ON THE PERFORMANCE OF THE UNITED STATES NUCLEAR POWER SECTOR: A BAYESIAN APPROACH

DAVID H. BERNSTEIN[1], CHRISTOPHER F. PARMETER[1], AND MIKE G. TSIONAS[2]

ABSTRACT. Concerns over climate change and global emissions has again placed attention on clean energy sources. Nuclear power plants are one of many sources of clean energy and yet few studies have examined the structure of technology exclusively in this area. We utilize Bayesian empirical likelihood methods to estimate a stochastic frontier model to examine scale economies, technical efficiency and technological change in the United States nuclear energy generation sector. We find decreasing scale economies, a fact consistent with the recent decline of the industry. Our results suggest that small nuclear reactors may benefit the sector as a whole.

Key Words: Nuclear Energy, Small Nuclear Reactor, Returns to Scale, Exponential Tilting, Asymmetric Laplace, Empirical Likelihood.

1. INTRODUCTION

A major hurdle towards building new nuclear generation is the ability to build it at a small enough scale to meet regional energy demand while keeping costs competitive with more traditional sources of energy. U.S. Secretary of Energy Jennifer Granholm noted at a small nuclear reactor (SMR) partnership meeting alongside the United Nations COP26 that the United States "have been supporting the development of SMRs for decades..." and "views nuclear energy as a pivotal technology in the global effort to ... ultimately combat

1: UNIVERSITY OF MIAMI; 2: LANCASTER UNIVERSITY
David Bernstein, Department of Economics, University of Miami, Miami, FL 33146; e-mail: dbernstein@bus.miami.edu. Christopher F. Parmeter, Corresponding Author, Department of Economics, University of Miami, Miami, FL 33146; e-mail: cparmeter@bus.miami.edu. Mike Tsionas, Department of Economics, Lancaster University; e-mail: m.tsionas@lancaster.ac.uk. The views expressed here are those of the authors and not necessarily those of the Federal Energy Regulatory Commission or the U.S. Government. This paper was previously circulated under the title "On Bayesian Quantile Estimation of the Stochastic Frontier Model". We thank participants at vNAPW XI for useful comments which improved the paper. The usual caveat applies.

climate change"(DOE, 2021). These comments and historic nuclear investments made in the Inflation Reduction Act of 2022 highlight the importance of examining the returns to scale of the U.S. nuclear fleet since it represents a crucial component of nuclear power costs.

Assessing returns to scale at the plant level could occur either through the size of the nuclear power generators themselves or through the siting of different units. For example, if focusing on scale economies based on the size of the generator, this would suggest new power plants should deploy **larger** generators to provide cheaper power. Alternatively, achieving scale economies by siting multiple generators at the same site (common in several European countries) would suggest providing cheaper energy costs by building more generators. In either setting, understanding how scale economies arise (if at all) is of tremendous importance for the nuclear power sector, and yet, outside of Krautmann & Solow (1988), there has been little research into the presence of returns to scale in the nuclear power industry.

Another concern with the production of nuclear energy (and electricity production more broadly) is if it is produced efficiently. Even with a lack of scale economies, cost savings relative to more traditional production could be realized if nuclear energy producers were more efficient than their competitors. Despite the importance of this issue, to our knowledge, no empirical evidence exists that studies technical efficiency in the nuclear power industry directly even though there is ample evidence on technical efficiency in the coal and gas powered electricity generation industry including Greene (1990), Knittel (2002), Rungsuriyawiboon & Stefanou (2007), Greene (2008), Chen, Schmidt & Wang (2014), Ghosh & Kathuria (2016), Filippini, Geissmann & Greene (2018) and Bernstein (2020).

Previous studies of technical efficiency of electricity generating nuclear power plants have always been in the context of a broad array of power plants (coal-fired, natural gas, hydro, nuclear). However, given the fundamental difference in heat formation at pressurized-water and boiling-water nuclear reactors, these plants clearly operate a different technology than say a combined cycle gas plant. This makes these types of comparisons dubious. Here we focus our attention exclusively on nuclear power plants. Further, use of a dummy variable to signify type of power generation to assess differences in technical efficiency is of little value when the technologies themselves differ across types of power generation.

The study of technical efficiency is commonly carried out using either data envelopment analysis (DEA, Simar & Wilson 2013) or stochastic frontier analysis (SFA, Parmeter & Kumbhakar 2014). Our focus will be with SFA, as these methods allow direct modelling of stochastic noise. This is important since nuclear power plants go down unexpectedly due to accidents, equipment failure and random safety shutdowns. Furthermore, the regulatory environment in which nuclear power plants operate also introduces uncertainties that are difficult to measure from an operational standpoint and are best captured with a random error. Utilities must make choices about input use in the face of uncertain regulatory impacts on the level of output, which may lead to optimization errors which produces additional noise in the production process.

It is likely that in the presence of these types of stochastic impacts, that the error term in our model is likely to be nonstandard. To allow for this case we construct a novel estimation approach that allows the error term in the production process to be asymmetric, rather than the commonly assumed symmetric noise.[1] This has many plausible benefits, one being the avoidance of the wrong skewness (Simar & Wilson 2010) that commonly plagues stochastic frontier modelling.

Asymmetry in the noise term is modeled through an assumption of Asymmetric Laplace, where it has recently been shown by Bera, Galvao, Montes-Rojas & Park (2016) that the first order conditions stemming from maximum likelihood using the Asymmetric Laplace distribution is equivalent to the solution stemming from a maximum entropy problem with

[1]Other papers that have explored this feature include Bonanno, De Giovanni & Domma (2015), Badunenko & Henderson (2023), Wei, Zhu & Wang (2021) and Horrace, Parmeter & Wright (2023).

moment constraints imposed. Further, the score functions from the Asymmetric Laplace provide estimators for the slope parameters along with a representative quantile. These results are not dependent on the error term actually having an Asymmetric Laplace distribution; in this sense the maximum likelihood estimator can be viewed as robust to misspecification of the error term. Along this lines, both Tsionas (2020) and Tsionas, Assaf & Andrikopoulos (2020) estimate the stochastic frontier model allowing the noise to be distributed as Asymmetric Laplace using standard Bayesian methods.[2] Our goal here is to consider the same specification but estimate the parameters using novel Bayesian exponentially tilted empirical likelihood (Schennach 2005, BETEL, hereafter). Use of BETEL also imparts robustness to potential outliers in the estimation itself, another common concern in empirical work involving the stochastic frontier model and also within the nuclear power generation setting.

Here we adapt the work of Lancaster & Jun (2010) (who in turn adapt Schennach's (2005) BETEL), so that one can estimate a 'quantile like' stochastic frontier model (see also Yang & He 2012). To provide Bayesian inference, we focus on the first-order conditions of maximum likelihood, similar to Komunjer's (2005) modifications to standard likelihood procedures. These first-order conditions involve orthogonality of regressors to certain generalized residuals and they are, in fact, moment conditions, a result that paves the way to application of generalized method of moments (GMM), empirical likelihood and, therefore, BETEL methods. To our knowledge this marks the first use of BETEL methods in the setting of the stochastic frontier model and is a marked departure from previous approaches that have relied exclusively on Asymmetric Laplace prior distributions in more standard Bayesian setups.

[2]Estimation of the stochastic frontier model using Bayesian approaches has become quite common in the energy field. See Chen, Barros & Borges (2015), Li et al. (2016), Makiela & Osiewalski (2018) and Haider & Mishra (2021) for recent examples.

We then apply these methods to an unbalanced panel of 58 US nuclear power plants over the 1994-2016 period. We find that estimated returns to scale are significantly less than 1 for a majority of the plants in the sample, that the modal plant is near fully efficient, and that technical change is nearly evenly distributed around 0. These results are consistent with the extant literature. As noted by Krautmann & Solow (1988), the basic technology underlying nuclear power generation has not changed much since the 1960s, even though the regulatory environment has. Thus, a finding of very little average technical change is not surprising. The finding of mainly decreasing returns to scale is in line with the work of Rungsuriyawiboon (2008), albeit for a different time period. Furthermore, our new approach delivers estimates that are quite different than a standard Bayesian setting, suggesting that the BETEL estimation approach is indeed imparting its impact on the empirical findings. Given our earlier discussion on how these nuclear plants are run it is not surprising to find such high levels of technical efficiency.

All of these results speak to an industry that is operating near the frontier, but the lack of updated technology or new plants coming on line suggests that further increases in nuclear energy generation are not likely coming unless new plants are opened. This is consistent with EIA's projection for nuclear energy production declining by 17% by 2025 (EIA, 2019). A key policy question, and one tied to our work here, is if new plants should be founded upon a single large generator or a multi-generator framework. Given the focus on reducing dependency on fossil-fuel generated electricity, the potential expansion for nuclear energy is of first order importance.

