satisfying $\sum_{i=1}^l p_t(y_{2,i}|y_1, x; \theta) = 1$, while $p_t(y_1|x; \theta)$ is a conditional density w.r.t. the Lebesgue measure. Also, let $p_t(y_{2,i}|x; \theta)$ denote the conditional probabilities of $y_{2t}|x_t = x$.

The asymptotics are derived for the kernel estimator given in equation (7) where

$$Y_{1t,i}^\theta := g_{1,t}(x_t, \varepsilon_i; \theta), \quad (11)$$

$$Y_{2t,i}^\theta := g_{2,t}(y_{2t}, x_t, \varepsilon_i; \theta), \quad (12)$$

for $i = 1, \dots, N$ and $t = 1, \dots, T$, where $\{\varepsilon_i\}_{i=1}^N$ are i.i.d. draws from F_ε , such that equation (6) holds. Recall that $Y_{2t,i}^\theta$ denotes a simulated value of the associated density, and not the outcome of the dependent variable. The condition in equation (6) is met when $Y_{2t,i}^\theta = \mathcal{K}^{(1)}(\tilde{Z}_{t,i}^\theta, y_{2t}|x_t)$ with $\mathcal{K}^{(1)}$ given in equation (8), while it only holds approximately for $\mathcal{K}^{(2)}$ and $\hat{\mathcal{K}}^{(2)}$ defined in equations (9) and (10) due to biases induced by the use of kernel smoothing. We handle these two cases in Theorem 3.4 where results for approximate simulations are given.

Note that we here use the same errors to generate the simulations over time. An alternative simulation scheme would be to draw a new batch of errors for each observation x_t , $Y_{t,i}^\theta = g_t(x_t, \varepsilon_{t,i}; \theta)$, $i = 1, \dots, \tilde{N}$, such that the total number of simulations would be $\tilde{N} \times T$, $\{\varepsilon_{i,t}\}_{i=1}^{\tilde{N}}$, $t = 1, \dots, T$. Under regularity conditions, the NPSMLE based on this simulation scheme would have similar asymptotic properties as the one based on the simulations in equations (11) and (12). However, as demonstrated in Lee (1992), choosing $N = \tilde{N}T$, the variance of the NPSMLE based on equations (11) and (12) will be smaller.⁶

In order for $\hat{\theta}$ to be asymptotically equivalent to $\tilde{\theta}$, we need $\hat{p} \xrightarrow{P} p$ sufficiently fast in some suitable function norm. To establish this, we verify the general conditions for uniform rates of kernel estimators found in Kristensen (2009). These general conditions are satisfied under the following set of regularity conditions regarding the model and its associated conditional density:

A.1 The functions $(x, t, \theta) \mapsto g_{1,t}(x, \varepsilon; \theta)$ and $(x, t, \theta) \mapsto g_{2,t}(y_2, x, \varepsilon; \theta)$ are continuously differentiable for all y_2 and ε such that for some function $\Lambda(\cdot)$ and constants $\lambda_{i,j} \geq 0$, $i, j = 1, 2$,

$$\|g_{1,t}(x, \varepsilon; \theta)\| \leq \Lambda(\varepsilon) \left[1 + \|x\|^{\lambda_{1,1}} + t^{\lambda_{1,2}} \right], \quad \|g_{2,t}(y_2, x, \varepsilon; \theta)\| \leq \Lambda(\varepsilon) \left[1 + \|x\|^{\lambda_{2,1}} + t^{\lambda_{2,2}} \right],$$

and $\mathbb{E}[\Lambda(\varepsilon)^s] < \infty$ for some $s > 2$. The derivatives of g_1 and g_2 w.r.t. (x, t, θ) satisfy the same bounds.

A.2 The conditional density $p_t(y_1, y_2|x; \theta)$ is continuous w.r.t. $\theta \in \Theta$, and $r \geq 2$ times continuously differentiable w.r.t. y_1 with the r -th derivative being uniformly continuous. There exists

⁶The results of Lee (1992) are for discrete choice models, but we conjecture that his results can be extended to general simulated MLE.

constants $\bar{B}_0 > 0$ and $\lambda_{0,1}, \lambda_{0,2} \geq 0$, such that the following bounds hold uniformly over (t, y_1, y_2, x, θ) with $\bar{B}(x, t) = \bar{B}_0 \left(1 + \|x\|^{\lambda_{0,1}} + t^{\lambda_{0,2}}\right)$:

$$\sum_{|\alpha|=r} \left| \frac{\partial^r p_t(y_1, y_2|x; \theta)}{\partial y_1^\alpha} \right| \leq \bar{B}(x, t), \quad \|y_1\|^k p_t(y_1, y_2|x; \theta) \leq \bar{B}(x, t). \quad (13)$$

A.3 $\theta \mapsto g_{1,t}(x, \varepsilon; \theta)$ and $\theta \mapsto g_{2,t}(x, y_2, \varepsilon; \theta)$ are twice continuously differentiable for all t, x, ε with their derivatives satisfying the same moment conditions as g_1 and g_2 in Assumption A.1.

A.4 $\partial p_t(y|x; \theta)/(\partial\theta)$ and $\partial^2 p_t(y|x; \theta)/(\partial\theta\partial\theta')$ are $r \geq 2$ times continuously differentiable w.r.t. y_1 with bounded derivatives such that they satisfy the same bounds in equation (13) as p .

Assumptions A.1 and A.2 are used to establish uniform convergence of \hat{p} (Lemma B.1). Assumption A.1 imposes restrictions on the two data-generating functions g_1 and g_2 . The smoothness conditions are rather weak, and satisfied by most models, while the polynomial bounds imposed on the two functions can be exchanged for other bounds, but will complicate some of the conditions imposed below. Note that the moment condition in Assumption A.1 does not concern the observed process $\{(y_t, x_t)\}$, only the errors ε that we draw when simulating. If for example, $\Lambda(\varepsilon) \propto \|\varepsilon\|^q$, then the moment condition is satisfied if $\mathbb{E}[\|\varepsilon\|^{qs}] < \infty$. Thus, in this case, the moment condition only rules out models driven by fat-tailed errors. If the model is time-homogenous, $\lambda_{i,2} = 0$, $i = 1, 2$.

Assumption A.2 restricts the conditional density that we are trying to estimate. The smoothness assumptions imposed on p in A.2 in conjunction with the use of higher-order kernels (which are introduced below) controls the bias of \hat{p} . The bounds are imposed to obtain a uniform bound of the variance of \hat{p} . Again, the assumptions are quite weak and are satisfied by many models. If the model is time-homogenous, $\lambda_{0,2} = 0$.

Assumptions A.3 and A.4 will only be used when examining the higher-order impact of simulations and kernel smoothing on our estimator. These two conditions yield uniform convergence of $\partial \hat{p}_t(y|x; \theta)/\partial\theta$ and $\partial^2 \hat{p}_t(y|x; \theta)/(\partial\theta\partial\theta')$, which in turn allows us to analyze the first and second derivatives of the simulated log-likelihood (Lemma B.2).

Our conditions are slightly stronger than the ones found in Fermanian and Salanié (2004, M.1–2 and L.1–3). There, weaker bounds and smoothness conditions are imposed on the function g , while their restrictions on the conditional density are very similar to ours.

The kernel K is assumed to belong to the following class of so-called higher-order or bias-reducing kernels.

K.1 The kernel K satisfies:

1. $\sup_{u \in \mathbb{R}^k} |K(u)| < \infty$ and $\int_{\mathbb{R}^k} |K(u)| du < \infty$. There exist $C, L < \infty$ such that either (i) $K(u) = 0$ for $\|u\| > L$ and $|K(u) - K(u')| \leq C \|u - u'\|$, or (ii) $K(u)$ is differentiable with $\sup_{u \in \mathbb{R}^k} |\partial K(u)/\partial u| \leq C$. For some $a > 1$, $|\partial^{|\alpha|} K(u)/\partial u^\alpha| \leq C \|u\|^{-a}$ for $\|u\| \geq L$ and all $1 \leq |\alpha| \leq 2$.

2. $\int_{\mathbb{R}^k} K(u) du = 1$ and for some $r \geq 1$: $\int_{\mathbb{R}^k} K(u) u^\alpha du = 0$, $1 \leq |\alpha| \leq r - 1$, and $\int_{\mathbb{R}^k} K(u) \|u\|^r du < \infty$.

K.2 The first and second derivatives of K also satisfy K.1.1.

This is a broad class of kernels allowing for unbounded support. For example, the Gaussian kernel satisfies K.1 with $r = 2$. When $r > 2$, K is a so-called higher-order kernel that reduces the bias of \hat{p} and its derivatives, and thereby obtains a faster rate of convergence. The smoothness of p as measured by its number of derivatives, r , determines the degree of bias reduction. The additional assumption K.2 is used in conjunction with Assumptions A.3 and A.4 to show that the first and the second derivatives of \hat{p} w.r.t. θ also converge uniformly.

Next, we impose regularity conditions on the model to ensure that the actual MLE is asymptotically well-behaved. We first introduce the relevant terms driving the asymptotics of the MLE. We first normalize the log-likelihood by some factor $\nu_T \rightarrow \infty$:

$$L_T(\theta) = \frac{1}{\nu_T} \sum_{t=1}^T \log p_t(y_t|x_t; \theta).$$

This normalizing factor ν_T is introduced to ensure that $L_T(\theta)$ is well-behaved asymptotically and that certain functions of data are suitably bounded, c.f. C.1–C.4 below. It is only important for the theoretical derivations, and not relevant for the actual implementation of our estimator since ν_T does not depend on θ . The choice of ν_T depends on the dynamics of the model. The standard choice is $\nu_T = T$ as is, for example, the case when the model is stationary. In order to allow for non-standard behavior of the likelihood due to, for example, stochastic and deterministic trends, we do not impose this restriction though.

We also redefine the simulated version of the likelihood: In order to obtain uniform convergence of $\log \hat{p}_t(y|x; \theta)$, we need to introduce trimming of the approximate log-likelihood as is standard in the literature on semiparametric estimators. The trimmed and normalized version of the simulated log-likelihood is given as

$$\hat{L}_T(\theta) = \frac{1}{\nu_T} \sum_{t=1}^T \tau_a(\hat{p}_t(y_t|x_t; \theta)) \log \hat{p}_t(y_t|x_t; \theta),$$

where $\tau_a(\cdot)$ is continuously differentiable trimming function satisfying $\tau_a(z) = 1$ if $|z| > a$, and 0 if $|z| < a/2$, with a trimming sequence $a = a(N) \rightarrow 0$. One could here simply use the indicator function for the trimming, but then $\hat{L}_T(\theta)$ would no longer be differentiable, and differentiability is useful when using numerical optimization algorithms to solve for $\hat{\theta}$.

Assuming that $L_T(\theta)$ is three times differentiable, c.f. Assumption C.3 below, we can define:

$$S_T(\theta) = \frac{\partial L_T(\theta)}{\partial \theta} = \frac{1}{\nu_T} \sum_{t=1}^T \frac{\partial \log p_t(y_t|x_t; \theta)}{\partial \theta} \in \mathbb{R}^d,$$

$$H_T(\theta) = \frac{\partial^2 L_T(\theta)}{\partial \theta \partial \theta'} = \frac{1}{\nu_T} \sum_{t=1}^T \frac{\partial^2 \log p_t(y_t|x_t; \theta)}{\partial \theta \partial \theta'} \in \mathbb{R}^{d \times d},$$

$$G_{T,i}(\theta) = \frac{\partial^3 L_T(\theta)}{\partial \theta \partial \theta' \partial \theta_i} = \frac{1}{\nu_T} \sum_{t=1}^T \frac{\partial^3 \log p_t(y_t|x_t; \theta)}{\partial \theta \partial \theta' \partial \theta_i} \in \mathbb{R}^{d \times d}.$$

The information is then defined as:

$$i_T(\theta) = \frac{1}{\nu_T} \sum_{t=1}^T \mathbb{E} \left[\frac{\partial \log p_t(y_t|x_t; \theta)}{\partial \theta} \frac{\partial \log p_t(y_t|x_t; \theta)}{\partial \theta'} \right] = \mathbb{E} [H_T(\theta)] \in \mathbb{R}^{d \times d}.$$

We also define the diagonal matrix $\mathcal{I}_T(\theta) = \text{diag} \{i_T(\theta)\} \in \mathbb{R}^{d \times d}$, where $\text{diag} \{i_T(\theta)\}$ denotes the diagonal elements of the matrix $i_T(\theta)$, and

$$U_T(\theta) = \mathcal{I}_T^{-\frac{1}{2}}(\theta) S_T(\theta), \quad V_T(\theta) = \mathcal{I}_T^{-\frac{1}{2}}(\theta) H_T(\theta) \mathcal{I}_T^{-\frac{1}{2}}(\theta), \quad W_{T,i}(\theta) = \mathcal{I}_T^{-\frac{1}{2}}(\theta) G_{T,i}(\theta) \mathcal{I}_T^{-\frac{1}{2}}(\theta). \quad (14)$$

With $\mathcal{I}_T \equiv \mathcal{I}_T(\theta_0)$, we then impose the following conditions on the actual log-likelihood function and the associated MLE which ensure consistency and a well-defined asymptotic distribution of the actual MLE, $\tilde{\theta}$:

C.1 The parameter space is given by a sequence of local neighborhoods,

$$\Theta_T = \left\{ \theta : \|\mathcal{I}_T^{1/2}(\theta - \theta_0)\| \leq \epsilon \right\} \subseteq \mathbb{R}^d,$$

for some $\epsilon > 0$ with $\mathcal{I}_T^{-1} = O_P(1)$.

C.2 For any $\eta > 0$, there exists a $\delta > 0$ such that

$$\lim_{T \rightarrow \infty} P \left(\sup_{\|\mathcal{I}_T^{1/2}(\theta_0 - \theta)\| > \eta} \{L_T(\theta_0) - L_T(\theta)\} \geq \delta \right) = 1.$$

C.3 $L_T(\theta)$ is three times continuously differentiable with its derivatives satisfying:

1. $(\sqrt{\nu_T} U_T(\theta_0), V_T(\theta_0)) \xrightarrow{d} (S_\infty, H_\infty)$ with $H_\infty < 0$ a.s.;
2. $\max_{j=1, \dots, d} \sup_{\theta \in \Theta_T} \|W_{j,T}(\theta)\| = O_P(1)$.

C.4 The following bounds hold for some $\delta, q > 0$:

1. $\sup_{\theta \in \Theta_T} \nu_T^{-q} \sum_{t=1}^T |\log p_t(y_t|x_t; \theta)|^{1+\delta} = O_P(1)$;
2. $\nu_T^{-q} \sum_{t=1}^T \|x_t\|^{1+\delta} = O_P(1)$ and $\nu_T^{-q} \sum_{t=1}^T \Lambda^2(\varepsilon_t) = O_P(1)$.

The above conditions C.1–C.4 are generalized versions of the conditions normally required for consistency and asymptotic normality of MLE's in stationary and ergodic models. For general non-ergodic models, simple conditions for C.2–C.4 are not available and they have to be verified

on a case-by-case basis. For the stationary case, they are implied by primitive conditions as found below in Corollary 3.2.

The specification of the parameter space in C.1 to be a sequence of non-increasing compact sets is introduced to allow for non-ergodic models. Given C.1, we re-define $\tilde{\theta}$ and $\hat{\theta}$ as the maximizers of the exact and approximate likelihood over Θ_T . This in particular means that we will only conduct our analysis of the NPSMLE in a (possibly shrinking) neighborhood of the true parameter value. Ideally, a global analysis would have been preferable, but currently, to the best of our knowledge, there exists no general results on the properties of MLEs for non-ergodic models over a fixed parameter space. For specific models, global results exist such as those found in Park and Phillips (2000, 2001), but these results appear difficult to extend to a more general setup and usually the analysis of non-ergodic models is done locally, see e.g. Kristensen and Rahbek (2010).

Assumption C.2 gives us consistency of the actual MLE, c.f. Lemma A.1. It is a combined uniform convergence and identification condition. It is for example implied by uniform convergence of the log-likelihood towards some population function which in turn identifies the true value of the parameter; see for example Amemiya (1983). Note that the condition may potentially hold even if the log-likelihood is not continuous, since its asymptotic limit will in great generality be so.

Assumption C.3 is a further strengthening of C.2 stating that the score and hessian converge in distribution after suitable normalization. This condition, in conjunction with C.1, implies both consistency and that the asymptotic distribution of the MLE is given as $\sqrt{\nu_T} \mathcal{I}_T^{1/2} (\tilde{\theta} - \theta_0) \xrightarrow{d} -H_\infty^{-1} S_\infty$, c.f. Lemma A.3.⁷

Assumption C.4 imposes bounds on a number of sample averages. They are used to show that the trimming of the simulated log-likelihood is asymptotically negligible for suitable choices of the trimming parameter a . Note that the factor ν_T in C.4 is the same as the one we normalized the log-likelihood with. The exponent $q > 0$ should be chosen to ensure that both the log-likelihood and the sample averages in C.4 are well-behaved.

In the ergodic case, we can appeal to standard results for stochastic equicontinuity—e.g. Proposition 1 in Kristensen and Rahbek (2005)—to obtain that C.4 holds with $\nu_T = T$ and $q = 1$, given that $\mathbb{E}[\|x_t\|^{1+\delta}] < \infty$ and $\mathbb{E}[\sup_{\theta \in \Theta} |\log p(y_t|x_t; \theta)|^{1+\delta}] < \infty$. See Corollary 3.2 below and its proof for further details. Furthermore, $i_T(\theta_0) = i(\theta_0) + o_P(1)$ with $i(\theta) = \mathbb{E}[\partial^2 \log p(y_t|x_t; \theta)/(\partial\theta\partial\theta')]$, such that \mathcal{I}_T can be chosen as the constant $\text{diag}\{i(\theta_0)\}$. This in turn implies that Θ_T is a fixed compact parameter set, and we get standard \sqrt{T} -convergence towards a normal distribution. Thus, in the case of stationarity, C.1–C.4 are more or less identical to the ones imposed in Fermanian and Salanié (2004, L.1–3).

In the general case, one should choose ν_T as the square of the slowest rate of convergence of the vector of MLEs. There is a tension between C.1 and C.4 in terms of the choice of ν_T . We cannot choose $\nu_T \rightarrow \infty$ too fast, since then $\|\mathcal{I}_T\| \rightarrow 0$ (in which case no information regarding θ_0 is available) and this is ruled out by C.1. On the other hand, we have to choose $\nu_T^q \rightarrow \infty$ sufficiently fast to ensure that the bounds in C.4 hold. By choosing $q > 0$ sufficiently large, C.1 and C.4

⁷Basawa and Scott (1983) and Jeganathan (1995) show what S_∞ and H_∞ look like in various cases.

will both be satisfied. However, a large value of q implies that we have to use a larger number of simulations for the NPSMLE to be asymptotically equivalent to the MLE, c.f. B.1 and B.2 below.

As an example of non-standard asymptotics of the MLE, consider a linear error-correction model, $\Delta y_t = \alpha\beta' y_{t-1} + \Omega^{1/2}\varepsilon_t$, where $\varepsilon_t \sim N(0, I_k)$. We can split the parameter vector into short-run, $\theta_1 = (\alpha, \text{vech}(\Omega))$, and long-run parameters, $\theta_2 = \beta$. The MLE $\tilde{\theta}_1$ converges with \sqrt{T} -speed towards a normal distribution, while $\tilde{\theta}_2$ is superconsistent with $T(\tilde{\theta}_2 - \theta_2)$ converging towards a Dickey-Fuller type distribution. In this case, we choose $\sqrt{\nu_T} = \sqrt{T}$, and so $i_T(\theta_0)$ and therefore \mathcal{I}_T , is not asymptotically constant. As demonstrated in Saikkonen (1995), this model satisfies C.2–C.4. Furthermore, $x_t = y_{t-1}$ satisfies $T^{-2} \sum_{t=1}^T \|x_t\|^{1+\delta} = O_P(1)$ so we can choose $q = 2$. We also refer to Park and Phillips (2001) and Kristensen and Rahbek (2010) where C.2–C.4 are verified for some non-linear, non-stationary models.

We impose the following restrictions on how the bandwidth h and trimming sequence a can converge to zero in conjunction with $N, T \rightarrow \infty$:

B. With $q, \delta > 0$ given in Condition C.4, $\bar{\lambda}_k = \lambda_{0,k} + \lambda_{1,k} + \lambda_{2,k}$, $k = 1, 2$, where $\lambda_{i,1}, \lambda_{i,2} \geq 0$, $i = 0, 1, 2$, are given in Assumptions A.1 and A.2 and for some $\gamma > 0$:

1. $|\log a| \nu_T^{q-1} N^{-\gamma(1+\delta)} \rightarrow 0$; $|\log(4a)|^{-\delta} \nu_T^{q-1} \rightarrow 0$; $T \nu_T^{-1} a^{-1} [N^{\gamma\bar{\lambda}_1} + T^{\bar{\lambda}_2}] \log(N)/\sqrt{N h^k} \rightarrow 0$; and $T \nu_T^{-1} a^{-1} [N^{\gamma\lambda_{0,1}} + T^{\lambda_{0,2}}] h^r \rightarrow 0$.
2. $|\log a| \nu_T^q N^{-\gamma(1+\delta)} \rightarrow 0$; $|\log(4a)|^{-\delta} \nu_T^q \rightarrow 0$; $T \nu_T^{-1/2} a^{-1} [N^{\gamma\bar{\lambda}_1} + T^{\bar{\lambda}_2}] \log(N)/\sqrt{N h^k} \rightarrow 0$; and $T \nu_T^{-1/2} a^{-1} [N^{\gamma\lambda_{0,1}} + T^{\lambda_{0,2}}] h^r \rightarrow 0$.

Condition B.1 is imposed when showing consistency of the NPSMLE, while B.2 will imply that the NPSMLE has the same asymptotic distribution as the MLE. The parameter $\gamma > 0$ can be chosen freely. We observe that large values of q and/or $\bar{\lambda}_1, \bar{\lambda}_2$ implies that N has to diverge at a faster rate relative to T . In practice, this means that a larger number of simulations have to be used for a given T to obtain a precise estimate. The joint requirements imposed on a , h and N are fairly complex, and it is not obvious how to choose these nuisance parameters for a given sample size T . This is a problem shared by, for example, semiparametric estimators that rely on a preliminary kernel estimator. We refer to Ichimura and Todd (2007) for an in-depth discussion of these matters. Fortunately, our simulation results indicate that standard bandwidth selection rules together with a bit of undersmoothing in general deliver satisfactory results.

Our strategy of proof is based on some apparently new results for approximate estimators, c.f. Appendix A. In particular, Theorems A.4 and A.5 establish that the NPSMLE and the MLE will be asymptotically first-order equivalent if $\hat{L}_T(\theta)$ converges uniformly towards $L_T(\theta)$ at a sufficiently fast rate. This makes our proofs considerably less burdensome than those found in other studies of simulation-based estimators—e.g. Fermanian and Salanié (2004) and Altissimo and Mele (2009)—since we do not need to analyze the simulated score and Hessian.

Theorem 3.1 *Assume that Assumptions A.1, A.2, and K.1 hold. Then the NPSMLE $\hat{\theta}$ based on (7) satisfies:*

- (i) *Under Conditions C.1, C.2, and C.4: $\mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0) = o_P(1)$ for any sequences $N \rightarrow \infty$, and $h, a \rightarrow 0$ satisfying B.1.*
- (ii) *Under Conditions C.1, C.3, and C.4: $\sqrt{\nu_T} \mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0) \xrightarrow{d} -H_\infty^{-1} S_\infty$ for any sequences $N \rightarrow \infty$, and $h, a \rightarrow 0$ satisfying B.2.*

When the data generating process is stationary and ergodic, the following more primitive conditions can be shown to imply C.1–C.4:

Corollary 3.2 *Assume that (y_t, x_t) is stationary and ergodic, and that Assumptions A.1, A.2, K.1, and B.1 hold with $q = 1$, $\nu_T = T$, $\bar{\lambda}_2 = \lambda_{0,2} = 0$ and:*

- (i) $\mathbb{E}[\|x_t\|^{1+\delta}] < \infty$, $|\log p(y|x; \theta)| \leq b_1(y|x)$, $\forall \theta \in \Theta$, with $\mathbb{E}[b_1(y_t|x_t)^{1+\delta}] < \infty$ and Θ compact;
- (ii) $\mathbb{E}[\log p(y_t|x_t; \theta)] < \mathbb{E}[\log p(y_t|x_t; \theta_0)]$, $\forall \theta \neq \theta_0$.

Then $\hat{\theta} \xrightarrow{P} \theta_0$.

If furthermore B.2 holds with $q = 1$, $\nu_T = T$ and $\bar{\lambda}_2 = \lambda_{0,2} = 0$ together with:

- (iii) $i(\theta_0) = \mathbb{E} \left[\frac{\partial \log p(y_t|x_t; \theta_0)}{\partial \theta} \frac{\partial \log p(y_t|x_t; \theta_0)}{\partial \theta'} \right]$ exists and is nonsingular;
- (iv) $\left\| \frac{\partial^2 \log p(y|x; \theta)}{\partial \theta \partial \theta'} \right\| \leq b_2(y|x)$ uniformly in a neighborhood of θ_0 with $\mathbb{E}[b_2(y_t|x_t)] < \infty$;

then $\sqrt{T}(\hat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(0, i(\theta_0)^{-1})$.

If for simplicity we set $\gamma = 0$ in B.2 and disregard the conditions on the trimming parameter a , then roughly speaking the NPSMLE will be first-order equivalent to the exact MLE in the stationary case if $Th^{2r} \rightarrow 0$ and $T/(Nh^k) \rightarrow 0$ reflecting the bias and variance due to kernel smoothing and simulations. The variance requirement seems to indicate that the usual curse of dimensionality inherent in kernel density estimation is present. However, this is caused by the initial error bounds used to establish Corollary 3.2 being overly conservative. In the following we will obtain more precise error bounds which show that the curse of dimensionality is significantly less severe. Moreover, these refined error bounds allow us to better gauge which additional biases and variances the NPSMLE suffers from due to simulations and kernel smoothing. These can potentially be used to adjust confidence bands based on the NPSMLE to take into account the additional simulation errors.

Since the higher-order analysis involves the first and second derivatives of $\hat{L}_T(\theta)$, we have to invoke the additional smoothness conditions on g and p stated in Assumptions A.3 and A.4. Under

the additional smoothness conditions, the first two derivatives of $g_{1,t}(x, \varepsilon; \theta)$ and $g_{2,t}(y_2, x, \varepsilon; \theta)$ w.r.t. θ exist, and so the first two derivatives of our density estimator are well-defined:

$$\frac{\partial \hat{p}_t(y_t|x_t; \theta)}{\partial \theta} = \frac{1}{N} \sum_{i=1}^N \left\{ \dot{Y}_{2t,i}^\theta K_h(Y_{1t,i}^\theta - y_1) + \dot{Y}_{1t,i}^\theta K_h^{(1)}(Y_{1t,i}^\theta - y_1) Y_{2t,i}^\theta \right\}, \quad (15)$$

$$\begin{aligned} \frac{\partial^2 \hat{p}_t(y_t|x_t; \theta)}{\partial \theta \partial \theta_j} &= \frac{1}{N} \sum_{i=1}^N \left\{ \ddot{Y}_{2t,i}^{\theta_j} K_h(Y_{1t,i}^\theta - y_1) + \dot{Y}_{2t,i}^\theta K_h^{(1)}(Y_{1t,i}^\theta - y_1)' \dot{Y}_{1t,i}^{\theta_j} \right\} \\ &+ \frac{1}{N} \sum_{i=1}^N \left\{ \ddot{Y}_{1t,i}^{\theta_j} K_h^{(1)}(Y_{1t,i}^\theta - y_1) Y_{2t,i}^\theta + \dot{Y}_{1t,i}^\theta K_h^{(2)}(Y_{1t,i}^\theta - y_1) \dot{Y}_{1t,i}^{\theta_j} Y_{2t,i}^\theta \right\} \\ &+ \frac{1}{N} \sum_{i=1}^N \dot{Y}_{1t,i}^\theta K_h^{(1)}(Y_{1t,i}^\theta - y_1) \dot{Y}_{2t,i}^{\theta_j}, \end{aligned} \quad (16)$$

for $j = 1, \dots, d$, where $K_h^{(i)}(y_1) = K^{(i)}(y_1/h)/h^{k+i}$, $i = 1, 2$, with $K^{(1)}(y_1) = \partial K(y_1)/(\partial y_1) \in \mathbb{R}^k$ and $K^{(2)}(y_1) = \partial^2 K(y_1)/(\partial y_1 \partial y_1') \in \mathbb{R}^{k \times k}$, while $\dot{Y}_{1t,i}^\theta = \left(\dot{Y}_{1t,i}^{\theta_1}, \dots, \dot{Y}_{1t,i}^{\theta_d} \right)' \in \mathbb{R}^{d \times k}$ with

$$\dot{Y}_{1t,i}^{\theta_j} = \frac{\partial g_{1,t}(x_t, \varepsilon_i; \theta)}{\partial \theta_j} \in \mathbb{R}^k, \quad \ddot{Y}_{1t,i}^{\theta_j} = \frac{\partial^2 g_{1,t}(x_t, \varepsilon_i; \theta)}{\partial \theta \partial \theta_j} \in \mathbb{R}^{d \times k},$$

and similarly for $\dot{Y}_{2t,i}^\theta \in \mathbb{R}^d$ and $\ddot{Y}_{2t,i}^{\theta_j} \in \mathbb{R}^d$. Lemma B.2 shows that these are uniformly consistent estimates of the actual derivatives of the conditional density p_t . The corresponding simulated version of the score is given by

$$\hat{S}_T(\theta) = \frac{1}{\nu_T} \sum_{t=1}^T \frac{\partial \hat{p}_t(y_t|x_t; \theta)}{\partial \theta} \left\{ \frac{\tau_a(\hat{p}_t(y_t|x_t; \theta))}{\hat{p}_t(y_t|x_t; \theta)} + \tau'_a(\hat{p}_t(y_t|x_t; \theta)) \log \hat{p}_t(y_t|x_t; \theta) \right\}; \quad (17)$$

an expression for the simulated version of the Hessian can be found in the proof of Theorem 3.3. We then follow Kristensen and Salanie (2010) and consider a second order functional Taylor expansion of $\hat{S}_T(\theta)$ w.r.t. \hat{p} . This takes the form:

$$\hat{S}_T(\theta_0) = S_T(\theta_0) + \nabla S_{T,N}[\hat{p} - p] + \nabla^2 S_{T,N}[\hat{p} - p, \hat{p} - p] + R_{T,N}, \quad (18)$$

where $\nabla S_{T,N}[\hat{p} - p]$ and $\nabla^2 S_{T,N}[\hat{p} - p, \hat{p} - p]$ are the first and the second order functional differentials w.r.t. p , while $R_{T,N}$ is the remainder term. The expressions of these can be found in the proof of Theorem 3.3, where the properties of the first- and second-order terms are analyzed.

To facilitate our analysis, which involves U -statistics, we restrict our attention to the stationary and β -mixing case. See e.g. Ango Nze and Doukhan (2004) for an introduction to this concept. We also assume that $p(y|x; \theta)$ is uniformly bounded away from zero thereby obviating trimming. Under these and other regularity conditions, we show that the two first terms in the expansion in equation (18) satisfy (c.f. the proof of Theorem 3.3):

$$\sqrt{T} S_T(\theta_0) + \sqrt{T} \nabla S_{T,N}[\hat{p} - p] \simeq \sqrt{T} h^r \mu_1 + Z_1 + \sqrt{\frac{T}{N}} Z_2, \quad (19)$$

where higher-order terms have been left out. Here, the first term is a bias component incurred by kernel smoothing, while the two remaining ones are variance components: Z_1 and Z_2 are two

independent variables where $Z_1 \sim \mathcal{N}(0, i(\theta_0))$ is the variance component of the observed data, while $Z_2 \sim \mathcal{N}(0, \text{Var}(\psi_2(\varepsilon_i)))$ is the variance component of the simulations. The variance of Z_2 is given by

$$\psi_2(\varepsilon_i) = \mathbb{E} \left[\frac{\dot{Y}_{2t,i}^{\theta_0}}{p(y_{2t}|x_t)} \middle| \varepsilon_i \right] - \mathbb{E} \left[\frac{s(Y_{1t,i}^{\theta_0}, y_{2t}|x_t) Y_{2t,i}^{\theta_0}}{p(y_{2t}|x_t)} \middle| \varepsilon_i \right], \quad (20)$$

where $s(y_1, y_2|x)$ denotes the score at $\theta = \theta_0$.

The second order term also contains a bias component which all non-linear, simulation-based estimators suffer from,

$$\sqrt{T} \nabla^2 S_{T,N} [\hat{p} - p, \hat{p} - p] \simeq \frac{\sqrt{T}}{Nh^{k+1}} \mu_2 + O_P(\sqrt{T}h^{2r}), \quad (21)$$

while the remainder term is of a lower order,

$$\sqrt{T} R_{T,N} = O_P(\sqrt{T}/(Nh^{k+2})^{3/2}) + O_P(\sqrt{T}h^{3r}).$$

The two leading bias terms in the above expressions, μ_1 and μ_2 , are given by:

$$\mu_1 = \sum_{|\alpha|=r} \int \int \left\{ \frac{\partial^{|\alpha|+1} p(y_t|x_t; \theta)}{\partial \theta \partial y_t^\alpha} - s(y_t|x_t; \theta) \frac{\partial^{|\alpha|} p(y_t|x_t; \theta)}{\partial y_t^\alpha} \right\} p(x_t) dx_t dy_t \in \mathbb{R}^d, \quad (22)$$

$$\mu_2 = \mathbb{E} \left[\frac{\dot{Y}_{1t,i}^{\theta_0} (Y_{2t,i}^{\theta_0})^2}{p(Y_{1t,i}^{\theta_0}, y_{2t}|x_t) p(y_{2t}|x_t)} \right] \int K(v) K^{(1)}(v) dv \in \mathbb{R}^d, \quad (23)$$

where $p(x_t)$ denotes the marginal density of x_t .