The remainder of the article is organized as follows. Section 2 provides background information on the nuclear power sector. Section 3 describes BETEL estimation of the stochastic frontier model through both asymmetric Laplace priors as well as BETEL and describes key implementation issues. Section 4 provides our nuclear powered electricity generation estimates. Section 5 contains concluding remarks and avenues for future research.

2. Background of Nuclear Power Generation in the United States

In 2017, 38% of all energy consumed in the United States was electric power (EIA, 2018) and of all electric energy produced in the U.S in 2018, the second largest source, after fossil fuels[3], came from nuclear with 19.3% (807 billion kWh) of total electricity generation (EIA, 2018). Nuclear power harnesses the energy of a fissile material (such as Uranium-235) through a process called nuclear fission whereby the splitting of atoms generates heat for the steam generation process. Nuclear fission emits no CO_2 whereas the burning of fossil fuels is responsible for 99% of CO_2 emissions for electricity (EIA, 2018).

As Davis (2012) notes "Nuclear power continues to generate enthusiasm based on its potential to reduce greenhouse gas emissions. A single pound of reactor-grade uranium oxide produces as much electricity as over 16,000 pounds of coal – enough to meet the needs of the average U.S. household for more than a year. While burning 16,000 pounds of coal generates thousands of pounds of carbon dioxide, sulfur dioxide, and nitrogen oxides nuclear power is virtually emissions-free." Yet, expansion of the nuclear energy generation market is virtually nonexistent. Only several new applications[4] for siting nuclear reactors has occurred over the last several decades while a large number of nuclear power plants either have already been taken off the grid or are planned to be removed in the next several years.[5] These features of the nuclear power generation industry are important as they speak to understanding how efficient the plants in operation are, given the focus on carbon emissions reduction and the substantial portion of the electricity market that nuclear currently accounts for.

[3]Fossil fuels include: natural gas, coal, petroleum and other gases.
[4]Two new nuclear generating units at Plant Vogtle in Georgia were scheduled to open in 2023 (https://www.southerncompany.com/innovation/vogtle-3-and-4.html).
[5]https://www.eia.gov/todayinenergy/detail.php?id=38792.

Moreover, Davis & Hausman (2016) document that even though nuclear power plants have low marginal costs, their prospects for long-term viability have been substantially reduced since 2009 given both the lower costs of coal and natural gas and the number of plants scheduled for closure. One report (UBS 2013) provided the following summary: "Nuclear units, with their high dispatch factors have among the greatest exposure to gas/power price volatility, as they are price takers. In tandem, nuclear generators have continued to see rising fuel and cost structures of late, with no anticipation for this to abate."

Davis & Wolfram (2012) find that deregulation and consolidation beginning in 1999 are associated with a 10% increase in plant operating performance.[6] This improvement was largely brought about by a reduction in outages. Along these lines, Hausman (2014) uses panel data from 1996-2009 to show that over the sample period, restructuring and relaxation of price regulation led to improved safety outcomes and increased output. Hence, much like the broader electricity literature (Fabrizio, Rose & Wolfram 2007, Rungsuriyawiboon & Stefanou 2007), there is consensus that the regulatory environment prior to deregulation was inferior to the current market.

Apergis, Payne, Menyah & Wolde-Rufael (2010) deploy a panel error correction model of 19 countries from 1984-2007 and find that nuclear energy consumption at the country level reduces CO_2 emissions in the long-run, while the opposite holds for renewable energy consumption. Given that renewable energy sources have historically required peaker plants to pick up the load at the tail end of the 'duck curve', while nuclear plants service baseload demand, this result is plausible from a technical viewpoint. Iwata, Okada & Samreth (2010) and Menyah & Wolde-Rufael (2010) also find that nuclear power plays a statistically significant role in reducing emissions. Davis & Hausman (2016) conclude that, "Current policies aimed at reducing carbon emissions tend to focus on wind, solar, and other renewables, but

[6] Karney (2019) replicates this result.

keeping existing nuclear plants open longer could mean hundreds of millions of tons of carbon abatement." Similarly, Roth & Jaramillo (2017) find that, "Preserving the existing nuclear power plant fleet, especially multi-reactor plants, is thus a cost effective carbon-avoidance strategy compared to the social cost of carbon." Apergis & Payne (2010) find a positive long-run relationship between nuclear energy consumption and real GDP. Another branch of research related to nuclear energy is principally concerned with nuclear accidents.

Zelenika-Zovko & Pearce (2011) argue that the indirect subsidy given to nuclear power plants, by the Price Anderson Act of 1957 and subsequently by the Energy Policy Act of 2005, is too large, given that nuclear facilities are not insured on the open market. The authors conclude that renewable energy sources should replace nuclear capacity. Conversely, Zhang, Mclellan, Tezuka & Ishihara (2012) argue that despite the Fukushima Daiichi accident of 2011, that nuclear power should not be discarded from an economic, environmental and energy independence viewpoint. Ultimately, Davis (2012) and Bistline, James & Sowder (2019) find that without carbon pricing, technological innovation, and business model innovation, the U.S. nuclear power sector will continue to decline.

Electric generation from nuclear power plants also produces bad outputs, primarily in the form of spent nuclear fuel and radioactive materials, which is regulated by the U.S. Nuclear Regulatory Commission. Given the high density of nuclear materials, nuclear waste is often stored on site.[7] Occasionally, nuclear leakage occurs, most notably at Three Mile Island in 1979, Chernobyl in 1986, and Fukushima in 2011.

Lastly, we note that a full accounting of the productive aspects of the nuclear power generation sector in the US have not been extensively studied. To our knowledge, only Krautmann & Solow (1988) and Rungsuriyawiboon (2008) have focused on technical change

[7]The U.S. Nuclear Regulatory Commission (NRC), "regulates spent fuel through a combination of regulatory requirements, licensing; safety and security oversight, including inspection, assessment of performance; and enforcement; operational experience evaluation; and regulatory support activities." (https://www.nrc.gov/waste/spent-fuel-storage.html, checked 03/2/2023).

and scale economies. However, neither of these papers has modeled if the plants operate efficiently. Thus, a full accounting using frontier methods is warranted.

3. The Stochastic Frontier Model

Our main interest focuses on the benchmark stochastic frontier model (SFM):

$$y_i = \boldsymbol{x}_i'\beta + v_i - u_i,\ i = 1, \ldots, n, \tag{1}$$

where u_i is a non-negative random variable representing technical inefficiency. We assume that v_i and u_i are independent of one another and of $\boldsymbol{x}_i$, though we will allow for some elements of $\boldsymbol{x}$ to be endogenous in our application. A straightforward approach to estimation of this model is to place assumptions on both of the error components and optimize the corresponding likelihood. However, if these distributional assumptions are incorrect this can then lead to statistical issues with the estimation of the parameters of the model. Moreover, if there are outliers in the data, this can also impact direct maximum likelihood estimation. We elect instead to use a Bayesian "quantile-like" approach to estimate the stochastic frontier model. This has the ability to impart robustness to the analysis.[8]

Interestingly, application of quantile methods have been slow to mature in the study of technical efficiency (Kumbhakar, Parmeter & Zelenyuk 2020). The extant approaches can be divided into those that simply deploy quantile regression and then classify firms as efficient based on the sign of their residual (Bernini, Freo & Gardini 2004, Knox, Blankmeyer & Stutzman 2007, Liu, Laporte & Ferguson 2008, Behr 2010), to more formal setups that attempt to parse noise from technical inefficiency while engaging in a quantile style framework (Jradi, Parmeter & Ruggiero 2019, Jradi, Parmeter & Ruggiero 2021, Tsionas 2020, Tsionas

[8]We avoid directly using the term "quantile stochastic frontier" since, as noted by Papadopoulos & Parmeter (2022), estimating quantiles in a stochastic frontier model is not clearly linked to any meaningful productive concept.

et al. 2020). It is this last group of papers that most closely align with the work presented here. Jradi et al. (2019) seek to determine the quantile that is 'consistent' with the idea of a frontier, what they call, for lack of a better term, the "optimal quantile".[9] Setting the quantile in this fashion explicitly acknowledges that we do not know at which probability level the corresponding conditional quantile agrees with the true production efficiency frontier and must be estimated. The work of both Tsionas (2020) and Tsionas et al. (2020) look at standard Bayesian methods to estimate the stochastic frontier model,[10] but do so within the confines of the Asymmetric Laplace distribution. Here our goal is to consider the models of Tsionas (2020) and Tsionas et al. (2020) but to use Bayesian exponentially tilted empirical likelihood (Schennach 2005, BETEL hereafter) methods to estimate the model.[11]