This shows that the overall bias of the estimator due to the use of simulations and kernel smoothing is $i^{-1}(\theta_0) \{h^r \mu_1 + 1/(Nh^{k+1}) \mu_2\}$ while an additional variance term relative to the exact MLE shows up and is given by $T/N \times i^{-1}(\theta_0) \text{Var}(\bar{\psi}_2(\varepsilon_i)) i^{-1}(\theta_0)$. Thus, if $\sqrt{T}h^r \rightarrow 0$ and $\sqrt{T}/(Nh^{k+1}) \rightarrow 0$, all bias terms vanish and $\sqrt{T}(\hat{\theta} - \theta_0)$ follows a normal distribution centered around zero. If furthermore $T/N \rightarrow 0$, no additional variance will be present and the NPSMLE is first-order equivalent to the true MLE. On the other hand, if either $\sqrt{T}h^r$ or $\sqrt{T}/(Nh^{k+1})$ does not vanish, a bias term will be present and the asymptotic distribution will not be centered around zero. Also, if $T/N \rightarrow 0$ there will be an increase in variance due to the presence of Z_2 . One could potentially reduce (or even remove) some of these bias and variance components by employing the techniques of Kristensen and Salanie (2010) who develop higher-order improvements of simulation-based estimators.

We collect the results in the following theorem:

Theorem 3.3 *Assume that:*

- (i) $\{(y_t, x_t)\}$ is stationary and β -mixing with geometrically decreasing mixing coefficients;
- (ii) A.1–A.4 and K.1–K.2 hold;

(iii) (i)–(iv) of Corollary 3.2 hold;

(iv) x_t is bounded and $\inf_{y_1, y_2, x, \theta} p(y_1, y_2 | x; \theta) > 0$.

Then, if $\sqrt{T}h^r \rightarrow c_1 \geq 0$, $\sqrt{T}/(Nh^{k+1}) \rightarrow c_2 \geq 0$ and $T/N \rightarrow c_3 \geq 0$,

$$\sqrt{T}(\hat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(\bar{c}, i^{-1}(\theta_0) [i(\theta_0) + c_3 \text{Var}(\psi_2(\varepsilon_i))] i^{-1}(\theta_0)),$$

where $\bar{c} = i^{-1}(\theta_0) \{c_1 \mu_1 + c_2 \mu_2\}$, with μ_1 and μ_2 as in equations (22) and (23).

The requirement (iv) is only imposed to simplify the proofs which otherwise would get overly long and complicated. We expect that the above result will still hold with the requirement (iv) replaced by additional restrictions on the trimming parameter a .

For the case where an unbiased estimator of the density is available and a new batch of simulations is used for each observation, Lee (1999) derives results similar to Theorem 3.3.

Estimation of Asymptotic Distribution To do any finite-sample inference, an estimator of the asymptotic distribution is needed. A general Monte Carlo method would be to simulate a large number of independent, long trajectories from the model and for each trajectory compute the corresponding score and hessian at $\theta = \hat{\theta}$. This would yield an approximation of the limiting distribution, $-H_\infty^{-1}S_\infty$. The computation of the score and Hessian can be done in several ways. If the model satisfies Assumption A.3, the estimators of the score and Hessian given in equation (17) and the proof of Theorem 3.3 are available. In the general case, a simple approach is to use numerical derivatives. Define:

$$\frac{\partial \hat{p}_t(y|x; \theta)}{\partial \theta_k} = \frac{\hat{p}_t(y|x; \theta + \delta e_k) - \hat{p}_t(y|x; \theta - \delta e_k)}{2\delta},$$

where e_k is the k th column of the identity matrix. We have:

$$\begin{aligned} & \frac{\partial \hat{p}_t(y|x; \theta)}{\partial \theta_k} - \frac{\partial p_t(y|x; \theta)}{\partial \theta_k} \\ &= \frac{\hat{p}_t(y|x; \theta + \delta e_k) - p_t(y|x; \theta + \delta e_k)}{2\delta} - \frac{\hat{p}_t(y|x; \theta - \delta e_k) - p_t(y|x; \theta - \delta e_k)}{2\delta} \\ &+ \left\{ \frac{p_t(y|x; \theta + \delta e_k) - p_t(y|x; \theta - \delta e_k)}{2\delta} - \frac{\partial p_t(y|x; \theta)}{\partial \theta_k} \right\}. \end{aligned}$$

Now letting $\delta = \delta(N) \rightarrow 0$ as $N \rightarrow \infty$ at a suitable rate, all three terms are $o_P(1)$. A consistent estimator of the second derivative can be obtained in a similar fashion. These can in turn be used to construct estimators of the information and score.

Approximate Simulations In many cases, the model in (1) is itself intractable, such that one cannot directly simulate from the exact model. Suppose that one on other hand has an approximation of the model at one's disposal. For example, solutions to dynamic programming problems

are typically approximated numerically, and sample paths of diffusions must be approximated by some discretization scheme.

We here derive the asymptotics of the approximate NPSMLE based on simulations from a sequence of approximate models. Assuming that the approximation error from using the approximate model relative to the true one can be made arbitrarily small, we demonstrate that the approximate NPSMLE has the same asymptotic properties as the actual MLE.

Suppose we only have the following approximations of g_{1t} and g_{2t} , $g_{M,1t}(x, \varepsilon; \theta)$ and $g_{M,2t}(y_2, x, \varepsilon; \theta)$ available, where $g_{M,kt} \rightarrow g_{kt}$, $k = 1, 2$, as $M \rightarrow \infty$ in L_1 -norm. We then redefine the simulated conditional density as:

$$\hat{p}_{M,t}(y_t|x_t; \theta) = \frac{1}{N} \sum_{i=1}^N K_h(\hat{Y}_{1t,i}^\theta - y_1) \hat{Y}_{2t,i}^\theta,$$

where $\hat{Y}_{t,i}^\theta$ is generated by the approximate model,

$$\hat{Y}_{1t,i}^\theta = g_{M,1t}(x_t, \varepsilon_i; \theta), \quad \hat{Y}_{2t,i}^\theta = g_{M,2t}(y_{2t}, x, \varepsilon_i; \theta), \quad i = 1, \dots, N.$$

Let $\hat{\theta}_M$ be the associated approximate NPSMLE,

$$\hat{\theta}_M = \arg \max_{\theta \in \Theta_T} \hat{L}_{M,T}(\theta), \quad \hat{L}_{M,T}(\theta) = \sum_{t=1}^T \tau_a(\hat{p}_{M,t}(y_t|x_t; \theta)) \log \hat{p}_{M,t}(y_t|x_t; \theta).$$

We give regularity conditions under which $\hat{\theta}_M$ has the same asymptotic properties as $\hat{\theta}$ which is based on simulations from the true model. We impose the following condition on the sequence of approximate models, and on the rates of N , h , a relative to the approximation error.

M.1 The sequence of approximate models $\{g_M\}$ satisfies for some constants $B_k, \lambda_{3,k}, \lambda_{4,k} \geq 0$, $k = 1, 2$:

$$\begin{aligned} \mathbb{E} \left[\sup_{\theta \in \Theta} \|g_{M,1t}(x, \varepsilon; \theta) - g_{1t}(x, \varepsilon; \theta)\| \right] &\leq B_1 \left(1 + \|x\|^{\lambda_{3,1}} + t^{\lambda_{3,2}} \right) \delta_{M,1}, \\ \mathbb{E} \left[\sup_{\theta \in \Theta} \|g_{M,2t}(y_2, x, \varepsilon; \theta) - g_{2t}(y_2, x, \varepsilon; \theta)\| \right] &\leq B_2 \left(1 + \|x\|^{\lambda_{4,1}} + t^{\lambda_{4,2}} \right) \delta_{M,2}, \end{aligned}$$

where $\delta_{M,k} \rightarrow 0$ as $M \rightarrow \infty$.

B.1' $T\nu_T^{-1}a^{-1}h^{-1} [N^{\gamma\lambda_{3,1}} + T^{\lambda_{3,2}}] \delta_{M,1} \rightarrow 0$ and $T\nu_T^{-1}a^{-1} [N^{\gamma\lambda_{4,1}} + T^{\lambda_{4,2}}] \delta_{M,2} \rightarrow 0$.

B.2' $T\nu_T^{-1/2}a^{-1}h^{-1} [N^{\gamma\lambda_{3,1}} + T^{\lambda_{3,2}}] \delta_{M,1} \rightarrow 0$ and $T\nu_T^{-1/2}a^{-1} [N^{\gamma\lambda_{4,1}} + T^{\lambda_{4,2}}] \delta_{M,2} \rightarrow 0$.

Assumption M.1 simply states that the approximation error is bounded as a polynomial in x and t . Assumptions B.1' and B.2' require that coefficients of this polynomial error bound, $\delta_{M,1}$ and $\delta_{M,2}$, go to zero sufficiently fast as $M \rightarrow \infty$ to control for the impact of the approximation error. Kloeden and Platen (1992) and Bruti-Liberati and Platen (2007) give primitive conditions under which the discrete-time approximation satisfies Assumption M.1; see also Detemple et al. (2006).

Under Assumptions M.1 and B.1' (B.2'), we show that the approximate NPSMLE is asymptotically equivalent to the NPSMLE up to order $\mathcal{I}_T^{1/2}$ ($\sqrt{\nu_T}\mathcal{I}_T^{1/2}$). Combining this result with Theorem 3.1, we obtain the following.

Theorem 3.4 *Assume that the conditions of Theorem 3.1 hold together with M.1. Then the approximate NPSMLE satisfies:*

1. $\mathcal{I}_T^{1/2}(\hat{\theta}_M - \theta_0) = o_P(1)$ for any sequences $M, N \rightarrow \infty$, and $h, a \rightarrow 0$ satisfying B.1 and B.1';
2. $\sqrt{\nu_T}\mathcal{I}_T^{1/2}(\hat{\theta}_M - \theta_0) \xrightarrow{d} -H_\infty^{-1}S_\infty$ for any sequences $M, N \rightarrow \infty$, and $h, a \rightarrow 0$ satisfying B.2 and B.2'.

One can exchange M.1 for

$$\mathbb{E} \left[\sup_{\theta \in \Theta} |\log p_{M,t}(y_{1t}, y_{2t}|x_t; \theta) - \log p_t(y_{1t}, y_{2t}|x_t; \theta)| \right] \leq \delta_M,$$

in which case Theorem 3.4 holds with B.1' and B.2' being replaced by the simpler conditions B.1'' : $\delta_M \rightarrow 0$ and B.2'' : $\sqrt{\nu_T}\delta_M \rightarrow 0$, respectively. However, since $p_{M,t}(y_t|x_t; \theta)$ and $p_t(y_t|x_t; \theta)$ in general are difficult to analyze, condition M.1 is easier to verify.

4 Implementing NPSML

One of the merits of NPSML is its general applicability. The applications include Markov decision processes, and discretely-sampled diffusions, where $p_t(y_t|x_t; \theta)$ typically does not have a closed-form representation but observations can still be simulated for NPSML.

The first example (Section 4.1.1) is the short-term interest rate model of Cox et al. (1985). This univariate diffusion has a known transition density, and therefore has been a popular benchmark of numerous diffusion estimation strategies (Durham and Gallant, 2002). We provide a detailed description on the implementation of NPSML in practice, and then test the validity of our approach by comparing it to the true MLE. In Section 4.1.2 we re-visit the jump-diffusion example of Section 2. The literature on estimating general jump diffusions has largely sidestepped maximum likelihood. In this context, this estimation exercise showcases the usefulness of NPSML. In Section 4.2, we briefly discuss how NPSML can be used for estimating generic Markov decision processes. We discuss diffusion models in detail here because they can be described more concisely than a typical Markov decision model, which requires a detailed description of the economic environment. We refer to Kristensen and Schjerning (2011) for implementation and analysis of smoothed maximum-likelihood estimators in discrete Markov decision model.

NPSML being for general purposes, other applications can be implemented in a similar way. At the implementation stage, only the part of the computer code that generates simulated observations needs to be modified.

4.1 Discretely-Observed Jump-Diffusions

4.1.1 Cox-Ingersoll-Ross Model

Cox et al. model short-term interest rates as a square-root diffusion process:

$$dy_t = \beta(\alpha - y_t)dt + \sigma\sqrt{y_t}dW_t.$$

We collect the unknown parameters in $\theta = (\alpha, \beta, \sigma)'$, which can be estimated with maximum likelihood. Conveniently, the transition density for a discretely-sampled path is known. When $t > s$,

$$p(y_t, t|y_s, s; \theta) = ce^{-w-v} \left(\frac{v}{w}\right)^{q/2} I_q(2\sqrt{wv}), \quad (24)$$

where $c = 2\beta / (\sigma^2 (1 - e^{-\beta(t-s)}))$, $w = cy_s e^{-\beta(t-s)}$, $v = cy_t$, $q = 2\alpha\beta/\sigma^2 - 1$, and $I_q(\cdot)$ is the modified Bessel function of the first kind of order q . Note that non-negativity of y_t requires $2\alpha\beta \geq \sigma^2$.

Following the benchmark in Durham and Gallant (2002), we generate 1,000 artificial sample paths of y_t according to the true transition density with $(\alpha, \beta, \sigma) = (0.06, 0.5, 0.15)$. Each path is then discretely observed $T = 300$ times, with time distance $t - s = 1/12$ between observations. These can be thought of monthly observations over 25 years. We estimate the diffusion parameters with true maximum likelihood and then with our NPSML.

In the implementation of NPSML, we forgo our knowledge of (24). Given a set of parameter values and y_s , we simulate paths using the Euler scheme—c.f. Kloeden and Platen (1992).⁸ We divide the interval $t - s$ into M subintervals, and recursively compute for $m = 0, \dots, M - 1$:

$$u_{m+1}^i = u_m^i + \beta(\alpha - u_m^i)\delta + \sigma\sqrt{u_m^i}\delta^{1/2}W_{m+1}^i,$$

where $u_0^i = y_s$, $\delta = \frac{t-s}{M}$, and W_{m+1}^i 's are i.i.d. standard normal random variables. Then, we set $\hat{Y}_{t,i}^\theta = u_M^i$ —the i th simulated observation of y_t conditional on y_s . Once we have generated $\hat{Y}_{t,i}^\theta$ for $i = 1, \dots, N$, then we can estimate the transition density by $\hat{p}_M(y_t, t|y_s, s) = \sum_{i=1}^N K_h(\hat{Y}_{t,i}^\theta - y_t) / N$, where K is chosen as a Gaussian kernel. It is now straightforward to construct $\hat{L}_T(\theta)$ and maximize it over Θ .

To study the finite-sample properties of our estimates as we increase the number of simulated observations per data point (N), we run NPSML for four different N 's ($N = 100, 250, 500, 625$). Another choice to be made in implementing NPSML is bandwidth (h). Silverman's rule of thumb gives us $h = 0.0035$, using the estimated standard deviation of $y_{t+1} - y_t$ and $N = 100$. In each estimation with a given h and N , we hold fixed h while we maximize over the parameter space. To assess the effect of bandwidth choice in NPSML, we consider three different h 's ($h = 0.0030, 0.0035, 0.0040$). In the simulation stage of NPSML, we use Euler scheme with $M = 10$

⁸We are approximating a continuous-time process using a discretization scheme, and hence need to appeal to Theorem 3.4.

		α	β	σ
True MLE	Bias	-0.0008	0.0805	0.0002
	Std. dev.	0.0069	0.1214	0.0034
	RMSE	0.0069	0.1456	0.0034
$h = 0.0030$	Bias	0.0036	0.0749	-0.0060
	Std. dev.	0.0143	0.1888	0.0066
	RMSE	0.0147	0.2031	0.0089
NPSMLE $N = 500$	$h = 0.0035$ Bias	0.0028	0.0778	-0.0093
	$h = 0.0035$ Std. dev.	0.0132	0.1857	0.0064
	$h = 0.0035$ RMSE	0.0135	0.2013	0.0113
$h = 0.0040$	Bias	0.0021	0.0793	-0.0129
	Std. dev.	0.0122	0.1819	0.0064
	RMSE	0.0124	0.1984	0.0144
$N = 100$	Bias	0.0071	0.0740	-0.0043
	Std. dev.	0.0180	0.2115	0.0076
	RMSE	0.0193	0.2241	0.0087
NPSMLE $h = 0.0035$	$N = 250$ Bias	0.0031	0.0817	-0.0078
	$N = 250$ Std. dev.	0.0127	0.1908	0.0068
	$N = 250$ RMSE	0.0131	0.2076	0.0103
$N = 625$	Bias	0.0006	0.0650	-0.0090
	Std. dev.	0.0076	0.1269	0.0043
	RMSE	0.0076	0.1426	0.0100

Table 1: Estimation Results for Cox-Ingersoll-Ross Model

to approximate the continuous-time process, and also adopt antithetic methods to reduce the simulation variance.

The results of the simulation study are shown in Table 1. The true MLE result is shown in the top panel. In the middle panel, we fix the number of simulated observations at $N = 500$, and vary the bandwidth for NPSML. In the bottom panel, we fix the bandwidth at $h = 0.0035$, and vary the number of simulations for NPSML. The true parameter values are $(\alpha, \beta, \sigma) = (0.06, 0.5, 0.15)$. We report the bias, standard deviation, and the root mean-squared error (RMSE) from the 1,000 estimations.

For all parameters, NPSML has larger biases and standard deviations (and hence RMSEs) than does the true MLE, with the exception being the smaller bias in the NPSML estimates of β . In the bottom panel, we see that the biases and standard deviations decrease as we increase the number of simulated observations N . The exception is the bias of the σ estimates, which increases and then decreases in absolute value with N .

Changing bandwidths produces a more complex pattern. For α , both biases and standard deviations decrease as we increase h . For β , biases increase but standard deviations fall as h goes up, although RMSE goes down like standard deviations. For σ , the magnitude of biases increases

with h , while standard deviations barely move leading to an overall increase in RMSE as h increases. However, the changes in bias and standard deviations are not very large, and we conclude that the NPSMLE appears to be quite robust towards the choice of bandwidth.

Comparing the magnitudes of biases and standard deviations from NPSML with those of the true MLE, one can get a sense of how the the additional biases and variances as stated in Theorem 3.3 are affected by our choice of bandwidth h and the number of simulated observations N .

4.1.2 Jump-Diffusion

We consider a bivariate version of the model in (3).

$$dy_{1,t} = \left(\mu - \frac{\exp(y_{2,t})}{2} \right) dt + \exp\left(\frac{y_{2,t}}{2}\right) dW_{1,t} + \log(1 + J_t)dQ_t, \quad (25)$$

$$dy_{2,t} = (\alpha_0 - \alpha_1 y_{2,t})dt + \alpha_2 dW_{2,t}. \quad (26)$$

This specification is used by Andersen et al. (2002) to model daily stock (S&P 500) returns. In their paper, $y_{2,t}$ is an unobservable stochastic volatility process, and they use EMM for estimation. Here we assume that both $y_{1,t}$ and $y_{2,t}$ are observable. One interpretation is that we infer the volatility from derivative prices as in Aït-Sahalia and Kimmel (2007). Note that it is not our intention to replicate either paper.

The factors $W_{1,t}$ and $W_{2,t}$ are standard one-dimensional Brownian motions with correlation ρ between them. Q_t is a pure jump process with jump size 1, independent of $W_{1,t}$ and $W_{2,t}$, and its jump intensity is given by λ_0 . The jump size J_t is assumed to be log-normally distributed:

$$\log(1 + J_t) \sim \mathcal{N}(-0.5\gamma^2, \gamma^2). \quad (27)$$

The parameter vector is $\theta = (\mu, \alpha_0, \alpha_1, \alpha_2, \gamma, \rho, \lambda_0)' \in \mathbb{R}^7$.

Ideally, we would like to give precise conditions under which the general jump diffusion (3) satisfies Assumptions A.1–A.4 and C.1–C.4. However, this proves very difficult without imposing strong conditions ruling out standard models considered in empirical finance, including the current example (25) and (26). Sufficient conditions for the existence of a twice-differentiable transition density for the general jump diffusion can be found in Bichteler et al. (1987) and Lo (1988), but these are rather restrictive and require, among other things, that the drift and diffusion terms be linearly bounded and infinitely differentiable. The asymptotic properties of the MLE of general jump diffusions are not very well-understood yet due to the problems of not having the transition density in closed form. Only in a few special cases, its properties can be derived; see e.g. Aït-Sahalia (2002).

In what follows, we first generate a sample path $\{(y_{1,t}, y_{2,t}) \in \mathbb{R}^2 : 0 \leq t \leq T\}$ from the true parameter values given in Table 2. We then assume that we observe this process only discretely, for $t = 0, 1, \dots, T$. Note that the discrete observations are temporally equidistant, with the interval length normalized to 1. We use these discrete observations $\{(y_{1,t}, y_{2,t}) : t = 0, 1, \dots, T\}$ as our

data. To generate this data series, we use the Euler scheme with the observation interval divided into 100 subintervals to approximate the jump-diffusion process.

Then we use NPSML without using our knowledge of the parameter values used for data generation. The first step of NPSML involves generating simulated observations from the model for any given θ , and we use the Euler scheme to approximate the data generating process. Given $(y_{1,s}, y_{2,s})$ for some period s , we divide the interval between $s + 1$ and s into M subintervals. In our benchmark estimation, we use $M = 10$. We recursively compute for $m = 1, \dots, M$:

$$u_{1,m}^i = u_{1,m-1}^i + \left(\mu - \frac{\exp(u_{2,m-1}^i)}{2} \right) \frac{1}{M} + \exp \left(\frac{u_{2,m-1}^i}{2} \right) \frac{\widetilde{W}_{1,m}^i}{\sqrt{M}} + \log(1 + J_m^i) U_m^i,$$

$$u_{2,m}^i = u_{2,m-1}^i + (\alpha_0 - \alpha_1 u_{2,m-1}^i) \frac{1}{M} + \alpha_2 \frac{W_{2,m}^i}{\sqrt{M}},$$

with $u_{1,0}^i = y_{1,s}$ and $u_{2,0}^i = y_{2,s}$ for all $i = 1, \dots, N$; J_m^i is an i.i.d. random variable with its distribution given in (27); U_m^i is an i.i.d. binomial random variable, with $Prob(U_m^i = 1) = \frac{\lambda_0}{M}$; $\widetilde{W}_{m,i}^1 = \sqrt{1 - \rho^2} W_{1,m}^i + \rho W_{2,m}^i$, where $W_{1,m}^i$ and $W_{2,m}^i$ are i.i.d. standard normal random variables. The subscript i indexes simulations. In our benchmark estimation, we use $N = 1,000$.

With the (approximate) simulated observations $\hat{Y}_{s+1,i}^\theta \equiv (u_{1,M}^i, u_{2,M}^i)$ for $i = 1, \dots, N$ with $N = 1,000$, we use (2) to obtain:

$$\hat{p}_M(y_{1,s+1}, y_{2,s+1} | y_{1,s}, y_{2,s}; \theta) = \frac{1}{N} \sum_{i=1}^N K_h \left(\hat{Y}_{s+1,i}^\theta - (y_{1,s+1}, y_{2,s+1}) \right).$$

where K is a multiplicative Gaussian kernel, $K_h(\cdot) = K_{h_1}(\cdot)K_{h_2}(\cdot)$. The bandwidths h_1 and h_2 (for y_1 and y_2 respectively) are chosen by the rule of thumb of Scott (1992, p.152). Again, we hold fixed h_1 and h_2 while we maximize over the parameter space.

With the estimated \hat{p}_t for $t = 1, 2, \dots, T$, we can evaluate the conditional likelihood, which is then maximized over the parameter space. As is typical for simulation-based estimations, when we maximize the likelihood function, we use the same set of random numbers for any θ .⁹

In our simulation study, we draw 100 sample paths of length $T = 1,000$ each, and estimate each sample path with NPSML. In column (1) of Table 2, we report the mean of the 100 point estimates for each parameter, and the 90% confidence interval constructed from the point estimates, with $N = 1,000$ and the rule-of-thumb bandwidths. The NPSML performs reasonably well, although the correlation coefficient ρ is systemically underestimated. One remarkable outcome is that the jump parameters (γ and λ_0) are rather precisely estimated, even though there are only 20 or so jump realizations in each sample path.¹⁰

To assess how sensitive the estimation results are to the choice of N (number of artificial observations used for density estimation) and the kernel bandwidths, we try different N and bandwidths.

⁹In the case of the binomial random variable U , we fix the realization of the underlying uniform random variable. For different $\theta - \lambda_0$, to be exact, U itself may have different realizations.

¹⁰We ran the same exercise with trimming of the approximate log-likelihood. The results, with N being as large as 1,000, were virtually the same as in column (1).

Parameter	True Value	(1)	(2)	(3)	(4)
μ	0.0304	0.0305 (0.0022,0.0524)	0.0305 (0.0162,0.0533)	0.0307 (0.0012,0.0661)	0.0306 (0.0084,0.0448)
α_0	-0.0120	-0.0148 (-0.0181,-0.0100)	-0.0152 (-0.0186,-0.0111)	-0.0144 (-0.0188,-0.0088)	-0.0149 (-0.0185,-0.0100)
α_1	0.0145	0.0161 (0.0116,0.2061)	0.0164 (0.0121,0.0207)	0.0160 (0.0114,0.0214)	0.0161 (0.0120,0.0214)
α_2	0.1153	0.1147 (0.1107,0.1168)	0.1139 (0.1092,0.1185)	0.1167 (0.1118,0.1198)	0.1127 (0.1089,0.1156)
γ	0.0150	0.0199 (0.0060,0.0542)	0.0310 (0.0000,0.0368)	0.0100 (0.0017,0.0158)	0.0121 (0.0085,0.0126)
ρ	-0.6125	-0.7291 (-0.7595,-0.6863)	-0.7526 (-0.7984,-0.7012)	-0.6933 (-0.7189,-0.6592)	-0.7740 (-0.8064,-0.7344)
λ_0	0.0200	0.0169 (0.0101,0.0213)	0.0133 (0.0086,0.0175)	0.0196 (0.0122,0.0197)	0.0166 (0.0104,0.0222)

Table 2: Estimation Results for Jump Diffusion. In each cell, the mean of the 100 point estimates in the simulation study is reported in the top half. In the bottom half, the 90% confidence interval constructed from the point estimates is reported. Column (1) is our benchmark with $N = 1,000$ and the rule-of-thumb bandwidths. Column (2) reports the results with $N = 750$. Column (3) is for $N = 1,000$ and bandwidths that are 20 percent narrower than those in the benchmark. Column (4) is for $N = 1,000$ and bandwidths that are 20 percent wider than those in the benchmark.

In column (2), we reduce the number of artificial observations to $N = 750$. In column (3), we use $N = 1,000$, but reduce both bandwidths by 20 percent. Finally, in column (4), we use $N = 1,000$ and bandwidths that are 20 percent greater than those in the benchmark.

When N is reduced to 750—column (2), the mean estimates move further away from the true parameter values. However, there is no clear increase or decrease in the dispersion of the estimates. The results in column (3) are of particular interest to us. Our theoretical results suggest that bandwidths should be chosen to go to zero at a faster rate than in the standard cases. With a little under-smoothing as in column (3), the mean estimates are closer to the true parameter values than in the benchmark. Note the estimates of ρ in particular. On the other hand, the results with over-smoothing as in column (4) do not compare favorably with the benchmark results with slightly more bias. We, in accordance with our theory, recommend a bandwidth narrower than what is given by the rule of thumb in actual implementations. At the same time, we would like to emphasize that the estimator in general is quite robust to the choice of bandwidth.

4.2 Markov Decision Processes and Dynamic Games

Another class of economic models that NPSML can readily be applied to is Markov decision processes: see Rust (1994) for an overview. In these models, the transition density is given by

$$p(y_t|x_t; \theta) = \int p(y_t|x_t, u_t)q(u_t)du_t,$$

where $p(y_t|x_t, u_t)$ is typically governed by an optimal decision rule of a dynamic programming problem. The integral on the right-hand side does not have a closed-form representation, except in few special cases. However, conditioning on x_t , one can simulate u_t and hence y_t , and use kernel methods to estimate $p(y_t|x_t)$. Therefore, NPSML is feasible; see Kristensen and Schjerning (2011);

NPSML can also be used to estimate a related class of economic models: Markov-perfect equilibria of dynamic games. Ericson and Pakes (1995) provide a canonical framework for this literature: a dynamic model of oligopolistic industry with entry and exit. The equilibrium transition probability of this model is given by

$$p_t(\omega_{t+1}|\omega_t; \theta), \quad \omega \in \mathcal{Z}^n,$$

where \mathcal{Z} is a finite set of states, and n is the number of state variables. The transition probability depends on individual firm-specific shocks, industry-wide shocks, and Markov-perfect strategies of firms regarding entry, exit and investment.¹¹ Firms' strategies represent an optimal decision rule of a dynamic programming problem. Clearly, the transition probability does not have a closed-form representation, but it is still possible to simulate observations from the model conditioning on ω_t .¹² Thus, NPSML is feasible. The computational burden of such models grow quickly with n . Doraszelski and Judd (2008) show how one can avoid this problem by casting the problems in continuous time. NPSML is readily applicable to such continuous-time dynamic stochastic games as well.

5 Concluding Remarks

We have generalized the NPSML of Fermanian and Salanié (2004) to deal with dynamic models, including nonstationary and time-inhomogeneous ones. Theoretical conditions in terms of the number of simulations and the bandwidth are given ensuring that the NPSMLE inherits the asymptotic properties of the infeasible MLE.

This method is applicable to general classes of models, and can be implemented with ease. Our finite-sample simulation study demonstrates that the method works well in practice.

One limitation of the paper is that we only consider the cases where it is possible to simulate the dependent variable conditional on finitely-many past observations. This excludes cases with

¹¹In this class of models, conditioning on ω_t , ω_{t+1} depends not only on individual actions but also on idiosyncratic and aggregate shocks. To obtain the transition probability, all the shocks need to be integrated out.

¹²In solving individual firms' dynamic programming problem, one needs to know their continuation value, and hence the transition probability. Therefore, for a given θ , one needs to compute a fixed point in $p_t(\omega_{t+1}|\omega_t)$.

latent dynamics. Extensions to methods with built-in nonlinear filters that explicitly account for latent variable dynamics are worked out in a companion paper (Brownlees et al., 2011) based on the main results given here.

Appendix

A Some General Results for Approximate Estimators

We first establish some general results for approximate MLEs. These results will then be applied to show the desired results for our proposed NPSMLE. In the rest of this part of the appendix, we consider the fully general situation where $\hat{L}_T(\theta) = \hat{L}_{T,N}(\theta)$ is a sequence of approximations to $L_T(\theta)$ (not necessarily the nonparametric simulated one proposed in the main text) of some infeasible (quasi-)log-likelihood function $L_T(\theta)$. We will analyze the impact of the approximation on the corresponding approximate estimator $\hat{\theta}$.

We first establish the asymptotic properties of the true MLE $\tilde{\theta}$ under Assumptions C.1–C.3. Next, we give a general set of conditions for the approximate estimator $\hat{\theta}$ to be asymptotically equivalent to $\tilde{\theta}$.

A.1 Asymptotics of True MLE

Lemma A.1 *Assume that Assumptions C.1 and C.2 hold. Then $\mathcal{I}_T^{1/2}(\tilde{\theta} - \theta_0) = o_P(1)$.*

Proof We introduce a normalized version of the parameter $\theta \in \Theta_T$, $\xi := \mathcal{I}_T^{1/2}(\theta - \theta_0)$, and define the corresponding likelihood in terms of this new parameterization, $Q_T(\xi) := L_T(\theta_0 + \mathcal{I}_T^{-1/2}\xi)$. The claim will now follow from $\tilde{\xi} := \mathcal{I}_T^{1/2}(\tilde{\theta} - \theta_0) \xrightarrow{P} 0$. However, under Assumptions C.1 and C.2, it is easily seen that $Q_T(\xi)$ satisfies the conditions of Theorem 3.4 of White (1994) with $\bar{Q}_T(\xi) = Q_T(\xi)$ from which the desired result follows. \square

Lemma A.2 *Assume that Assumption C.1 holds. Then Assumption C.3 implies Assumption C.2.*

Proof Use a second order Taylor expansion to obtain for any bounded sequence $\xi_T \in \mathbb{R}^d$ such that $\theta_0 + \mathcal{I}_T^{-1/2}\xi_T \in \Theta_T$,

$$L_T(\theta_0 + \mathcal{I}_T^{-1/2}\xi_T) - L_T(\theta_0) = U_T(\theta_0)\xi_T + \frac{1}{2}\xi_T'V_T(\bar{\theta})\xi_T,$$

for some $\bar{\theta} \in [\theta_0, \theta_0 + \mathcal{I}_T^{-1/2}\xi_T] \in \Theta_T$, and with $U_T(\theta)$ and $V_T(\theta)$ defined in equation (14). By another application of Taylor's Theorem,

$$\begin{aligned} |\xi_T'V_T(\bar{\theta})\xi_T - \xi_T'V_T(\theta_0)\xi_T| &= \left| \xi_T'\mathcal{I}_T^{-1/2} [H_T(\bar{\theta}) - H_T(\theta_0)] \mathcal{I}_T^{-1/2}\xi_T \right| \\ &\leq \left\{ \max_{i=1,\dots,d} \sup_{\theta \in \Theta_T} \xi_T'W_{T,i}(\theta)\xi_T \right\} \times \|\mathcal{I}_T^{-1/2}\xi_T\| = o_P(1), \end{aligned}$$

where we have used Assumption C.3.2 and the fact that $\|\mathcal{I}_T^{-1/2}\xi_T\| = o_P(1)$. Thus,

$$L_T(\theta_0 + \mathcal{I}_T^{-1/2}\xi_T) - L_T(\theta_0) = U_T(\theta_0)\xi_T + \frac{1}{2}\xi_T'V_T(\theta_0)\xi_T + o_P(1) = \frac{1}{2}\xi_T'H_\infty\xi_T + o_P(1),$$

where the second equality follows from C.3.1. Since $\xi_T'H_\infty\xi_T < 0$ a.s., $L_T(\theta)$ is concave with a unique maximum at $\theta = \theta_0$ in Θ_T (with probability tending to one). In particular, C.2 holds. \square

Lemma A.3 *Assume that Assumptions C.1 and C.3 hold. Then the MLE $\tilde{\theta}$ satisfies*

$$\sqrt{\nu_T} \mathcal{I}_T^{1/2}(\tilde{\theta} - \theta_0) \xrightarrow{d} -H_\infty^{-1} S_\infty.$$

Proof By Lemmas A.1 and A.2, we know that $\tilde{\theta}$ is consistent. A first-order Taylor expansion of the score and C.3.2 together with the same arguments as in the proof of Lemma A.2 yield

$$\sqrt{\nu_T} U_T(\theta_0) = -V_T(\tilde{\theta}) \sqrt{\nu_T} \mathcal{I}_T^{1/2}(\tilde{\theta} - \theta_0) = -V_T(\theta_0) \sqrt{\nu_T} \mathcal{I}_T^{1/2}(\tilde{\theta} - \theta_0) + o_P(1),$$

and the result now follows from C.3.1. \square

A.2 Asymptotics of Approximate MLE

Theorem A.4 *Assume that C.1 and C.2 hold and $\sup_{\theta \in \Theta_T} |\hat{L}_T(\theta) - L_T(\theta)| = o_P(1)$ as $T \rightarrow \infty$ for a sequence $N = N(T) \rightarrow \infty$. Then $\mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0) = o_P(1)$.*