While the application of BETEL may be viewed as straightforward for the stochastic frontier model, previous uses of this approach in the quantile regression setting (for example Lancaster & Jun 2010) have done so with a *known* quantile (such as the median). In the stochastic frontier setting however, only one quantile is consistent with the exact location of the frontier (Papadopoulos & Parmeter 2022), which is almost universally *unknown* in applications. Further, as shown in Horrace et al. (2023), τ cannot be interpreted as a quantile; rather it captures a degree of asymmetry of the noise distribution. Thus, differently from earlier applications of BETEL, we embed a prior on τ to facilitate data-driven estimation of the asymmetry.[12]

Our extension of the BETEL approach to estimate the stochastic frontier model is not without additional complications however. The moment conditions we consider here involve unobserved variables (more specifically technical inefficiencies due to the composed error

[9]This terminology is not ideal but is the one in the current argot of the field. See Papadopoulos & Parmeter (2022) for a detailed treatment of this entire literature.
[10]See also Makiela & Osiewalski (2018) and Haider & Mishra (2021).
[11]See Appendix A for a full discussion of BETEL in general.
[12]See Appendix B for a detailed treatment of our MCMC algorithm.

of the stochastic frontier model). We circumvent this problem via application of entropic latent variable integration simulation (Schennach 2014, ELVIS hereafter). Through ELVIS, we average these moment conditions, using a least-favorable distribution and entropy maximization, via simulation. This results in conventional moment conditions involving only observable variables and, therefore, one can use GMM or BETEL, as we do in this paper.[13] The method requires the specification of a (dominating conditional) measure for the distribution of the unobservables given the observables. The exact choice has no direct effect on the results as long as it satisfies some basic properties stated in Definition 2.2 and Proposition 2.1 of Schennach (2014). Using a measure in the Exponential family satisfies the so-called "least favorable" property which roughly says that the dominating measure is not simpler compared to the actual distribution of the unobservables given the observed variables.

Conditional on u_i, our analysis goes through if we replace y_i by $y_i + u_i$. This is hardly any progress as the u_is are unknown. Suppose, however, that we assign to each u_i (assumed to be iid) a marginal distribution[14] whose density is $f(u_i|\boldsymbol{\theta})$ and $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^{d_\theta}$ is defined as a vector that includes β and parameters in the distribution of u_i. Then, with some abuse of notation, we have

$$p(\boldsymbol{\theta}, \boldsymbol{u}|\mathcal{D}) \propto p(\boldsymbol{\theta}) \prod_{i=1}^{n} w_i^*(\boldsymbol{\theta}, \boldsymbol{u}) \cdot \prod_{i=1}^{n} f(u_i|\boldsymbol{\theta}), \tag{2}$$

where $\boldsymbol{u} = [u_1, \ldots, u_n]'$, $p(\boldsymbol{\theta})$ is a prior on $\boldsymbol{\theta}$,

$$w_i^*(\boldsymbol{\theta}, \boldsymbol{u}) = \frac{\exp\left\{\boldsymbol{\lambda}(\boldsymbol{\theta}, \boldsymbol{u})' \varrho_\tau(y_i + u_i - \boldsymbol{x}_i'\beta)\boldsymbol{x}_i\right\}}{\sum_{i'=1}^{n} \exp\left\{\boldsymbol{\lambda}(\boldsymbol{\theta}, \boldsymbol{u})' \varrho_\tau(y_{i'} + u_{i'} - \boldsymbol{x}_{i'}'\beta)\boldsymbol{x}_{i'}\right\}}, \; i = 1, \ldots, n, \tag{3}$$

and the Lagrange multipliers, $\boldsymbol{\lambda}(\boldsymbol{\theta}, \boldsymbol{u})$, solve:

$$\boldsymbol{\lambda}(\boldsymbol{\theta}, \boldsymbol{u}) = \arg\min_{\boldsymbol{\varphi}} n^{-1} \sum_{i=1}^{n} \exp\left\{\boldsymbol{\varphi}' \left[\varrho_\tau(y_i + u_i - \boldsymbol{x}_i'\beta)\boldsymbol{x}_i\right]\right\}. \tag{4}$$

[13]Example 1.5 of Schennach (2014) describes the measurement error case.

[14]Here an advantage is that the u_is are not necessarily assumed to be independent of the regressors $\boldsymbol{x}_i$ or the error term v_i.

The expression in Equation (2) can be used to define the final posterior for the SFM:

$$
\begin{aligned}
p(\boldsymbol{\theta}|\mathcal{D}) &\propto \int_{\mathbb{R}_+^n} p(\boldsymbol{\theta}, \boldsymbol{u}|\mathcal{D})\,\mathrm{d}\boldsymbol{u} \qquad (5)\\
&\propto p(\boldsymbol{\theta}) \cdot \int_{\mathbb{R}_+^n} \prod_{i=1}^{n} w_i^*(\boldsymbol{\theta}, \boldsymbol{u}) \cdot \prod_{i=1}^{n} f(u_i|\boldsymbol{\theta})\,\mathrm{d}\boldsymbol{u}.
\end{aligned}
$$

Therefore, the problem can be understood in two stages: In the first stage, we use a specification for $y_i|\boldsymbol{x}_i, \boldsymbol{\theta}, u_i$ which is based on the derivative of the check function (which can be used to provide the first-order conditions for MLE). In the second stage, we assign a distribution to inefficiency, u_i, to yield (2). We assume that $u_i \overset{iid}{\sim} \mathcal{N}_+(0, \sigma_u^2)$, viz. a Half-Normal distribution whose density is $f(u|\sigma_u) = \left(\frac{\pi\sigma_u^2}{2}\right)^{-1/2} e^{-u^2/(2\sigma_u^2)}$, $u \geq 0$.

In practice, to explore the posterior we use MCMC, drawing from the posterior conditional distributions $\boldsymbol{\theta}|\boldsymbol{u}, Y$ and $\boldsymbol{u}|\boldsymbol{\theta}, Y$. For the first distribution we utilize first and second-order derivative information from the log posterior using Girolami & Calderhead's (2011) algorithm. For the second distribution, the procedure is more involved. For the conditional posterior of interest, $p(\boldsymbol{u}|\boldsymbol{\theta}, Y)$, we approximate it making the guess that

$$
p(\boldsymbol{u}|\boldsymbol{\theta}, Y) \simeq q(\boldsymbol{u}; \boldsymbol{\alpha}) \propto \prod_{i=1}^{n} q_i(u_i; \boldsymbol{\alpha}), \qquad (6)
$$

where $\boldsymbol{\alpha} \in \mathcal{A} \subseteq \mathbb{R}^{d_\alpha}$ is a parameter vector that we calibrate to make $q(\boldsymbol{u}; \boldsymbol{\alpha})$ as close as possible to $p(\boldsymbol{u}|\boldsymbol{\theta}, Y)$. We assume that each $q_i(u_i; \boldsymbol{\alpha})$ is the density of a Truncated Normal distribution, denoted $\mathcal{N}_+(\alpha_{i1}, \alpha_{i2}^2)$, where α_{i1} and α_{i2}^2 are respectively location and scale parameters ($i = 1, \ldots, n$). The calibration is performed by minimizing Kullback-Leibler divergence:

$$
KL = \int_{\mathcal{A}} q(\boldsymbol{u}; \boldsymbol{\alpha}) \log \frac{q(\boldsymbol{u}; \boldsymbol{\alpha})}{p(\boldsymbol{u}|\boldsymbol{\theta}, Y)}\,\mathrm{d}\boldsymbol{u}, \qquad (7)
$$

for a given $\boldsymbol{\theta} \in \Theta$. Given a random sample $\left\{\tilde{\boldsymbol{u}}^{(d)}, d = 1, \ldots, D\right\}$ from the distribution with density $q(\boldsymbol{u}; \boldsymbol{\alpha})$, the integral can be approximated by

$$KL \simeq D^{-1} \sum_{d=1}^{D} \log \frac{q(\tilde{\boldsymbol{u}}^{(d)}; \boldsymbol{\alpha})}{p(\tilde{\boldsymbol{u}}^{(d)}|\theta, Y)}. \tag{8}$$

As the Kullback-Leibler divergence provides a lower bound to the "marginal likelihood" $p(\boldsymbol{u}|\mathcal{D})$ its use here, is natural. In turn, we use standard numerical optimization techniques to solve for the optimal values, $\hat{\boldsymbol{\alpha}}$ of $\boldsymbol{\alpha}$. What makes this process computationally intensive is that we need to have access to the posterior in the denominator of Equation (8); this requires repeatedly solving for the Lagrange multipliers in Equation (4). However, this problem is much easier than Equation (8) as the dimensionality of $\boldsymbol{\varphi}$ in Equation (4) matches the dimensionality of the moment conditions.