Proof We wish to show that for any $\eta > 0$,

$$P(\|\mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0)\| > \eta) \rightarrow 0, \quad T \rightarrow \infty.$$

Let $\eta > 0$ be given. Then by C.2 there exists a $\delta > 0$ such that, $L_T(\theta_0) - L_T(\hat{\theta}) \geq \delta$ with probability tending to 1. Thus, as $T \rightarrow \infty$,

$$P(\|\mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0)\| > \eta) \leq P(L_T(\theta_0) - L_T(\hat{\theta}) \geq \delta).$$

We then have to show that the right-hand side converges to zero. To this end, write

$$L_T(\theta_0) - L_T(\hat{\theta}) = \left\{ L_T(\theta_0) - \hat{L}_T(\theta_0) \right\} + \left\{ \hat{L}_T(\theta_0) - L_T(\hat{\theta}) \right\},$$

where,

$$L_T(\theta_0) - \hat{L}_T(\theta_0) \leq \sup_{\theta \in \Theta_T} |L_T(\theta) - \hat{L}_T(\theta)| = o_P(1),$$

while, by the definition of $\hat{\theta}$,

$$\hat{L}_T(\theta_0) - L_T(\hat{\theta}) \leq \hat{L}_T(\hat{\theta}) - L_T(\hat{\theta}) \leq \sup_{\theta \in \Theta_T} |L_T(\theta) - \hat{L}_T(\theta)| = o_P(1). \quad \square$$

Next, we state two results for the approximate estimator to have the same asymptotic distribution as the actual MLE. Theorem A.5 establishes this result only requiring that the approximate likelihood function satisfies $\sup_{\theta \in \Theta_T} |\hat{L}_T(\theta) - L_T(\theta)| = o_P(1/\sqrt{\nu_T})$. Theorem A.6 imposes stronger smoothness conditions, requiring that $\hat{L}_T(\theta)$ be twice differentiable with derivatives $\hat{S}_T(\theta)$ and $\hat{H}_T(\theta)$; on the other hand we only require $\|\hat{S}_T(\theta_0) - S_T(\theta_0)\| = o_P(1/\|\sqrt{\nu_T} \mathcal{I}_T^{1/2}\|)$ and $\sup_{\theta \in \Theta_T} \|\hat{H}_T(\theta) - H_T(\theta)\| = o_P(\|\mathcal{I}_T^{-1}\|)$ which are weaker convergence restrictions than $o_P(1/\sqrt{\nu_T})$, since $\|\mathcal{I}_T^{-1/2}\| = O(1)$. So there is a trade-off between smoothness and the convergence rate.

Theorem A.5 Assume that Assumptions C.1 and C.3 hold, and $\sup_{\theta \in \Theta_T} |\hat{L}_T(\theta) - L_T(\theta)| = o_P(1/\sqrt{\nu_T})$ for some sequence $N = N(T) \rightarrow \infty$. Then, $\sqrt{\nu_T} \mathcal{I}_T^{1/2}(\hat{\theta} - \tilde{\theta}) = o_P(1)$. In particular, $\sqrt{\nu_T} \mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0) \xrightarrow{d} -H_\infty^{-1} S_\infty$.

Proof Define $\tilde{\xi}_T = \mathcal{I}_T^{1/2}(\tilde{\theta} - \theta_0)$ and $\hat{\xi}_T = \mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0)$. By a second-order Taylor expansion of $L_T(\hat{\theta})$ around $\tilde{\theta}$ together with the fact that $S_T(\tilde{\theta}) = 0$ yields, for some $\bar{\theta} \in [\tilde{\theta}, \hat{\theta}]$,

$$L_T(\tilde{\theta}) - L_T(\hat{\theta}) = L_T(\theta_0 + \mathcal{I}_T^{-1/2} \tilde{\xi}_T) - L_T(\theta_0 + \mathcal{I}_T^{-1/2} \hat{\xi}_T) = \frac{1}{2} (\tilde{\xi}_T - \hat{\xi}_T)' \mathcal{I}_T^{-1/2} H_T(\bar{\theta}) \mathcal{I}_T^{-1/2} (\tilde{\xi}_T - \hat{\xi}_T).$$

By the same arguments as in the proof of Lemma A.2, we obtain

$$L_T(\tilde{\theta}) - L_T(\hat{\theta}) = \frac{1}{2} (\tilde{\xi}_T - \hat{\xi}_T)' H_\infty (\tilde{\xi}_T - \hat{\xi}_T) + o_P(1/\sqrt{\nu_T}).$$

Write the left-hand side as

$$\sqrt{\nu_T} \{L_T(\tilde{\theta}) - L_T(\hat{\theta})\} = \sqrt{\nu_T} \{L_T(\tilde{\theta}) - \hat{L}_T(\tilde{\theta})\} + \sqrt{\nu_T} \{\hat{L}_T(\tilde{\theta}) - L_T(\hat{\theta})\}$$

where

$$\sqrt{\nu_T} \{L_T(\tilde{\theta}) - \hat{L}_T(\tilde{\theta})\} \leq \sqrt{\nu_T} \sup_{\theta \in \Theta_T} |\hat{L}_T(\theta) - L_T(\theta)| = o_P(1),$$

and, using that $\hat{\theta}$ is the maximizer of $\hat{L}_T(\theta)$,

$$\sqrt{\nu_T} \{\hat{L}_T(\tilde{\theta}) - L_T(\hat{\theta})\} \leq \sqrt{\nu_T} \{\hat{L}_T(\hat{\theta}) - L_T(\hat{\theta})\} \leq \sqrt{\nu_T} \sup_{\theta \in \Theta_T} |\hat{L}_T(\theta) - L_T(\theta)| = o_P(1).$$

Thus,

$$\|\sqrt{\nu_T} \mathcal{I}_T^{1/2}(\hat{\theta} - \tilde{\theta})\|^2 \leq \|H_\infty^{-1}\| \sqrt{\nu_T} (\tilde{\xi}_T - \hat{\xi}_T)' H_\infty (\tilde{\xi}_T - \hat{\xi}_T) = o_P(1). \quad \square$$

Theorem A.6 Assume that Assumptions C.1 and C.3 hold together with:

(i) $\theta \mapsto \hat{L}_T(\theta)$ is twice differentiable in Θ_T .

(ii) There exists a sequence $N = N(T) \rightarrow \infty$ such that $\left\| \mathcal{I}_T^{1/2} \left\{ \hat{S}_T(\theta_0) - S_T(\theta_0) \right\} \right\| = o_P(1/\sqrt{\nu_T})$,
and $\left\| \mathcal{I}_T^{-1} \left\{ H_T(\theta_0) - \hat{H}_T(\theta_0) \right\} \right\| = o_P(1)$.

Then $\sqrt{\nu_T} \mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0) \xrightarrow{d} -H_\infty^{-1} S_\infty$ for this sequence N .

Proof By standard Taylor expansions,

$$0 = \sqrt{\nu_T} \mathcal{I}_T^{-1/2} S_T(\tilde{\theta}) = \sqrt{\nu_T} \mathcal{I}_T^{-1/2} S_T(\theta_0) + \mathcal{I}_T^{-1/2} H_T(\theta_0) \mathcal{I}_T^{-1/2} \sqrt{\nu_T} \mathcal{I}_T^{1/2}(\tilde{\theta} - \theta_0) + o_P(1),$$

and

$$0 = \sqrt{\nu_T} \mathcal{I}_T^{-1/2} \hat{S}_T(\hat{\theta}) = \sqrt{\nu_T} \mathcal{I}_T^{-1/2} \hat{S}_T(\theta_0) + \mathcal{I}_T^{-1/2} \hat{H}_T(\theta_0) \mathcal{I}_T^{-1/2} \sqrt{\nu_T} \mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0) + o_P(1).$$

Subtracting the two equations, and using that $\mathcal{I}_T^{-1/2} H_T(\theta_0) \mathcal{I}_T^{-1/2} \xrightarrow{d} H_\infty > 0$, we obtain

$$0 = \sqrt{\nu_T} \mathcal{I}_T^{1/2} \left\{ S_T(\theta_0) - \hat{S}_T(\theta_0) \right\} + H_\infty \sqrt{\nu_T} \mathcal{I}_T^{1/2}(\tilde{\theta} - \hat{\theta}) + \mathcal{I}_T^{-1/2} \left\{ H_T(\theta_0) - \hat{H}_T(\theta_0) \right\} \mathcal{I}_T^{-1/2} \sqrt{\nu_T} \mathcal{I}_T^{1/2}(\hat{\theta} - \theta_0)$$

The result now follows from condition (ii) and Lemma A.3. \square

B Properties of Simulated Conditional Density

We here establish uniform convergence of \hat{p}_t given in equation (7) and its derivatives w.r.t. θ .

Lemma B.1 *Assume that A.1, A.2 and K.1 hold. Then \hat{p}_t in (7) satisfies for all $y_2 \in \mathcal{Y}_2$ and any compact set Θ :*

$$\sup_{1 \leq t \leq T} \sup_{y_1 \in \mathbb{R}^k} \sup_{\|x\| \leq d_n} \sup_{\theta \in \Theta} |\hat{p}_t(y_1, y_2|x; \theta) - p_t(y_1, y_2|x; \theta)| = O_P \left(\left[d_n^{\bar{\lambda}_1} + T^{\bar{\lambda}_2} \right] \log(N) / \sqrt{N h^k} \right) + O_P \left(\left[d_n^{\lambda_{0,1}} + T^{\lambda_{0,2}} \right] h^r \right),$$

where $\bar{\lambda}_i = \lambda_{0,i} + \lambda_{1,i} + \lambda_{2,i}$, $i = 1, 2$.

Proof Define $\gamma = (x, \theta, t) \in \Gamma = \mathcal{X}_t \times \Theta \times \{1, 2, 3, \dots\}$. Write $\hat{p}(y_1, y_2; \gamma) = \hat{p}(y_1, y_2|x; \theta)$ and $p(y_1, y_2; \gamma) = p(y_1, y_2|x; \theta)$. We split up into a bias and a variance component:

$$\begin{aligned} \hat{p}(y_1, y_2; \gamma) - p(y_1, y_2; \gamma) &= \{\mathbb{E}[\hat{p}(y_1, y_2; \gamma)] - p(y_1, y_2; \gamma)\} + \{\hat{p}(y_1, y_2; \gamma) - \mathbb{E}[\hat{p}(y_1, y_2; \gamma)]\} \\ &=: \text{Bias}(y_1, y_2; \gamma) + \text{Var}(y_1, y_2; \gamma). \end{aligned}$$

Using standard arguments for kernel estimators, the bias term can be shown to satisfy

$$|\text{Bias}(y_1, y_2; \gamma)| \leq h^r \int |K(v)| |v|^r dv \times \left| \frac{\partial^r p_t(y_1, y_2; \gamma)}{\partial y_1^r} \right| + o(h^r).$$

Thus, using the bound imposed on the r -th derivative, $|\text{Bias}(y_1, y_2; \gamma)| = O \left(\left[d_n^{\lambda_{0,1}} + T^{\lambda_{0,2}} \right] h^r \right)$ uniformly over (y_1, γ) . To establish the uniform rate of the variance term, we apply the result of Kristensen (2009, Theorem 1) for averages of the form

$$\hat{\Psi}(x; \gamma) = \frac{1}{nh^d} \sum_{i=1}^n Y_i(\gamma) G \left(\frac{X_i(\gamma) - x}{h} \right), \quad (28)$$

for some kernel-type function G . With $Y_i(\gamma) = g_{2,t}(y_2, x, \varepsilon_i; \theta)$, $X_i(\gamma) = g_{1,t}(x, \varepsilon_i; \theta)$ and $G = K$, our simulated density can be written in this form. We then verify that his conditions A.1–A.6 in Kristensen (2009) are satisfied under our assumptions. His A.1 is trivially satisfied since we have i.i.d. draws, while his Condition A.6 imposed on G is implied by our K.1. The bounds in his A.4 and A.5 becomes in our case, using his Remark 2.2:

$$\begin{aligned} \tilde{B}_0 &= p(y_1; \gamma), \quad \tilde{B}_1 = \|y_1\|^k \mathbb{E} [|Y_i(\gamma)| |X_i(\gamma) = y_1] p(y_1; \gamma), \\ \tilde{B}_2 &= \|y_1\|^k \mathbb{E} [|\partial_\gamma Y_i(\gamma)| |X_i(\gamma) = y_1] p(y_1; \gamma), \\ \tilde{B}_3 &= \|y_1\|^k \mathbb{E} [|Y_i(\gamma)| |\partial_\gamma X_i(\gamma)| |X_i(\gamma) = y_1] p(y_1; \gamma), \end{aligned}$$

where $\partial_\gamma Y_i(\gamma)$ and $\partial_\gamma X_i(\gamma)$ denote their derivatives w.r.t. $\gamma = (x, \theta, t)$. By Assumption A.2, $\tilde{B}_0 = O \left(1 + \|x\|^{\lambda_{0,1}} + t^{\lambda_{0,2}} \right)$ while, using A.1,

$$\begin{aligned} \mathbb{E} [|Y_i(\gamma)| |X_i(\gamma) = y_1] &= \int_{\{e: g_{1,t}(x, e; \theta) = y_1\}} |g_{2,t}(y_2, x, e; \theta)| dF_\varepsilon(e) \leq \int |g_{2,t}(y_2, x, e; \theta)| dF_\varepsilon(e) \\ &= \mathbb{E} [|g_{2,t}(y_2, \varepsilon; \gamma)|] \leq \mathbb{E} [\Lambda(\varepsilon)] \left[1 + \|x\|^{\lambda_{2,1}} + t^{\lambda_{2,2}} \right], \end{aligned}$$

and similarly for the two other conditional expectations in \tilde{B}_2 and \tilde{B}_3 . Thus,

$$\tilde{B}_k \leq \mathbb{E}[\Lambda(\varepsilon)] \left(1 + \|x\|^{\lambda_1} + t^{\lambda_2}\right) \|y_1\|^q p(y_1; \gamma) = O\left(1 + \|x\|^{\lambda_{0,1} + \lambda_{2,1}} + t^{\lambda_{0,2} + \lambda_{2,2}}\right),$$

for $k = 1, 2$, where the second equality follows from A.2, while $\tilde{B}_3 = O\left(1 + \|x\|^{\bar{\lambda}_1} + t^{\bar{\lambda}_2}\right)$. \square

Lemma B.2 *Assume that A.1–A.4 and K.1 hold. Then $\partial^i \hat{p}_t / \partial \theta^i$, $i = 1, 2$, given in (15) and (16) satisfy for all $y_2 \in \mathcal{Y}_2$ and any compact set Θ :*

$$\begin{aligned} & \sup_{1 \leq t \leq T} \sup_{y_1 \in \mathbb{R}^k} \sup_{\|x\| \leq d_n} \sup_{\theta \in \Theta} \left| \frac{\partial^i \hat{p}_t(y_1, y_2 | x; \theta)}{\partial \theta^i} - \frac{\partial^i p_t(y_1, y_2 | x; \theta)}{\partial \theta^i} \right| \\ &= O_P\left(\left[d_n^{\bar{\lambda}_1} + T^{\bar{\lambda}_2}\right] \log(N) / \sqrt{N h^{k+i}}\right) + O_P\left(\left[d_n^{\lambda_{0,1}} + T^{\lambda_{0,2}}\right] h^r\right). \end{aligned}$$

Proof We only give a proof for the first derivative. The proof for the second one follows along the same lines. We proceed as in the proof of Lemma B.1. With $Y_{1t,i}^\theta = g_{1,t}(x_t, \varepsilon_i; \theta)$ and $Y_{2t,i}^\theta = g_{2,t}(y_{2t}, x_t, \varepsilon_i; \theta)$, it follows from the expression in equation (15) that conditional on (y_{1t}, y_{2t}, x_t) :

$$\mathbb{E}\left[\frac{\partial \hat{p}_t(y_{1t}, y_{2t} | x_t; \theta)}{\partial \theta}\right] = \frac{1}{h^k} \int \dot{Y}_{2t,i}^\theta K\left(\frac{Y_{1t,i}^\theta - y_{1t}}{h}\right) dF_\varepsilon(\varepsilon) + \frac{1}{h^{k+1}} \int \dot{Y}_{1t,i}^\theta K^{(1)}\left(\frac{Y_{1t,i}^\theta - y_{1t}}{h}\right) Y_{2t,i}^\theta dF_\varepsilon(\varepsilon),$$

where, uniformly over (t, x_t, θ) ,

$$\begin{aligned} \frac{1}{h^k} \int \dot{Y}_{2t,i}^\theta K\left(\frac{Y_{1t,i}^\theta - y_{1t}}{h}\right) dF_\varepsilon(\varepsilon) &= \int K(v) p(y_{1t} + vh | x_t; \theta) \frac{\partial p_t(y_{2t} | y_{1t} + vh, x_t; \theta)}{\partial \theta} dv \\ &= \frac{\partial p_t(y_{2t} | y_{1t}, x_t; \theta)}{\partial \theta} p_t(y_{1t} | x_t; \theta) + O\left(\left[d_n^{\lambda_1} + T^{\lambda_2}\right] h^r\right), \\ \frac{1}{h^{k+1}} \int \dot{Y}_{1t,i}^\theta K^{(1)}\left(\frac{Y_{1t,i}^\theta - y_{1t}}{h}\right) Y_{2t,i}^\theta dF_\varepsilon(\varepsilon) &= p_t(y_{2t} | y_{1t}, x_t; \theta) \frac{\partial p_t(y_{1t} | x_t; \theta)}{\partial \theta} + O\left(\left[d_n^{\lambda_1} + T^{\lambda_2}\right] h^r\right). \end{aligned}$$

For the variance component, we again apply the results of Kristensen (2009). With $\gamma = (x, \theta, t)$ and $X_{n,i}(\gamma) = g_{1,t}(x, \varepsilon_i; \theta)$, $\partial \hat{p}_t / \partial \theta$ can be written as the sum of two kernel averages, each of the form (28); the first with $Y_{n,i}(\gamma) = \dot{g}_{1,t}(x, \varepsilon_i; \theta) g_{2,t}(y_2, x, \varepsilon_i; \theta)$ and $G = K^{(1)}$, and the second with $Y_{n,i}(\gamma) = \dot{g}_{2,t}(y_2, x, \varepsilon_i; \theta)$ and $G = K$. Under the conditions imposed on our model, his A.1–A.5 hold. \square

C Proofs

Proof of Theorem 3.1 The first part of the result will follow if we can verify the conditions in Theorem A.4. In order to do this, we introduce an additional trimming function, $\tilde{\tau}_{a,t} = \tau_a(\hat{p}_t(y_t | x_t; \theta)) \mathbb{I}\{\|x_t\| \leq N^\gamma\}$, where $\mathbb{I}\{\cdot\}$ is the indicator function and $\gamma > 0$ is chosen as in Assumption B.1, and two trimming sets,

$$A_{1,t}(\varepsilon) = \{\hat{p}_t(y_t | x_t; \theta) \geq \varepsilon a, \|x_t\| \leq N^\gamma\}, \quad A_{2,t}(\varepsilon) = \{p_t(y_t | x_t; \theta) \geq \varepsilon a, \|x_t\| \leq N^\gamma\},$$

for any $\varepsilon > 0$. Defining $A_t(\varepsilon) = A_{1,t}(\varepsilon) \cap A_{2,t}(\varepsilon)$, it follows by the same arguments as in Andrews (1995, p. 588), $A_{2,t}(2\varepsilon) \subseteq A_{1,t}(\varepsilon) \subseteq A_t(\varepsilon/2)$ w.p.a.1 as $N \rightarrow \infty$ under Assumption B.1. Thus, $\mathbb{I}_{A_{2,t}(4)} \leq \mathbb{I}_{A_{1,t}(2)} \leq \tilde{\tau}_{a,t} \leq \mathbb{I}_{A_{1,t}(1/2)} \leq \mathbb{I}_{A_t(1/4)}$.

We then split up $\hat{L}_T(\theta) - L_T(\theta)$ into three terms,

$$\begin{aligned} \hat{L}_T(\theta) - L_T(\theta) &= \frac{1}{\nu_T} \sum_{t=1}^T [\tau_a(\hat{p}_t(y_t|x_t; \theta)) - \tilde{\tau}_{a,t}] \log \hat{p}_t(y_t|x_t; \theta) \\ &\quad + \frac{1}{\nu_T} \sum_{t=1}^T \tilde{\tau}_{a,t} [\log \hat{p}_t(y_t|x_t; \theta) - \log p_t(y_t|x_t; \theta)] + \frac{1}{\nu_T} \sum_{t=1}^T [\tilde{\tau}_{a,t} - 1] \log p_t(y_t|x_t; \theta) \\ &=: B_1(\theta) + B_2(\theta) + B_3(\theta), \end{aligned}$$

and then show that $\sup_{\theta \in \Theta_T} |B_i(\theta)| = o_P(1)$, $i = 1, 2, 3$. By Assumption C.4,

$$|B_1(\theta)| \leq \frac{|\log a|}{\nu_T} \sum_{t=1}^T \mathbb{I}\{\|x_t\| > N^\gamma\} \leq \frac{|\log a| \nu_T^{q-1}}{N^{\gamma(1+\delta)}} \frac{1}{\nu_T^q} \sum_{t=1}^T \|x_t\|^{1+\delta} \leq \frac{|\log a| \nu_T^{q-1}}{N^{\gamma(1+\delta)}} \times O_P(1),$$

while,

$$\begin{aligned} |B_2(\theta)| &\leq \frac{1}{\nu_T} \sum_{t=1}^T \mathbb{I}_{A_t(1/4)} |\log \hat{p}_t(y_t|x_t; \theta) - \log p_t(y_t|x_t; \theta)| \\ &\leq \frac{T}{a\nu_T} \times \sup_{1 \leq t \leq T} \sup_{\theta \in \Theta} \sup_{y \in \mathbb{R}^k} \sup_{\|x\| \leq N^\gamma} |\hat{p}_t(y_t|x_t; \theta) - p_t(y_t|x_t; \theta)|. \end{aligned}$$

The final term is bounded by

$$\begin{aligned} |B_3(\theta)| &\leq \frac{1}{\nu_T} \sum_{t=1}^T |\tilde{\tau}_{a,t} - 1| |\log p_t(y_t|x_t; \theta)| \\ &\leq \frac{1}{\nu_T} \sum_{t=1}^T \mathbb{I}\{p_t(y_t|x_t; \theta) < 4a\} |\log p_t(y_t|x_t; \theta)| + \frac{1}{\nu_T} \sum_{t=1}^T \mathbb{I}\{\|x_t\| > N^\gamma\} |\log p_t(y_t|x_t; \theta)| \\ &=: B_{3,1}(\theta) + B_{3,2}(\theta). \end{aligned}$$

First, as $a \rightarrow 0$,

$$\begin{aligned} |B_{3,1}(\theta)| &\leq \frac{1}{\nu_T} \sum_{t=1}^T \mathbb{I}\{p_t(y_t|x_t; \theta) < 4a\} |\log p_t(y_t|x_t; \theta)| \\ &= \frac{1}{\nu_T} \sum_{t=1}^T \mathbb{I}\{|\log p_t(y_t|x_t; \theta)| > |\log(4a)|\} |\log p_t(y_t|x_t; \theta)| \\ &\leq |\log(4a)|^{-\delta} \nu_T^{q-1} \frac{1}{\nu_T^q} \sum_{t=1}^T |\log p_t(y_t|x_t; \theta)|^{1+\delta} = |\log(4a)|^{-\delta} \nu_T^{q-1} \times O_P(1). \end{aligned}$$

where we have used Assumption C.5. Similarly, again by C.5,

$$\begin{aligned} |B_{3,2}(\theta)| &\leq \frac{1}{\nu_T} \sum_{t=1}^T \mathbb{I}\{\|x_t\| > N^\gamma\} |\log p_t(y_t|x_t; \theta)| \\ &\leq \left\{ \frac{1}{\nu_T} \sum_{t=1}^T \mathbb{I}\{\|x_t\| > N^\gamma\} \right\}^{\delta/(1+\delta)} \left\{ \frac{1}{\nu_T} \sum_{t=1}^T |\log p_t(y_t|x_t; \theta)|^{1+\delta} \right\}^{1/(1+\delta)} \end{aligned}$$

$$\begin{aligned}
&\leq \frac{\nu_T^{q-1}}{N^{\gamma(1+\delta)}} \left\{ \frac{1}{\nu_T^q} \sum_{t=1}^T \|x_t\|^{1+\delta} \right\}^{\delta/(1+\delta)} \left\{ \frac{1}{\nu_T^q} \sum_{t=1}^T |\log p_t(y_t|x_t; \theta)|^{1+\delta} \right\}^{1/(1+\delta)} \\
&= \frac{\nu_T^{q-1}}{N^{\gamma(1+\delta)}} \times O_P(1).
\end{aligned}$$

The consistency result follows from Theorem A.4 together with Lemma B.1 and Assumption B.1.

To show the second result, we merely have to strengthen the convergence of $\hat{L}_T(\theta)$ to take place with rate ν_T , c.f. Theorem A.5. One can still apply the above bounds which now have to go to zero with rate ν_T . This is ensured by Assumption B.2. \square

Proof of Corollary 3.2 We verify that Assumption C.1–C.4 hold with $\nu_T = T$ and $q = 1$ under the conditions imposed in the corollary. First, we obtain by LLN for mixing sequences that $i_T(\theta_0) = i(\theta_0) + o_P(1)$ with $i(\theta_0) = \mathbb{E} [\partial^2 \log p(y_t|x_t; \theta_0) / (\partial\theta\partial\theta')]$, such that \mathcal{I}_T can be chosen as the constant $\mathcal{I} = \text{diag} \{i(\theta)\}$. There is a one-to-one deterministic correspondence between the mapping $\xi \mapsto L_T(\theta_0 + \mathcal{I}_T^{-1/2}\xi)$ and $L_T(\theta)$ and we can restrict our attention to the latter.

To check Assumption C.2, first note that by the uniform LLN $\sup_{\theta \in \Theta} |L_T(\theta) - L(\theta)| = o_P(1)$ with $L(\theta) = \mathbb{E}[\log p(y_t|x_t; \theta)]$ being continuous under Condition (i); see, for example, Kristensen and Rahbek (2005, Proposition 1). Thus,

$$\begin{aligned}
L_T(\tilde{\theta}) - L_T(\theta) &= \{L(\tilde{\theta}) - L(\theta)\} + \{L_T(\tilde{\theta}) - L(\tilde{\theta})\} - \{L_T(\theta) - L(\theta)\} \\
&= \{L(\tilde{\theta}) - L(\theta)\} + o_P(1) = \{L(\theta_0) - L(\theta)\} + o_P(1).
\end{aligned}$$

Since Θ is compact and θ_0 is the unique maximum of $L(\theta)$, for any given $\delta > 0$, there exists $\eta > 0$ such that $\sup_{\|\theta_0 - \theta\| > \eta} \{L(\theta_0) - L(\theta)\} \geq \delta$. This proves that C.2 holds.

To verify C.3, appeal to the Martingale CLT for mixing sequences to obtain:

$$\sqrt{\nu_T} U_T(\theta_0) = \mathcal{I}^{-1/2} \times T^{-1/2} \sum_{t=1}^T \frac{\partial \log p_t(y_t|x_t; \theta)}{\partial \theta} \Big|_{\theta=\theta_0} \xrightarrow{d} N\left(0, \mathcal{I}^{-1/2} i(\theta_0) \mathcal{I}^{-1/2}\right),$$

while by the LLN,

$$V_T(\theta_0) = \mathcal{I}^{-1/2} \times T^{-1} \sum_{t=1}^T \frac{\partial^2 \log p_t(y_t|x_t; \theta)}{\partial \theta \partial \theta'} \Big|_{\theta=\theta_0} \times \mathcal{I}^{-1/2} \xrightarrow{P} -\mathcal{I}^{-1/2} i(\theta_0) \mathcal{I}^{-1/2}.$$

Finally,

$$\max_{j=1, \dots, d} \sup_{\theta \in \Theta} \|W_{j,T}(\theta)\| \leq C \sup_{\theta \in \Theta} \left\| T^{-1} \sum_{t=1}^T \frac{\partial^2 \log p_t(y_t|x_t; \theta)}{\partial \theta \partial \theta'} - i(\theta_0) \right\| + C i(\theta_0) = O_P(1).$$

Condition C.4 holds by another application of the (uniform) LLN:

$$\begin{aligned}
&\sup_{\theta \in \Theta_T} \left| T^{-1} \sum_{t=1}^T |\log p_t(y_t|x_t; \theta)|^{1+\delta} - \mathbb{E} \left[|\log p(y_t|x_t; \theta)|^{1+\delta} \right] \right| \xrightarrow{P} 0, \\
&T^{-1} \sum_{t=1}^T \|x_t\|^{1+\delta} \xrightarrow{P} \mathbb{E} \left[\|x_t\|^{1+\delta} \right], \quad T^{-1} \sum_{t=1}^T \Lambda_1^2(\varepsilon_t) \xrightarrow{P} \mathbb{E} \left[\Lambda_1^2(\varepsilon_t) \right]. \quad \square
\end{aligned}$$

Proof of Theorem 3.3 Since $p(y|x;\theta)$ is bounded away from zero, we can leave out the trimming and re-define the simulated likelihood as $\hat{L}_T(\theta) = T^{-1} \sum_{t=1}^T \log \hat{p}(y_t|x_t;\theta)$. The associated simulated score takes the form

$$\hat{S}_T(\theta) = T^{-1} \sum_{t=1}^T \frac{1}{\hat{p}(y_t|x_t;\theta)} \frac{\partial \hat{p}(y_t|x_t;\theta)}{\partial \theta}.$$

By the mean value theorem, for some $\bar{\theta}$ on the line segment between $\hat{\theta}$ and θ_0 ,

$$0 = \hat{S}_T(\theta_0) + \hat{H}_T(\bar{\theta})(\hat{\theta} - \theta_0).$$

We then analyze the two terms, $\hat{S}_T(\theta_0)$ and $\hat{H}_T(\bar{\theta})$, in turn. Define $p_t(\theta) = p(y_t|x_t;\theta)$, $\hat{p}_t(\theta) = \hat{p}(y_t|x_t;\theta)$ and $s_t(\theta) = \partial \log p_t(\theta) / (\partial \theta)$. Also, we will use the notation $\partial_{\theta} p_t(\theta) = \partial p_t(\theta) / (\partial \theta) \in \mathbb{R}^d$, $\partial_{y_1, \theta}^{\alpha, 1} p_t(\theta) = \partial p_t(\theta) / (\partial y_1^{\alpha} \partial \theta) \in \mathbb{R}^d$, and similarly for other functions. When a function is evaluated at θ_0 , we will frequently suppress the dependence on θ .

We first analyze $\hat{S}_T(\theta_0)$: By the same arguments as in Lee (1999, Proposition 1), the expansion given in equation (18) holds with

$$\begin{aligned} \nabla S_{T,N} [\hat{p} - p] &= \frac{1}{T} \sum_{t=1}^T \left[\frac{1}{p_t} \{ \partial_{\theta} \hat{p}_t - \partial_{\theta} p_t \} - \frac{s_t}{p_t} \{ \hat{p}_t - p_t \} \right], \\ \nabla^2 S_{T,N} [\hat{p} - p, \hat{p} - p] &= \frac{1}{T} \sum_{t=1}^T \left[-\frac{1}{p_t^2} \{ \partial_{\theta} \hat{p}_t - \partial_{\theta} p_t \} \{ \hat{p}_t - p_t \} + \frac{s_t}{p_t^2} \{ \hat{p}_t - p_t \}^2 \right], \\ R_{T,N} &= \frac{1}{T} \sum_{t=1}^T \left[-\frac{1}{\hat{p}_t p_t^2} \{ \partial_{\theta} \hat{p}_t - \partial_{\theta} p_t \} \{ \hat{p}_t - p_t \}^2 + \frac{s_t}{\hat{p}_t p_t^2} \{ \hat{p}_t - p_t \}^3 \right]. \end{aligned}$$

We split up the first-order differential into a bias and a variance component,

$$\nabla S_{T,N} [\hat{p} - p] = \nabla S_{T,N} [\mathbb{E}[\hat{p}|\mathcal{Z}_T] - p] + \nabla S_{T,N} [\hat{p} - \mathbb{E}[\hat{p}|\mathcal{Z}_T]],$$

where $\mathbb{E}[\cdot|\mathcal{Z}_T]$ denotes expectations conditional on data, $\mathcal{Z}_T = \{z_1, \dots, z_T\}$ where $z_t = (y_t, x_t)$. Using standard bias expansions for kernel estimators,

$$\begin{aligned} p(y_{1,t}, y_{2,t}|x_t) - \mathbb{E}[\hat{p}(y_{1,t}, y_{2,t}|x_t)|z_t] &= h^r \sum_{|\alpha|=r} \partial_{y_1}^{\alpha} p(y_{1,t}, y_{2,t}|x_t) + o(h^r), \\ \partial_{\theta} p(y_{1,t}, y_{2,t}|x_t) - \mathbb{E}[\partial_{\theta} \hat{p}(y_{1,t}, y_{2,t}|x_t)|z_t] &= h^r \sum_{|\alpha|=r} \partial_{y_1, \theta}^{\alpha, 1} p(y_{1,t}, y_{2,t}|x_t) + o(h^r), \end{aligned}$$

implying that the bias component satisfies

$$\nabla S_{T,N} [\mathbb{E}[\hat{p}|\mathcal{Z}_T] - p] = h^r T^{-1} \sum_{t=1}^T \sum_{|\alpha|=r} \left[\frac{\partial_{y_1, \theta}^{\alpha, 1} p_t}{p_t} - \frac{s_t}{p_t} \partial_y^{\alpha} p_t \right] + o_P(h^r) = \mu_1 h^r + o_P(h^r),$$

with μ_1 given in equation (22). Next, define

$$\psi_N(z_t, \varepsilon_i) = \frac{1}{p_t} \delta_{N,1}(z_t, \varepsilon_i) - \frac{s_t}{p_t} \delta_{N,2}(z_t, \varepsilon_i),$$

$$\begin{aligned}\delta_{N,1}(z_t, \varepsilon_i) &= \frac{1}{h^{k+1}} \dot{Y}_{1t,i} K^{(1)} \left(\frac{Y_{1t,i} - y_{1t}}{h} \right) Y_{2t,i} + \frac{1}{h^k} \dot{Y}_{2t,i} K \left(\frac{Y_{1t,i} - y_{1t}}{h} \right), \\ \delta_{N,2}(z_t, \varepsilon_i) &= \frac{1}{h^k} K \left(\frac{Y_{1t,i} - y_{1t}}{h} \right) Y_{2t,i}.\end{aligned}$$