To make the approach even simpler, we do not solve for each new Lagrange multiplier vector that results from a new evaluation of $p(\tilde{\boldsymbol{u}}^{(d)}|\theta, Y)$ in Equation (8). Instead, we fix $\boldsymbol{\varphi}$ for each $d \in \{1, \ldots, D\}$ and we assume that α_{i1}, and α_{i2}^2 are the same across i so, in effect, there are only two parameters for the optimization in Equation (8). Assuming that the α_{i1}s are all different (but the scale is the same) did not result in any material differences. Specifically, we assume $D = 1,000$ (making the results remain the same for $D = 2,000$ and $D = 5,000$ and we use S =150,000 MCMC iterations, omitting the first 50,000 to mitigate possible start up effects. The optimization in Equation (8) is performed only during the burn-in phase which results in additional computational savings. The prior we have used is:

$$p(\theta) \propto p(\beta, \sigma_u, \tau) \propto p(\beta|\sigma_u, \tau)\sigma_u^{-(\underline{n}+1)} e^{-\underline{q}/(2\sigma_u^2)} p(\tau), \; \sigma_u > 0. \tag{9}$$

Our prior $p(\beta|\sigma_u, \tau)$ is flat in the set $\mathfrak{B}$ which is defined as the set where the monotonicity conditions of the production function hold at the mean of the data and another randomly chosen nine data points. We set $\underline{n} = 0$ and $\underline{q} = 10^{-6}$ which is improper and imposes little prior information. Our prior for the quantile parameter is $\tau \sim \mathcal{B}(a, b)$, where $\mathcal{B}(a, b)$ denotes the Beta distribution, viz. $p(\tau) \propto \tau^{a-1}(1-\tau)^{b-1}, 0 < \tau < 1, a, b > 0$. We set the hyperparameters a and b both to 5. As we will see for our application in Figure 2, the prior is proper but not particularly informative; the prior mean is $\frac{1}{2}$ and places positive probability density over all values in $(0, 1)$. In our MCMC, we reparametrize $\tau^* = \text{arctanh}(\tau)$ or $\tau^* = \frac{1}{2} \log \frac{1+\tau}{1-\tau}$, and $\sigma_u^* = \log \sigma_u$ so that the resulting parameters for τ and σ_u fall in the needed domain. The arctanh transformation is Fisher's transform for the correlation coefficient.

As the MCMC provides draws $\left\{\boldsymbol{u}^{(s)}, s = 1, \ldots, S\right\}$ that converge in distribution (as $S \to \infty$) to $p(\boldsymbol{u}|\mathcal{D})$, these draws can be used in a straightforward way to compute posterior moments and/or marginal densities of technical inefficiencies. For example, posterior mean technical efficiency is estimated as $r_i = S^{-1} \sum_{s=1}^{S} e^{-u_i^{(s)}}$, $i = 1, \ldots, n$.

4. The Performance of Electricity Generation from Nuclear Power

We collect output data for a subset of nuclear power plants in the United States. The data originates from the FERC Form 1 data requirements and span the period of 1994 to 2016. These data represent plant-level observations for major[15] steam-electric generating facilities that use nuclear fuel. The variables considered in the analysis include a single output, net generation exclusive of plant use in kilowatt hours (Y) four inputs: the total installed capacity or maximum generator name plate ratings - MW (K), average number of employees (L), barrels of oil (Oil) and units of nuclear fuel in grams of Uranium (Nf) indicated by each plant, and two plant level controls: plant hours connected to the load (Hours) and net peak demand on Plant in MW per 60 minutes (Peak).

[15]FERC defines major as: "(1) one million Megawatt hours or more; (2) 100 megawatt hours of annual sales for resale; (3) 500 megawatt hours of annual power exchange delivered; or (4) 500 megawatt hours of annual wheeling for others (deliveries plus losses)."

Table 1. Summary statistics for nuclear electric generating plants from 1994-2016

	Y	K	L	Oil	Nf	Peak	Hours
mean	868,572	1,080	765	39	18,558	988	7,449
sd	611,121	693	337	333	37,729	697	2,662
min	4,176	3	22	0	1	0	0
max	2,137,776	2,667	2,277	4,139	193,950	2,605	17,507

Notes: The sample contains 481 observations.

Table 1 displays the mean, standard deviation, maximum and minimum for each variable. Y, Oil and Nf are each divided by 1,000 for fit. The sample contains 481 observations across 58 power plants.

Our sample consists of plants in nearly every state with nuclear reactors, from New York to California to Texas to Minnesota. The states not represented in the sample are almost entirely contained to areas outside of regional transmission organizations (RTO) in the Western Electricity Coordinating Council area where there are no nuclear power plants. As nuclear plants serve regions of large load demand, this makes our sample highly representative of nuclear plants in the the U.S. The distribution of our data is fairly even across the remaining states, as can be seen in Figure 1.

We model output as a translog function:

$$\begin{aligned}\ln y_{it} =& \beta_0 + \sum_{m=1}^{4} (\beta_m + \rho_m t) \ln x_{mit} + 0.5 \sum_{p=1}^{4} \sum_{m=1}^{4} \beta_{mp} \ln x_{mit} \ln x_{pit} \\ &+ (\delta_1 + \delta_2 t) Peak_{it} + (\gamma_1 + \gamma_2 t) Hours_{it} + \lambda_1 t + (1/2)\lambda_2 t^2 + v_{it} - u_{it}, \end{aligned} \tag{10}$$

where t is a time trend and our four conventional inputs (x_m) are K, L, Nf and Oil. We treat Peak and Hours as exogenous controls. We estimate two versions of this production function using quantile BETEL. Quantile BETEL I treats all inputs, the time trend and their squares and interactions as exogenous. Quantile BETEL II treats all lagged inputs as

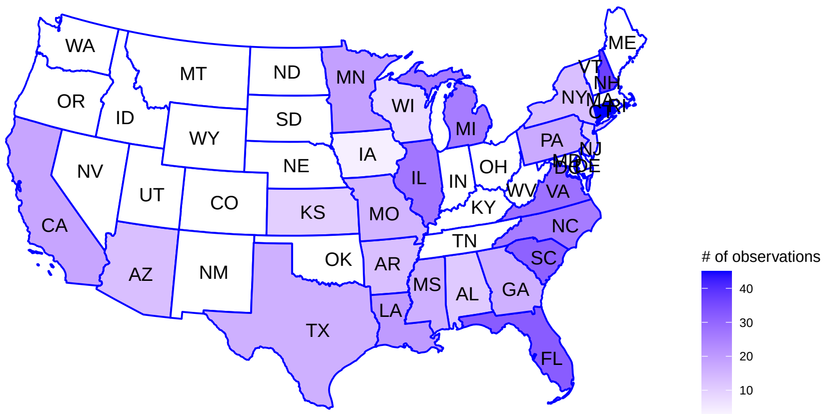

FIGURE 1. Distribution of observations in the analysis.

well as their squares and interactions as exogenous, thus allowing contemporaneous inputs to be endogenous.

In Figure 2 we provide marginal posterior densities of the parameter τ along with its prior for different values of the prior parameter $\underline{q}$. The marginal posterior densities are not materially sensitive to the choice of this parameter. The posterior mean and median are, respectively, 0.817 and 0.894, the posterior standard deviation is 0.190 and the posterior mean absolute deviation is 0.152. As the marginal posteriors seem to be bimodal with a mode around the median (0.50) we repeat the analysis using a flat prior for τ. The resulting marginal posteriors are reported in panel b.

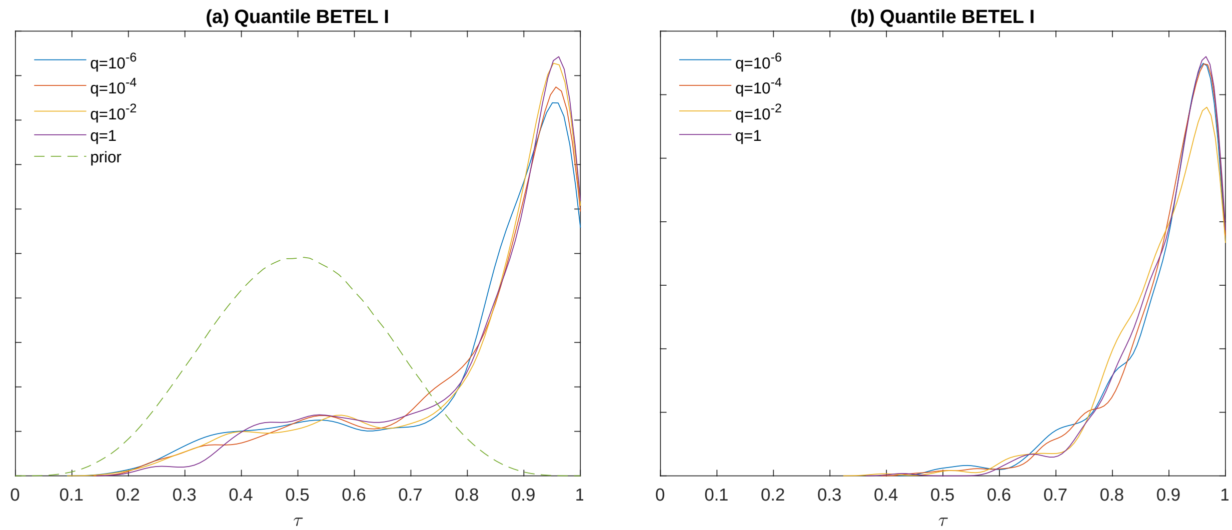

FIGURE 2. Marginal posterior densities of τ – quantile BETEL.

The new marginal posterior densities are similar to those in panel (a) around the dominant ones and lose the local mode near unity. Specifically, the posterior mean and median are, respectively, 0.897 and 0.923, the posterior standard deviation is 0.090 and the posterior mean absolute deviation is 0.070. We take this as evidence that there is some prior sensitivity to the form of the prior (beta versus flat) but the effect is, for the most part, relatively insignificant, as the conclusion remains the same in that the "optimal quantile" (which has been sought out in the previous literature, Jradi et al. 2021) is in the neighborhood of 0.90. It goes without saying that if the researcher wants to be dogmatic about the value of τ this can be easily accommodated in our approach. The results are virtually identical when we allow for potential endogeneity of inputs albeit with slightly less probability mass in the lower tail of the posterior distributions.

Our prior for the Bayes Normal-Half-Normal model is $p(\beta, \sigma_v, \sigma_u) \propto \sigma_v^{-(\bar{n}+1)} e^{-\bar{q}/(2\sigma_v^2)} \sigma_u^{-1}$, where $\bar{n} = 0$ and $\bar{q} = 10^{-6}$ which is improper and imposes little prior information. We use

150,000 Gibbs sampling iterations, omitting the first 50,000 to mitigate possible start up effects.

Figure 3 presents the estimated input elasticities for the production of electricity across the three specifications estimated. We see that the support of the distribution of estimated elasticities for all inputs are relatively small, with the modal input elasticities of nuclear fuel being the largest (between 0.2 - 0.3), as expected. The modal input elasticities for labor and capital are close in magnitude, while oil has the smallest estimated modal input elasticity. Most plants do not use oil in production; hence, it is re-affirming that all three models find oil to be the least important factor in production.

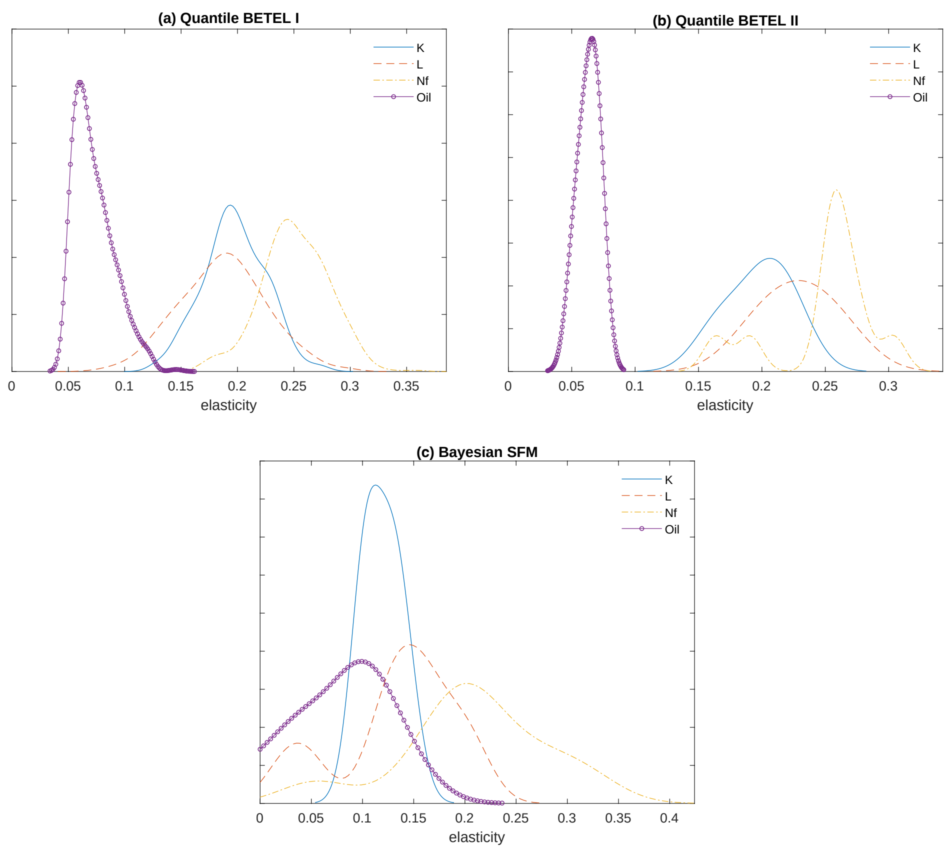

Figure 3. Estimated Input Elasticities.

It is difficult to compare these estimated input elasticities to the extant literature since so few papers have focused *exclusively* on nuclear power generation. Perhaps the closest is Rungsuriyawiboon (2008). Rungsuriyawiboon (2008) estimates price and output elasticities (found through a cost function) with the results suggesting that fuel, labor and capital had statistically different price elasticities (with capital being the most price responsive in the long run) and fuel having the largest output elasticity. This last result of Rungsuriyawiboon (2008) trues with our finding here that fuel appears to have the most sensitivity. We also observe that the model value of labor elasticity depends upon the estimation method (with the BETEL II posterior estimates having the largest model value for the input elasticity).

Figure 4 presents the posterior distribution of estimated returns to scale. The distribution is close to symmetric and centered on 0.8 for quantile BETEL I with a similar shape for quantile BETEL II (albeit a mode that is shifted to the left, near 0.75), suggesting the presence of decreasing RTS. Further, traditional Bayesian SFA produces estimated RTS much smaller than either of the BETEL models, with a modal RTS of 0.58. Our BETEL estimates (both I and II) align closely with those of Rungsuriyawiboon (2008), who found a short run RTS of 0.895 and a long run RTS of 0.776, albeit for an earlier time period than we are analyzing here.[16]

It is clear that the vast difference in modal/average RTS between our BETEL estimates (both BETEL I and II) and standard Bayesian estimation suggest that, between the allowed asymmetry of the noise term and the robust estimation that the new approach we adopt, has important implications for what is observed regarding optimal plant size. The finding of substantial decreasing RTS is important and widely known. Beginning in the early 2000s, a decline in baseload electricity generation relative to peaking and intermittent capacity occurred as a result of the Shale Revolution, and more recently, growth in wind. Intermittent

[16] Kahouli (2011), using a cross-country dataset from the years 1971 to 1997, found increasing returns to scale using a statistical approach much different from that deployed here.

capacity poses a challenge to the baseload model because in times of high production of renewable sources, baseload plants lack flexibility, whereas natural gas plants can generally adjust their output levels rapidly.