With $\psi_{N,1}(z_t) = \mathbb{E}[\psi_N(z_t, \varepsilon_i) | z_t]$, we can then write the variance component as

$$\nabla S_{T,N}[\hat{p} - \mathbb{E}[\hat{p} | \mathcal{Z}_n]] = \frac{1}{TN} \sum_{t=1}^T \sum_{i=1}^N \{\psi_N(z_t, \varepsilon_i) - \psi_{N,1}(z_t)\},$$

which, since \mathcal{Z}_T and $\mathcal{E}_N = \{\varepsilon_1, \dots, \varepsilon_N\}$ are mutually independent, we recognize as a so-called two-sample U -statistic; see, for example, Lehmann (1951). With $\psi_{N,2}(\varepsilon_i) := \mathbb{E}[\psi_N(z_t, \varepsilon_i) | \varepsilon_i]$, and

$$w_N(z_t) := \frac{1}{N} \sum_{i=1}^N \{\psi_N(z_t, \varepsilon_i) - \psi_{N,1}(z_t) - \psi_{N,2}(\varepsilon_i) + \mathbb{E}[\psi(z_t, \varepsilon_i)]\},$$

we can decompose the statistic as:

$$\nabla S_{T,N}[\hat{p} - \mathbb{E}[\hat{p}]] = \frac{1}{N} \sum_{i=1}^N \{\psi_{N,2}(\varepsilon_i) - \mathbb{E}[\psi(z_t, \varepsilon_i)]\} + \frac{1}{T} \sum_{t=1}^T w_N(z_t) =: U_{1,T,N} + W_{1,T,N}.$$

The second term will now be shown to be negligible. Conditioned on \mathcal{E}_N , $\{w_N(z_t) : t = 1, \dots, T\}$ is a stationary and mixing sequence for any given N . Moreover,

$$\begin{aligned}\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T w_N(z_t) | \mathcal{E}_N \right] &= \mathbb{E}[w_N(z_t) | \mathcal{E}_N] \\ &= \frac{1}{N} \sum_{i=1}^N \{\mathbb{E}[\psi_N(z_t, \varepsilon_i) | \varepsilon_i] - \mathbb{E}[\psi_{N,1}(z_t)] - \psi_{N,2}(\varepsilon_i) + \mathbb{E}[\psi(z_t, \varepsilon_i)]\} \\ &= \frac{1}{N} \sum_{i=1}^N \{\psi_{N,2}(\varepsilon_i) - \psi_{N,2}(\varepsilon_i)\} = 0.\end{aligned}$$

This implies that $\mathbb{E}[W_{1,T,N}] = 0$, and, by the law of total variance,

$$\text{Var}(W_{1,T,N}) = \mathbb{E}[\text{Var}(W_{1,T,N} | \mathcal{E}_N)] + \text{Var}(\mathbb{E}[W_{1,T,N} | \mathcal{E}_N]) = \mathbb{E}[\text{Var}(W_{1,T,N} | \mathcal{E}_N)].$$

Here, by standard results for the variance of mixing sequences—see, for example, Kristensen and Salanie (2010, Lemma 5),

$$\text{Var}(W_{1,T,N} | \mathcal{E}_N) \leq \frac{C}{T} \mathbb{E} \left[\|w_N(z_t)\|^{2+\delta} | \mathcal{E}_N \right]^{2/(2+\delta)},$$

for some constant $C > 0$ which only depends on the mixing coefficients of $\{z_t\}$. By standard arguments for i.i.d. sample averages of kernel smoothers,

$$\mathbb{E} \left[\|w_N(z_t)\|^{2+\delta} \right] \leq \frac{C}{N^{1+\delta/2}} \mathbb{E} \left[\|\psi_N(z_t, \varepsilon_i)\|^{2(1+\delta)} \right] = O \left(\frac{1}{N^{1+\delta/2} h^{k+2\delta}} \right).$$

In total, $W_{1,T,N} = O_P \left(1/\sqrt{TNh^{2(k+2\delta)/(2+\delta)}} \right)$. To analyze $U_{1,T,N}$, first note that

$$\mathbb{E} \left[\frac{\delta_{N,1}(z_t, \varepsilon_i)}{p(y_t | x_t)} \middle| \varepsilon_i \right] = \sum_{y_{2t} \in \mathcal{Y}_2} \int \dot{Y}_{1t,i} Y_{2t,i} \left\{ \frac{1}{h^{k+1}} \int K^{(1)} \left(\frac{Y_{1t,i} - y_{1t}}{h} \right) dy_{1t} \right\} dF_x(x_t)$$

$$\begin{aligned}
& + \sum_{y_{2t} \in \mathcal{Y}_2} \int \dot{Y}_{2t,i} \left\{ \frac{1}{h^k} \int K \left(\frac{Y_{1t,i} - y_{1t}}{h} \right) dy_{1t} \right\} dF_x(x_t) \\
& = \mathbb{E} \left[\frac{\dot{Y}_{2t,i}}{p(y_{2t}|x_t)} \middle| \varepsilon_i \right], \tag{29}
\end{aligned}$$

where we have used $h^{-(k+1)} \int K^{(1)} \left(\frac{Y_{1t,i} - y_{1t}}{h} \right) dy_{1t} = h^{-k} \int K^{(1)}(v) dv = 0$. Second,

$$\begin{aligned}
\mathbb{E} \left[\frac{s(y_t|x_t)}{p(y_t|x_t)} \delta_{N,2}(z_t, \varepsilon_i) \middle| \varepsilon_i \right] & = \sum_{y_{2t} \in \mathcal{Y}_2} \int \left\{ \frac{1}{h^k} \int s(y_{1t}, y_{2t}|x_t) K \left(\frac{Y_{1t,i} - y_{1t}}{h} \right) dy_{1t} \right\} Y_{2t,i} dF_x(x_t) \\
& = \sum_{y_{2t} \in \mathcal{Y}_2} \int s(Y_{1t,i}, y_{2t}|x) Y_{2t,i} dF_x(x_t) + O(h^r) \\
& = \mathbb{E} \left[\frac{s(Y_{1t,i}, y_{2t}|x_t) Y_{2t,i}}{p(y_{2t}|x_t)} \middle| \varepsilon_i \right] + O(h^r). \tag{30}
\end{aligned}$$

Thus, $U_{1,T,N} = \sum_{i=1}^N \{\psi_2(\varepsilon_i) - \mathbb{E}[\psi_2(\varepsilon_i)]\} / N + O_P(h^r)$, with $\psi_2(\varepsilon_i)$ defined in equation (20). By the CLT for mixing sequences, $\sqrt{N}U_{1,T,N} \xrightarrow{d} N(0, \text{Var}(\psi_2(\varepsilon_i)))$. This establishes equation (19).

Next, we analyze the second order differential. Write

$$\begin{aligned}
\nabla^2 S_{T,N} [\hat{p} - p, \hat{p} - p] & = \nabla^2 S_{T,N} [\mathbb{E}[\hat{p}|\mathcal{Z}_n] - p, \mathbb{E}[\hat{p}|\mathcal{Z}_n] - p] + \nabla^2 S_{T,N} [\hat{p} - \mathbb{E}[\hat{p}|\mathcal{Z}_n], \hat{p} - \mathbb{E}[\hat{p}|\mathcal{Z}_n]] \\
& \quad + 2\nabla^2 S_{T,N} [\mathbb{E}[\hat{p}|\mathcal{Z}_n] - p, \hat{p} - \mathbb{E}[\hat{p}|\mathcal{Z}_n]].
\end{aligned}$$

Since the cross-term is of a smaller order than the first two ones, we can ignore this. Again using the bias expansion for kernel estimators and appealing to the LLN for stationary and ergodic sequences, the bias component satisfies

$$\begin{aligned}
& \nabla^2 S_{T,N} [\mathbb{E}[\hat{p}|\mathcal{Z}_n] - p, \mathbb{E}[\hat{p}|\mathcal{Z}_n] - p] \\
& = \frac{1}{T} \sum_{t=1}^T \left[-\frac{1}{p_t^2} \{ \mathbb{E}[\partial_\theta \hat{p}_t | z_t] - \partial_\theta p_t \} \{ \mathbb{E}[\hat{p}_t | z_t] - p_t \} + \frac{s_t}{p_t^2} \{ \mathbb{E}[\hat{p}_t | z_t] - p_t \}^2 \right] = O_P(h^{2r}).
\end{aligned}$$

The variance component can be written as

$$\begin{aligned}
\nabla^2 S_{T,N} [\hat{p} - \mathbb{E}[\hat{p}|\mathcal{Z}_n], \hat{p} - \mathbb{E}[\hat{p}|\mathcal{Z}_n]] & = \frac{1}{TN^2} \sum_{t=1}^T \sum_{i=1}^N \phi_N(z_t, \varepsilon_i, \varepsilon_i) + \frac{1}{TN^2} \sum_{t=1}^T \sum_{i \neq j} \phi_N(z_t, \varepsilon_i, \varepsilon_j) \\
& = : \frac{1}{N} U_{2,T,N} + W_{2,T,N},
\end{aligned}$$

with

$$\begin{aligned}
\phi_N(z_t, \varepsilon_i, \varepsilon_j) & = -\frac{1}{p_t^2} \{ \delta_{N,1}(z_t, \varepsilon_i) - \mathbb{E}[\delta_{N,1}(z_t, \varepsilon_i) | z_t] \} \{ \delta_{N,2}(z_t, \varepsilon_j) - \mathbb{E}[\delta_{N,2}(z_t, \varepsilon_j) | z_t] \} \\
& \quad + \frac{s_t}{p_t^2} \{ \delta_{N,2}(z_t, \varepsilon_i) - \mathbb{E}[\delta_{N,2}(z_t, \varepsilon_i) | z_t] \} \{ \delta_{N,2}(z_t, \varepsilon_j) - \mathbb{E}[\delta_{N,2}(z_t, \varepsilon_j) | z_t] \}.
\end{aligned}$$

The first term, $U_{2,T,N}$, is again a second order two-sample U -statistic while $W_{2,T,N}$ is a third order one. To analyze $U_{2,T,N}$, we proceed in the same manner as with $\nabla S_{T,N}[\hat{p} - \mathbb{E}[\hat{p}]]$. By the Hoeffding decomposition,

$$U_{2,T,N} = \mathbb{E}[\phi_N(z_t, \varepsilon_i, \varepsilon_i)] + \frac{1}{N} \sum_{i=1}^N \{ \phi_{N,1}(\varepsilon_i) - \mathbb{E}[\phi_{N,1}(\varepsilon_i)] \} + \frac{1}{T} \sum_{t=1}^T \{ \phi_{N,2}(z_t) - \mathbb{E}[\phi_{N,2}(z_t)] \} + \tilde{R}_{T,N},$$

where $\phi_{N,1}(\varepsilon_i) = \mathbb{E}[\phi_N(z_t, \varepsilon_i, \varepsilon_i) | \varepsilon_i]$, $\phi_2(z_t) = \mathbb{E}[\phi_N(z_t, \varepsilon_i, \varepsilon_i) | z_t]$ and $\tilde{R}_{T,N}$ is the remainder term. By the same arguments as used for $W_{1,T,N}$, it follows that $\tilde{R}_{T,N} = O_P\left(1/\sqrt{TN}h^{k+\delta}\right)$, while by the CLT for stationary and mixing sequences the two sample averages in the above expression are of order $O_P\left(1/\sqrt{T}\right)$ and $O_P\left(1/\sqrt{N}\right)$ respectively. Using the same arguments as before,

$$\begin{aligned} \mathbb{E}\left[\frac{1}{p_t^2}\delta_{N,1}(z_t, \varepsilon_i)\delta_{N,2}(z_t, \varepsilon_i)\right] &\simeq \frac{1}{h^{k+1}}\mathbb{E}\left[\frac{\dot{Y}_{1t,i}Y_{2t,i}^2}{p(Y_{1t,i}, y_{2t}|x_t)p(y_{2t}|x_t)}\Bigg|\varepsilon_i\right]\int K^*(v)dv \\ &\quad + \frac{1}{h^k}\mathbb{E}\left[\frac{Y_{2t,i}\dot{Y}_{2t,i}}{p(Y_{1t,i}, y_{2t}|x_t)p(y_{2t}|x_t)}\Bigg|\varepsilon_i\right]\int K^2(v)dv, \end{aligned}$$

where $K^*(v) = K(v)K^{(1)}(v)$, and

$$\mathbb{E}\left[\frac{s_t}{p_t^2}\delta_{N,2}^2(z_t, \varepsilon_i) | \varepsilon_i\right] \simeq \frac{1}{h^k}\mathbb{E}\left[\frac{Y_{2t,i}\dot{Y}_{2t,i}}{p(Y_{1t,i}, y_{2t}|x_t)p(y_{2t}|x_t)}\Bigg|\varepsilon_i\right]\int K^2(v)dv,$$

where we have left out higher order terms. Thus,

$$\mathbb{E}[\phi_N(z_t, \varepsilon_i, \varepsilon_i)] = h^{-(k+1)}\mu_2 + h^{-k}\tilde{\mu}_2 + o\left(h^{-(k+1)}\right),$$

where μ_2 is given in equation (23) and $\tilde{\mu}_2$ is the sum of the two other expectations above. Next,

$$\mathbb{E}[\phi_N(z_t, \varepsilon_i, \varepsilon_j) | z_t, \varepsilon_i] = \mathbb{E}[\phi(z_t, \varepsilon_i, \varepsilon_j) | z_t, \varepsilon_j] = 0,$$

while $\mathbb{E}[\phi_N(z_t, \varepsilon_i, \varepsilon_j) | \varepsilon_i, \varepsilon_j] = o(1)$ as $h \rightarrow 0$ such that the corresponding U -statistic $W_{2,T,N}$ is second-order degenerate, implying $W_{2,T,N} = O_P\left(1/(N^{3/2}h^{k+1})\right)$. This establishes equation (21).

Finally, appealing to Lemma B.2 and the LLN, the remainder term of the expansion of $\hat{S}_T(\theta_0)$ satisfies

$$\begin{aligned} \sqrt{T}\|R_{T,N}\| &\leq T^{-1/2}\sum_{t=1}^T\frac{1}{|\hat{p}_t|p_t^2}\|\partial_\theta\hat{p}_t - \partial_\theta p_t\|\{\hat{p}_t - p_t\}^2 + T^{-1/2}\sum_{t=1}^T\frac{s_t}{|\hat{p}_t|p_t^2}|\hat{p}_t - p_t|^3 \\ &= O_P\left(\sqrt{T}h^{3r}\right) + O_P\left(\sqrt{T}/(Nh^{k+2})^{3/2}\right). \end{aligned}$$

Next, we consider the simulated Hessian:

$$\hat{H}_T(\theta) = \frac{1}{T}\sum_{t=1}^T\frac{\partial_{\theta\theta}^2\hat{p}_t}{\hat{p}_t} + \frac{1}{T}\sum_{t=1}^T\frac{\partial_\theta\hat{p}_t\partial_\theta\hat{p}_t'}{\hat{p}_t^2}.$$

We write:

$$\|\hat{H}_T(\bar{\theta}) - i(\theta_0)\| \leq \sup_{\theta \in \Theta}\|\hat{H}_T(\theta) - H_T(\theta)\| + \sup_{\theta \in \Theta}\|H_T(\theta) - i(\theta)\| + \|i(\bar{\theta}) - i(\theta_0)\|.$$

As shown in the proof of Corollary 3.2, the second and third terms are both $o_P(1)$. Regarding the first term, write

$$\hat{H}_T(\theta) - H_T(\theta) = \frac{1}{T}\sum_{t=1}^T\left\{\frac{\partial_{\theta\theta}^2\hat{p}_t(\theta)}{\hat{p}_t(\theta)} - \frac{\partial_{\theta\theta}^2 p_t(\theta)}{p_t(\theta)}\right\} + \frac{1}{T}\sum_{t=1}^T\left\{\frac{\partial_\theta\hat{p}_t(\theta)\partial_\theta\hat{p}_t(\theta)'}{\hat{p}_t^2(\theta)} - \frac{\partial_\theta p_t(\theta)\partial_\theta p_t(\theta)'}{p_t^2(\theta)}\right\}$$

$$=: A_1(\theta) + A_2(\theta).$$

The first term, $A_1(\theta)$, satisfies uniformly in $\theta \in \Theta$:

$$\begin{aligned} \|A_1(\theta)\| &\leq \frac{1}{T} \sum_{t=1}^T \frac{1}{p_t(\theta)} \left\| 1 - \frac{1}{\hat{p}_t(\theta)} (\hat{p}_t(\theta) - p_t(\theta)) \right\| \left\| \left\{ \partial_{\theta\theta}^2 \hat{p}_t(\theta) - \partial_{\theta\theta}^2 p_t(\theta) \right\} - \frac{\partial_{\theta\theta}^2 p_t(\theta)}{p_t(\theta)} \{ \hat{p}_t(\theta) - p_t(\theta) \} \right\| \\ &\leq \frac{1}{T} \sum_{t=1}^T \frac{1 + o_P(1)}{p_t(\theta)} \left\| \partial_{\theta\theta}^2 \hat{p}_t(\theta) - \partial_{\theta\theta}^2 p_t(\theta) \right\| + \frac{1}{T} \sum_{t=1}^T \frac{\left\| \partial_{\theta\theta}^2 p_t(\theta) \right\| (1 + o_P(1))}{p_t^2(\theta)} |\hat{p}_t(\theta) - p_t(\theta)| \\ &= (1 + o_P(1)) \left\{ \frac{1}{T} \sum_{t=1}^T \frac{1}{p_t(\theta)} \right\} \sup_{1 \leq t \leq T} \left\| \partial_{\theta\theta}^2 \hat{p}_t(\theta) - \partial_{\theta\theta}^2 p_t(\theta) \right\| \\ &\quad + (1 + o_P(1)) \left\{ \frac{1}{T} \sum_{t=1}^T \frac{\left\| \partial_{\theta\theta}^2 p_t(\theta) \right\|}{p_t^2(\theta)} \right\} \sup_{1 \leq t \leq T} |\hat{p}_t(\theta) - p_t(\theta)| = o_P(1), \end{aligned}$$

where we have used the LLN together with $|\hat{p}_t(\theta) - p_t(\theta)| = o_P(1)$, $\left\| \partial_{\theta\theta}^2 \hat{p}_t(\theta) - \partial_{\theta\theta}^2 p_t(\theta) \right\| = o_P(1)$ and $p_t(\theta) > 0$ uniformly over $1 \leq t \leq T$ and $\theta \in \Theta$. We can show that $\sup_{\theta \in \Theta} \|A_2(\theta)\| = o_P(1)$ by the same arguments. The claimed result now follows by combining the results obtained for the simulated score and Hessian. \square