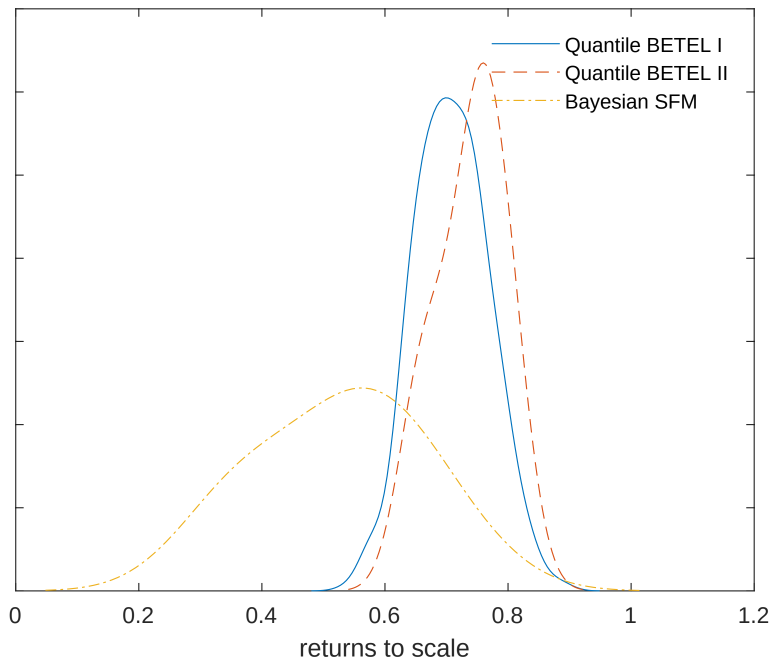

FIGURE 4. Estimated Returns to Scale.

The nuclear industry is well aware of the lack of scale economies. As noted in Johnson (2013), "Proposals for smaller, modular reactors – from less than one-tenth to one-third the size of a traditional reactor could speed construction, reducing costs and financial risk. SMRs, as they're known, would be built using modular components made in factories and shipped to the site. Promoters hope that process would bring economies of scale to a business where few have existed."

Figure 5 displays the distribution of estimated production efficiency. While RTS may be below one, we see that the majority of our plant-year observations have high levels of efficiency (for both quantile BETEL I and II), suggesting little room for improvement in overall nuclear energy production. This is consistent with the idea that increasing electricity production at a nuclear plant is a relatively straightforward process. Contrast this with the standard Bayesian SFM estimates of inefficiency, which suggest that the industry is more inefficient. We note that we cannot make comparisons with Rungsuriyawiboon (2008) as that approach did not allow for technical efficiency. However, with such high levels of technical efficiency, it is surmised that this might not have had a large impact on those empirical findings.

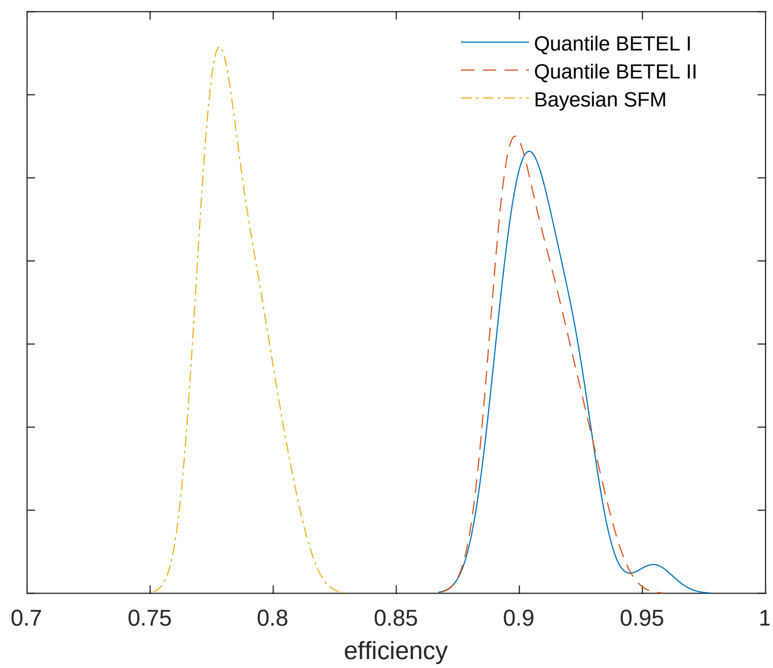

FIGURE 5. Estimated Efficiency.

Productive efficiency in this context is explicitly linked to the capacity factor of nuclear plants via the numerator. The capacity factor is defined as the actual output relative to potential output given capacity, while productive efficiency is actual output over model output. In 2016, nuclear plants had a capacity factor of close to 92% relative to 34% percent for wind and roughly 55% for coal and natural gas (EIA). The disparity in capacity factors can be understood for two reasons. First, nuclear plants require less maintenance than natural gas or coal plants. Second, while solar and wind energy sources can be highly intermittent hour to hour, nuclear fuel lasts for over a year.

We do, however, see that some nuclear plants operate well below frontier levels. Another substantial difference between the methods developed here and more traditional methods are the estimated efficiency scores. The modal efficiency for quantile BETEL I and II is right at 0.90 while standard SFA produces a model mean efficiency of 0.78, a 14% difference. This is an economically meaningful difference in the implications between the models. For comparison, using a jointly derived dataset on natural gas plants, Bernstein (2020) finds (transient) productive efficiency to be 0.71 percent on average for standard SFA with much wider distributions than the present work.

Figure 6 presents the estimated distributions of technical change,[17] efficiency change[18] and productivity growth[19] for the nuclear plants in our database. Technical change is, on average slightly positive but distributed around 0. This is not surprising as once a nuclear plant is built, the mechanism by which to generate electricity is fixed. This magnitude is also consistent with the earlier findings of Rungsuriyawiboon (2008, pg. 13) who notes that "The magnitudes of the technological change are relatively small, but significant. These results suggest that technological progress in nuclear power generation may have slowed over the sample period."

Both efficiency change and productivity growth are negative, at about 2% annually at the modal value. Again, this is not surprising as the massive fixed costs of building a nuclear plant make changes over time costly. In fact, despite the closure of several nuclear power plants in the U.S. over recent years, the EIA reports that from 2010 to 2018, nuclear power

[17]See Sengupta (1995) for a definition.
[18]Badunenko & Kumbhakar (2017) for a formal definition.
[19]Caves, Christensen & Swanson (1981) for a definition.

generation actually increased slightly. In order to increase a plant's output, additional capital expenditures are required to use more highly enriched fuel or the use of newer fuel, according to the U.S. Nuclear Regulatory Commission.

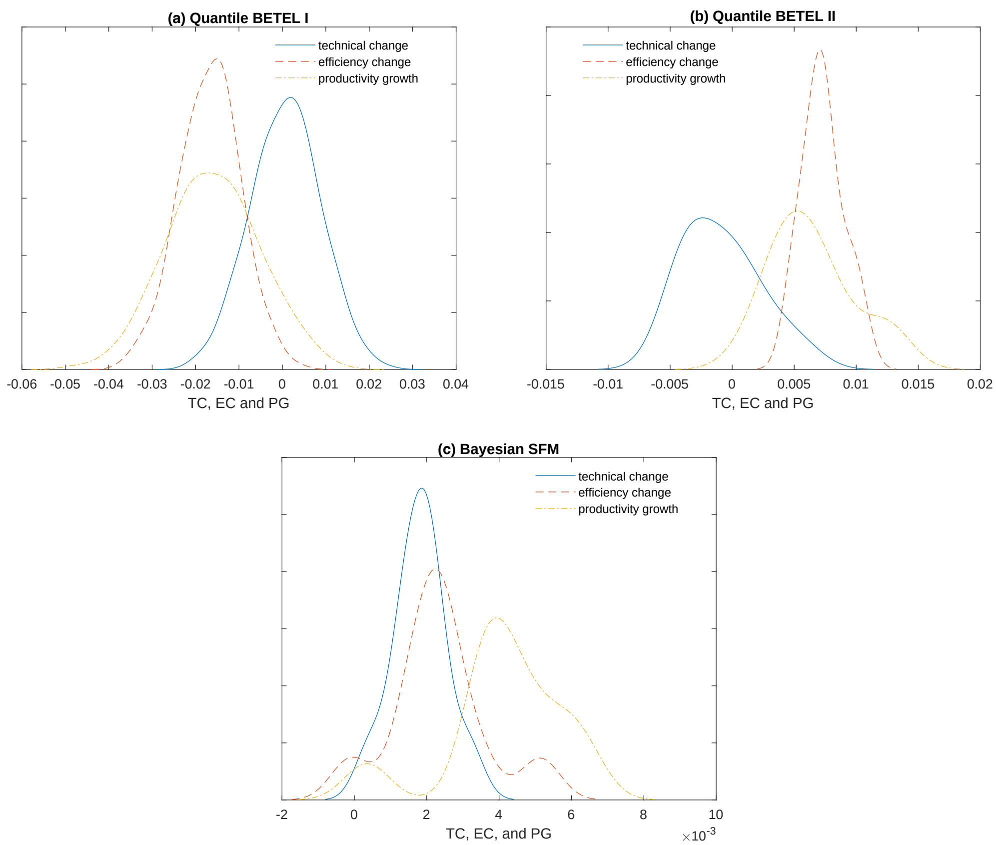

FIGURE 6. Estimated Technical change, efficiency change and productivity growth.