Proof of Theorem 3.4 We follow the exact same arguments as in the proof of Theorem 3.3, except that we now have

$$|\hat{p}_{M,t}(y_t|x_t; \theta) - p_t(y_t|x_t; \theta)| \leq |\hat{p}_{M,t}(y_t|x_t; \theta) - \hat{p}_t(y_t|x_t; \theta)| + |\hat{p}_t(y_t|x_t; \theta) - p_t(y_t|x_t; \theta)|$$

where $\hat{p}_t(y|x; \theta)$ is the kernel estimator based on simulations from the true model. The uniform rate of convergence of the second term is given by Lemma B.1, while the first term satisfies by M.1,

$$\begin{aligned} |\hat{p}_{M,t}(y_t|x_t; \theta) - \hat{p}_t(y_t|x_t; \theta)| &\leq \frac{1}{N} \sum_{i=1}^N \left| K_h(\hat{Y}_{1t,i}^\theta - y_{1t}) \hat{Y}_{2t,i}^\theta - K_h(Y_{1t,i}^\theta - y_{1t}) Y_{2t,i}^\theta \right| \\ &\leq \frac{1}{Nh} \sum_{i=1}^N \left| K_h^{(1)}(\bar{Y}_{1t,i}^\theta - y_{1t}) \right| \left| \hat{Y}_{2t,i}^\theta \right| \left\| \hat{Y}_{1t,i}^\theta - Y_{1t,i}^\theta \right\| \\ &\quad + \frac{1}{N} \sum_{i=1}^N \left| K_h(Y_{1t,i}^\theta - y_{1t}) \right| \left| \hat{Y}_{2t,i}^\theta - Y_{2t,i}^\theta \right| \\ &\leq \left\{ \frac{\delta_{M,1}}{h} B_1 \left(1 + \|x_t\|^{\lambda_{3,1}} + t^{\lambda_{3,2}} \right) + \delta_{M,2} B_2 \left(1 + \|x_t\|^{\lambda_{4,1}} + t^{\lambda_{4,2}} \right) \right\} O_P(1), \end{aligned}$$

where $\bar{Y}_{1t,i}^\theta$ lies on the line segment between $\hat{Y}_{1t,i}^\theta$ and $Y_{1t,i}^\theta$. Thus, uniformly over $y_t \in \mathbb{R}^k$, $\|x_t\| \leq N^\gamma$ and $\theta \in \Theta$,

$$\frac{T}{a\nu_T} \times |\hat{p}_{M,t}(y_t|x_t; \theta) - \hat{p}_t(y_t|x_t; \theta)| = O_P \left(\frac{T}{a\nu_T} \frac{\delta_{M,1}}{h} \left[N^{\gamma\lambda_{3,1}} + T^{\lambda_{3,2}} \right] \right) + O_P \left(\frac{T}{a\nu_T} \delta_{M,2} \left[N^{\gamma\lambda_{4,1}} + T^{\lambda_{4,2}} \right] \right),$$

where the right-hand side has to go to zero to obtain consistency. This holds by B.1'. For first-order equivalence, we require that the right-hand side vanish with rate $\sqrt{\nu_T}$. This holds by B.2'. \square

References

- AÏT-SAHALIA, Y. (2002): “Maximum Likelihood Estimation of Discretely Sampled Diffusions: A Closed-Form Approximation Approach,” *Econometrica*, 70, 223–262.
- (2008): “Closed-Form Likelihood Expansions for Multivariate Diffusions,” *Annals of Statistics*, 36, 906–937.
- AÏT-SAHALIA, Y. AND R. KIMMEL (2007): “Maximum Likelihood Estimation of Stochastic Volatility Models,” *Journal of Financial Economics*, 83, 413–452.
- ALTISSIMO, F. AND A. MELE (2009): “Simulated Nonparametric Estimation of Dynamic Models,” *Review of Economic Studies*, 76, 413–450.
- AMEMIYA, T. (1983): *Advanced Econometrics*, Massachusetts: Harvard University Press.
- ANDERSEN, T. G., L. BENZONI, AND J. LUND (2002): “An Empirical Investigation of Continuous-Time Equity Return Models,” *Journal of Finance*, 57, 1239–1284.
- ANDREWS, D. W. K. (1995): “Nonparametric Kernel Estimation for Semiparametric Models,” *Econometric Theory*, 11, 560–596.
- ANGO NZE, P. AND P. DOUKHAN (2004): “Weak Dependence: Models and Applications to Econometrics,” *Econometric Theory*, 20, 995–1045.
- BANDI, F. M. AND P. C. B. PHILLIPS (2003): “Fully Nonparametric Estimation of Scalar Diffusion Models,” *Econometrica*, 71, 241–283.
- BASAWA, I. V. AND D. J. SCOTT (1983): *Asymptotic Optimal Inference for Non-Ergodic Models*, New York: Springer-Verlag.
- BICHTELER, K., J.-B. GRAVERAUX, AND J. JACOD (1987): *Mallivin Calculus for Processes with Jumps*, Gordon and Breach Science Publishers.
- BRANDT, M. W. AND P. SANTA-CLARA (2002): “Simulated Likelihood Estimation of Diffusions with an Application to Exchange Rates Dynamics in Incomplete Markets,” *Journal of Financial Economics*, 63, 161–210.
- BROWNLEES, C., D. KRISTENSEN, AND Y. SHIN (2011): “Smooth Filtering and Maximum-Likelihood Estimation of Dynamic Latent Variable Models,” Manuscript, Department of Economics, Columbia University.
- BRUTI-LIBERATI, N. AND E. PLATEN (2007): “Approximation of Jump Diffusions in Finance and Economics,” *Computational Economics*, 29, 283–312.
- CAI, Z., Q. YAO, AND W. ZHANG (2001): “Smoothing for Discrete-Valued Time Series,” *Journal of the Royal Statistical Society (B)*, 63, 357–375.
- CARRASCO, M., M. CHERNOV, J.-P. FLORENS, AND E. GHYSELS (2007): “Efficient estimation of general dynamic models with a continuum of moment conditions,” *Journal of Econometrics*, 140, 529–573.

- COX, J. C., J. E. INGERSOLL, JR., AND S. A. ROSS (1985): “A Theory of the Term Structure of Interest Rates,” *Econometrica*, 53, 385–408.
- CREEL, M. AND D. KRISTENSEN (2009): “Estimation of Dynamic Latent Variable Models Using Simulated Nonparametric Moments,” Working Paper 792.09, UFAE and IAE.
- (2011): “Indirect Likelihood Inference,” Manuscript, Department of Economics, Columbia University.
- DETEMPLE, J., R. GARCIA, AND M. RINDISBACHER (2006): “Asymptotic Properties of Monte Carlo Estimators of Diffusion Processes,” *Journal of Econometrics*, 134, 1–68.
- DONOHO, D. L., I. M. JOHNSTONE, G. KERKYACHARIAN, AND D. PICARD (1996): “Density Estimation by Wavelet Tresholding,” *Annals of Statistics*, 24, 508–539.
- DORASZELSKI, U. AND K. JUDD (2008): “Avoiding the Curse of Dimensionality in Dynamic Stochastic Games,” Manuscript, Stanford University.
- DORASZELSKI, U. AND A. PAKES (2007): “A Framework for Applied Dynamic Anaysis in IO,” in *Handbook of Industrial Organization*, ed. by M. Armstrong and R. Porter, Amsterdam: Elsevier B.V., vol. 3, 1887–1966.
- DUFFIE, D. AND K. J. SINGLETON (1993): “Simulated Moments Estimation of Markov Models of Asset Prices,” *Econometrica*, 61, 929–952.
- DURHAM, G. B. AND A. R. GALLANT (2002): “Numerical Techniques for Maximum Likelihood Estimation of Continuous-Time Diffusion Processes,” *Journal of Business and Economic Statistics*, 20, 297–316.
- ELERIAN, O., S. CHIB, AND N. SHEPHARD (2001): “Likelihood Inference for Discretely Observed Non-linear Diffusions,” *Econometrica*, 69, 959–993.
- ERICSON, R. AND A. PAKES (1995): “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62, 53–82.
- FENTON, V. M. AND A. R. GALLANT (1996): “Notes and Comments: Convergence Rates of SNP Density Estimators,” *Econometrica*, 64, 719–727.
- FERMANIAN, J.-D. AND B. SALANIÉ (2004): “A Nonparametric Simulated Maximum Likelihood Estimation Method,” *Econometric Theory*, 20, 701–734.
- GALLANT, A. R. AND D. W. NYCHKA (1987): “Semi-Nonparametric Maximum Likelihood Estimation,” *Econometrica*, 55, 363–390.
- GALLANT, A. R. AND G. TAUCHEN (1996): “Which Moments to Match?” *Econometric Theory*, 12, 657–681.
- GOURIÉROUX, C., A. MONFORT, AND É. RENAULT (1993): “Indirect Inference,” *Journal of Applied Econometrics*, 8, S85–S118.

- HAJIVASSILIOU, V. A. AND D. L. MCFADDEN (1998): “The Method of Simulated Scores for the Estimation of LDV Models,” *Econometrica*, 66, 863–896.
- ICHIMURA, H. AND P. E. TODD (2007): “Implementing Nonparametric and Semiparametric Estimators,” in *Handbook of Econometrics*, ed. by J. J. Heckman and E. E. Leamer, Amsterdam: North Holland, vol. 6, 5369–5468.
- JEGANATHAN, P. (1995): “Some Aspects of Asymptotic Theory with Applications to Time Series Models,” *Econometric Theory*, 11, 818–887.
- JONES, M., J. MARRON, AND S. SHEATHER (1996): “A Brief Survey of Bandwidth Selection for Density Estimation,” *Journal of the American Statistical Association*, 91, 401–407.
- KARLSEN, H. A. AND D. TJØSTHEIM (2001): “Nonparametric Estimation in Null Recurrent Time Series,” *Annals of Statistics*, 29, 372–416.
- KLOEDEN, P. E. AND E. PLATEN (1992): *Numerical Solution of Stochastic Differential Equations*, New York: Springer-Verlag.
- KRISTENSEN, D. (2009): “Uniform Convergence Rates of Kernel Estimators with Heterogeneous, Dependent Data,” *Econometric Theory*, 25, 1433–1445.
- (2010): “Pseudo-Maximum-Likelihood Estimation in Two Classes of Semiparametric Diffusion Models,” *Journal of Econometrics*, 156, 239–259.
- KRISTENSEN, D. AND A. RAHBK (2005): “Asymptotics of the QMLE for a Class of ARCH(q) Models,” *Econometric Theory*, 21, 946–961.
- (2010): “Likelihood-Based Inference for Cointegration with Nonlinear Error-Correction,” *Journal of Econometrics*, 158, 78–94.
- KRISTENSEN, D. AND B. SALANIE (2010): “Higher Order Improvements for Approximate Estimators,” CAM Working Paper 2010-04, Department of Economics, University of Copenhagen.
- KRISTENSEN, D. AND B. SCHJERNING (2011): “Implementation and Estimation of Discrete Markov Decision Models By Sieve Approximations,” Tech. rep., University of Copenhagen.
- LEE, B.-S. AND B. F. INGRAM (1991): “Simulation Estimation of Time-Series Models,” *Journal of Econometrics*, 47, 197–205.
- LEE, D. AND K. SONG (2009): “Simulated MLE for Discrete Choices using Transformed Simulated Frequencies,” Manuscript, Department of Economics, University of Pennsylvania.
- LEE, L.-F. (1992): “On Efficiency of Methods of Simulated Moments and Maximum Simulated Likelihood Estimation of Discrete Response Models,” *Econometric Theory*, 8, 518–552.
- (1995): “Asymptotic Bias in Simulated Maximum Likelihood Estimation of Discrete Choice Models,” *Econometric Theory*, 11, 437–483.
- (1999): “Statistical Inference with Simulated Likelihood Functions,” *Econometric Theory*, 15, 337–360.

- LEE, S.-W. AND B. E. HANSEN (1994): “Asymptotic Theory for the GARCH(1,1) Quasi-Maximum Likelihood Estimator,” *Econometric Theory*, 10, 29–52.
- LEHMANN, E. (1951): “Consistency and Unbiasedness of Certain Nonparametric Tests,” *Annals of Mathematical Statistics*, 22, 165–179.
- LI, Q. AND J. RACINE (2007): *Nonparametric Econometrics: Theory and Practice*, New Jersey: Princeton University Press.
- LO, A. W. (1988): “Maximum Likelihood Estimation of Generalized Itô Processes with Discretely Sampled Data,” *Econometric Theory*, 4, 231–247.
- MANRIQUE, A. AND N. SHEPHARD (1998): “Simulation-Based Likelihood Inference for Limited Dependent Processes,” *Econometrics Journal*, 1, 174–202.
- PARK, J. Y. AND P. C. B. PHILLIPS (2000): “Nonstationary Binary Choice Models,” *Econometrica*, 69, 117–161.
- (2001): “Nonlinear Regressions with Integrated Time Series,” *Econometrica*, 68, 1249–1280.
- PEDERSEN, A. R. (1995a): “Consistency and Asymptotic Normality of an Approximate Maximum Likelihood Estimator for Discretely Observed Diffusion Processes,” *Bernoulli*, 1, 257–279.
- (1995b): “A New Approach to Maximum-Likelihood Estimation for Stochastic Differential Equations Based on Discrete Observations,” *Scandinavian Journal of Statistics*, 22, 55–71.
- PHILLIPS, P. C. B. (1983): “ERA’s: A New Approach to Small Sample Theory,” *Econometrica*, 51, 1505–1525.
- RUST, J. (1994): “Structural Estimation of Markov Decision Processes,” in *Handbook of Econometrics*, ed. by R. F. Engle and D. L. McFadden, Amsterdam: North Holland, vol. 4, 3081–3143.
- SAIKKONEN, P. (1995): “Problems with the Asymptotic Theory of Maximum Likelihood Estimation in Integrated and Cointegrated Systems,” *Econometric Theory*, 11, 888–911.
- SANDMANN, G. AND S. J. KOOPMAN (1998): “Estimation of Stochastic Volatility Models via Monte Carlo Maximum Likelihood,” *Journal of Econometrics*, 87, 271–301.
- SCHAUMBURG, E. (2001): “Maximum Likelihood Estimation of Lévy Type SDEs,” Ph.D. thesis, Princeton University.
- SCOTT, D. W. (1992): *Multivariate Density Estimation: Theory, Practice, and Visualization*, New York: Wiley.
- SMITH, JR., A. (1993): “Estimating Nonlinear Time-Series Models Using Simulated Vector Autoregressions,” *Journal of Applied Econometrics*, 8, S63–S84.
- STONE, C. J. (1990): “Large-Sample Inference for Log-Spline Models,” *Annals of Statistics*, 18, 717–741.

- SUNDARESAN, S. M. (2000): “Continuous-Time Methods in Finance: A Review and an Assessment,” *Journal of Finance*, 55, 1569–1622.
- TAUCHEN, G. (1997): “New Minimum Chi-Square Methods in Empirical Finance,” in *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress*, ed. by D. M. Kreps and K. F. Wallis, Cambridge: Cambridge University Press, 279–317.
- WAHBA, G. (1981): “Data-Based Optimal Smoothing of Orthogonal Series Density Estimates,” *Annals of Statistics*, 9, 146–156.
- WHITE, JR., H. L. (1984): “Maximum Likelihood Estimation of Misspecified Dynamic Models,” in *Misspecification Analysis*, ed. by T. K. Dijkstra, New York: Springer-Verlag, 1–19.
- (1994): *Estimation, Inference and Specification Analysis*, Cambridge: Cambridge University Press.
- YU, J. (2007): “Closed-Form Likelihood Approximation and Estimation of Jump-Diffusions with an Application to the Realignment Risk of the Chinese Yuan,” *Journal of Econometrics*, 141, 1245–1280.