5. Concluding Remarks

This paper has proposed adaptation of Schennach's (2005) BETEL approach to estimate the stochastic frontier model which posed additional challenges as the inefficiencies are unobserved. We deployed this new approach to novel data studying the generation of electricity through nuclear power in the U.S. This industry has seen a steady decline in the U.S. over the past several decades and some have questioned its long term viability. Our results here suggest that the industry is characterized by decreasing returns to scale, negative technical change and high levels of efficiency.

Krautmann & Solow (1988) explain that predicting the future of the nuclear sector is difficult due to the ever evolving regulatory environment. They note that, "...it is not clear what impact the regulations imposed after Three Mile Island have had on economies of scale...". More recently, the future of nuclear power is similarly difficult to assess. On the one hand, the closures of plants, and growing public opposition to nuclear power (according to a recent Gallup poll), could spell a downturn. Conversely, the ability of modern small-scale modular reactors to increase the current RTS measured herein, and the potential of nuclear energy in quickly fighting climate change, are developments that could revitalize the sector over the medium-run. Future work should continue to study this critically important area of electricity generation as the U.S. moves further towards a fully clean energy grid.

References

Apergis, N. & Payne, J. E. (2010), 'A panel study of nuclear energy consumption and economic growth', *Energy Economics* **32**(3), 545–549.

Apergis, N., Payne, J. E., Menyah, K. & Wolde-Rufael, Y. (2010), 'On the causal dynamics between emissions, nuclear energy, renewable energy, and economic growth', *Ecological Economics* **69**(11), 2255–2260.

Badunenko, O. & Henderson, D. J. (2023), 'Production analysis with asymmetric error', *Journal of Productivity Analysis* . Forthcoming.

Badunenko, O. & Kumbhakar, S. C. (2017), 'Economies of scale, technical change and persistent and time-varying cost efficiency in Indian banking: Do ownership, regulation and heterogeneity matter?', *European Journal of Operational Research* **260**, 789–803.

Behr, A. (2010), 'Quantile regression for robust bank efficiency score estimation', *European Journal of Operational Research* **200**, 568–581.

Bera, A. K., Galvao, A. F., Montes-Rojas, G. V. & Park, S. Y. (2016), 'Asymmetric Laplace regression: Maximum likelihood, maximum entropy and quantile regression', *Journal of Econometric Methods* **5**, 79–101.

Bernini, C., Freo, M. & Gardini, A. (2004), 'Quantile estimation of frontier production function', *Empirical Economics* **29**, 373–381.

Bernstein, D. H. (2020), 'An updated assessment of technical efficiency and returns to scale for U.S. electric power plants', *Energy Policy* **147**, 111896.

Bistline, J., James, R. & Sowder, A. (2019), 'Technology, policy, and market drivers of (and barriers to) advanced nuclear reactor deployment in the United States after 2030', *Nuclear Technology* **205**(8), 1075–1094.

Bonanno, G., De Giovanni, D. & Domma, F. (2015), 'The 'wrong skewness' problem: a re-specification of stochastic frontiers', *Journal of Productivity Analysis* **47**(1), 49–64.

Caves, D., Christensen, L. & Swanson, J. A. (1981), 'Productivity growth, scale economies, and capacity utilization in U.S. railroads, 1955-74', *American Economic Review* **71**(4), 994–1002.

Chen, Y.-Y., Schmidt, P. & Wang, H.-J. (2014), 'Consistent estimation of the fixed effects stochastic frontier model', *Journal of Econometrics* **181**(2), 65–76.

Chen, Z., Barros, C. P. & Borges, M. R. (2015), 'A Bayesian stochastic frontier analysis of Chinese fossil-fuel electricity generation companies', *Energy Economics* **48**, 136–144.

Davis, L. & Hausman, C. (2016), 'Market impacts of a nuclear power plant closure', *American Economic Journal: Applied Economics* **8**(2), 92–122.

Davis, L. W. (2012), 'Prospects for nuclear power', *Journal of Economic Perspectives* **26**(1), 49–66.

Davis, L. W. & Wolfram, C. (2012), 'Deregulation, consolidation, and efficiency: Evidence from US nuclear power', *American Economic Journal: Applied Economics* **4**(4), 194–225.

Fabrizio, K. R., Rose, N. L. & Wolfram, C. D. (2007), 'Do markets reduce costs? assessing the impact of regulatory restructuring on US electric generation efficiency', *American Economic Review* **97**(4), 1250–1277.

Filippini, M., Geissmann, T. & Greene, W. H. (2018), 'Persistent and transient cost efficiency—an application to the Swiss hydropower sector', *Journal of Productivity Analysis* **49**(1), 65–77.

Ghosh, R. & Kathuria, V. (2016), 'The effect of regulatory governance on efficiency of thermal power generation in India: A stochastic frontier analysis', *Energy Policy* **89**, 11 – 24.

Girolami, M. & Calderhead, B. (2011), 'Riemann manifold Langevin and Hamiltonian Monte Carlo methods', *Journal of the Royal Statistical Society, Series B* **73**(1), 123–214.

Gourieroux, C., Monfort, A., Renault, E. & Trognon, A. (1987), 'Generalised residuals', *Journal of Econometrics* **34**(1), 5–32.

Greene, W. H. (1990), 'A gamma-distributed stochastic frontier model', *Journal of Econometrics* **46**(1-2), 141–164.

Greene, W. H. (2008), 'The econometric approach to efficiency analysis', *The measurement of productive efficiency and productivity growth* **1**(1), 92–250.

Haider, S. & Mishra, P. P. (2021), 'Does innovative capability enhance the energy efficiency of Indian Iron and Steel firms? a Bayesian stochastic frontier analysis', *Energy Economics* **95**, 105128.

Hausman, C. (2014), 'Corporate incentives and nuclear safety', *American Economic Journal: Economic Policy* **6**(3), 178–206.

Horrace, W. C., Parmeter, C. F. & Wright, I. A. (2023), 'On asymmetry and quantile estimation of the stochastic frontier model', *Journal of Productivity Analysis* . Forthcoming.

Iwata, H., Okada, K. & Samreth, S. (2010), 'Empirical study on the environmental Kuznets curve for CO2 in France: The role of nuclear energy', *Energy Policy* **38**(8), 4057–4063.

Johnson, K. (2013), 'What's holding back nuclear energy', *Wall Stret Journal* .
URL: *https://www.wsj.com/articles/what8217s-holding-back-nuclear-energy-1384202091*

Jradi, S., Parmeter, C. F. & Ruggiero, J. (2019), 'Quantile estimation of the stochastic frontier model', *Economics Letters* **182**, 15–18.

Jradi, S., Parmeter, C. F. & Ruggiero, J. (2021), 'Quantile estimation of stochastic frontiers with the Normal-Exponential specification', *European Journal of Operational Research* **295**, 475–483.

Kahouli, S. (2011), 'Effects of technological learning and uranium price on nuclear cost: Preliminary insights from a multiple factors learning curve and uranium market modeling', *Energy Economics* **33**(5), 840–852.

Karney, D. H. (2019), 'Electricity market deregulation and environmental regulation: Evidence from U.S. nuclear power', *Energy Economics* p. 104500.

Knittel, C. R. (2002), 'Alternative Regulatory Methods and Firm Efficiency: Stochastic Frontier Evidence from the U.S. Electricity Industry', *The Review of Economics and Statistics* **84**(3), 530–540.

Knox, K. J., Blankmeyer, E. C. & Stutzman, J. R. (2007), 'Technical efficiency in Texan nursing facilities: a stochastic production forntier approach', *Journal of Economics and Finance* **31**(1), 75–86.

Koenker, R. (2005), *Quantile Regression*, Cambridge University Press.

Koenker, R. & Bassett, G. (1978), 'Regression quantiles', *Econometrica* **46**(1), 33–50.

Koenker, R. & Hallock, K. (2001), 'Quantile regression', *Journal of Economic Perspectives* **15**, 143–156.

Komunjer, I. (2005), 'Quasi-maximum likelihood estimation for conditional quantiles', *Journal of Econometrics* **128**(1), 137–164.

Krautmann, A. C. & Solow, J. L. (1988), 'Economies of scale in nuclear power generation', *Southern Economic Jounal* **55**(1), 70–85.

Kumbhakar, S. C., Parmeter, C. F. & Zelenyuk, V. (2020), Stochastic frontier analysis: Foundations and advances II, *in* S. Ray, R. Chambers & S. C. Kumbhakar, eds, 'Handbook of Production Economics', Springer, Singapore.

Lancaster, T. & Jun, S. J. (2010), 'Bayesian quantile regression methods', *Journal of Applied Econometrics* **25**(2), 287–307.

Li, H.-Z., Kopsakangas-Savolainen, M., Xiao, X.-Z., Tian, Z.-Z., Yang, X.-Y. & Wang, J.-L. (2016), 'Cost efficiency of electric grid utilities in China: A comparison of estimates from SFA-MLE, SFA-Bayes and StoNED-CNLS', *Energy Economics* **55**, 272–283.

Liu, C., Laporte, A. & Ferguson, B. S. (2008), 'The quantile regression approach to efficiency measurement: Insights from Monte Carlo simulations', *Health Economics* **17**, 1073–1087.

Makiela, K. & Osiewalski, J. (2018), 'Cost efficiency analysis of electricity distribution', *The Energy Journal* **39**, 31–56.

Menyah, K. & Wolde-Rufael, Y. (2010), 'CO2 emissions, nuclear energy, renewable energy and economic growth in the US', *Energy Policy* **38**(6), 2911–2915.

Papadopoulos, A. & Parmeter, C. F. (2022), 'Quantile Methods for Stochastic Frontier Analysis', *Foundations and Trends in Econometrics* **12**(1), 1–132.

Parmeter, C. F. & Kumbhakar, S. C. (2014), 'Efficiency Analysis: A Primer on Recent Advances', *Foundations and Trends in Econometrics* **7**(3-4), 191–385.

Roth, M. B. & Jaramillo, P. (2017), 'Going nuclear for climate mitigation: An analysis of the cost effectiveness of preserving existing U.S. nuclear power plants as a carbon avoidance strategy', *Energy* **131**, 67–77.

Rungsuriyawiboon, S. (2008), 'Estimating cost structures in the U.S. nuclear power industry', *Energy Exploration & Exploitation* **26**(1), 1–22.

Rungsuriyawiboon, S. & Stefanou, S. E. (2007), 'Dynamic Efficiency Estimation: An Application to U.S. Electric Utilities', *Journal of Business & Economic Statistics* **25**(2), 226–238.

Schennach, S. M. (2005), 'Bayesian exponentially tilted empirical likeliood', *Biometrika* pp. 31–46.

Schennach, S. M. (2014), 'Entropic latent variable integration via simulation', *Econometrica* **82**(1), 345–385.

Sengupta, J. K. (1995), *Technical Change and Efficiency*, Springer Netherlands, Dordrecht, pp. 86–132.

Simar, L. & Wilson, P. W. (2010), 'Inferences from cross-sectional, stochastic frontier models', *Econometric Reviews* **29**(1), 62–98.

Simar, L. & Wilson, P. W. (2013), 'Estimation and inference in nonparametric frontier models: Recent developments and perspectives', *Foundations and Trends in Econometrics* **5**(2), 183–337.

Tsionas, E. G., Assaf, A. G. & Andrikopoulos, A. (2020), 'Quantile stochastic frontier models with endogeneity', *Economics Letters* **188**, 108964.

Tsionas, M. G. (2020), 'Quantile stochastic frontiers', *European Journal of Operational Research* **282**(3), 1177–1184.

Wei, Z., Zhu, X. & Wang, T. (2021), 'The extended skew-normal-based stochastic frontier model with a solution to 'wrong skewness' problem', *Statistics* pp. 1–20.

Yang, Y. & He, X. (2012), 'Bayesian empirical likelihood for quantile regression', *Annals of Statistics* **40**(2), 1102–1131.

Zelenika-Zovko, I. & Pearce, J. (2011), 'Diverting indirect subsidies from the nuclear industry to the photovoltaic industry: Energy and financial returns', *Energy Policy* **39**(5), 2626–2632.

Zhang, Q., Mclellan, B. C., Tezuka, T. & Ishihara, K. N. (2012), 'Economic and environmental analysis of power generation expansion in Japan considering Fukushima nuclear accident using a multi-objective optimization model', *Energy* **44**(1), 986–995.

Appendix A. Bayesian Exponentially Tilted Empirical Likelihood

The appeal of studying conditional quantiles is manifest for practitioners. Faced with clear underlying heterogeneity, the conditional mean can obscure differences that arise in the conditional distribution. Moreover, quantile estimators are robust in the presence of outliers that are known to influence and undermine their counterparts constructing an average. Given the recognized importance of studying quantiles of conditional distributions, a wide range of estimators have been proposed for an even wider array of statistic and econometric models.

Among available estimators of conditional quantiles, quantile regression is undoubtedly the most popular and easily applied (Koenker & Hallock 2001, Koenker 2005). However, a wide variety of alternative approaches have recently begun to appear in the literature. For example, Komunjer (2005) proposes quasi-maximum likelihood estimation based on the 'Tick-Exponential' family, which is demonstrated to be equivalent to quantile regression when the Tick-Exponential family equals the Asymmetric Laplace family. This natural likelihood generalization is an important development in the progression of likelihood estimation in the quantile setting. The work of Bera et al. (2016) has demonstrated that the first order conditions from maximum likelihood using the Asymmetric Laplace distribution are equivalent to the solution stemming from the maximum entropy problem with moment constraints imposed.

Komunjer (2005) proposed a class of estimators for conditional quantiles in possibly misspecified nonlinear models belonging to the family of quasi-maximum likelihood estimators. In this paper, we propose a general class of quantile ML estimators, precisely as Koenker & Bassett (1978) defined a general class of quantile estimators for the linear model. To begin consider the linear model:

$$y_i = \boldsymbol{x}_i'\beta + v_i,\ i = 1, \dots, n, \tag{A.1}$$

where $\boldsymbol{x}_i \in \mathbb{R}^k$ is the vector of explanatory variables in a linear model whose dependent variable is $y_i \in \mathbb{R}$, $\beta \in \mathbb{R}^k$ is a parameter vector, and v_i is two-sided measurement error. We denote the data by $\mathcal{D} = \{(y_i, \boldsymbol{x}_i)\}_{i=1}^n$. Suppose $v_i \perp \boldsymbol{x}_i,\ i = 1, \ldots, n$.

Suppose we have a log likelihood function:

$$\log L(\theta; \mathcal{D}) = \sum_{i=1}^{n} l(\mathcal{D}_i; \theta), \tag{A.2}$$

where $\theta \in \Theta \subseteq \mathbb{R}^d$ is the parameter vector. The ML estimator $\hat{\theta}$ maximizes the above expression in θ. Our intention is to define a quantile maximum likelihood estimator (qMLE), $\hat{\theta}_\tau$, for every quantile $\tau \in (0, 1)$.

To do so, we find it more convenient to start from the first-order conditions:

$$\frac{\partial \log L(\theta; \mathcal{D})}{\partial \theta} = \sum_{i=1}^{n} \frac{\partial l(\mathcal{D}_i; \hat{\theta}_\tau)}{\partial \theta} = 0. \tag{A.3}$$

These first-order conditions are essentially moment conditions of the form:

$$\sum_{i=1}^{n} g(\mathcal{D}_i, \hat{\theta}_\tau) = \mathbf{0}, \tag{A.4}$$

where $g : \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$, which can often be rewritten as:

$$\sum_{i=1}^{n} h(\hat{\varepsilon}_i, \hat{\theta}_\tau) \boldsymbol{x}_i = 0, \tag{A.5}$$

with $\hat{\varepsilon}_i = y_i - \boldsymbol{x}_i'\hat{\beta}_\tau$ and $h(y_i - \boldsymbol{x}_i'\beta)$ is the generalized residual (Gourieroux, Monfort, Renault & Trognon 1987). An empirical quantile likelihood can then be defined as:

$$\begin{gathered} \max \prod_{i=1}^n p_i, \\ \text{subject to} \\ \sum_{i=1}^n p_i h\left(\hat{\varepsilon}_i, \hat{\theta}_\tau\right) \boldsymbol{x}_i = \mathbf{0}, \\ \sum_{i=1}^n p_i = 1,\ p_i \geq 0,\ i = 1, \ldots, n. \end{gathered} \tag{A.6}$$

where the weights differ from traditional likelihood analysis $p_i = 1/n\ \forall i$.[20]

Schennach (2005) has shown that a Bayesian version of empirical likelihood is BETEL which relies on the following posterior:

$$p(\theta|\mathcal{D}) \propto p(\theta) \prod_{i=1}^n p_i^*(\theta), \tag{A.7}$$

where $p(\theta)$ is a prior and $p_i^*, i = 1, \ldots, n$ are solutions to the following optimization problem:

$$\begin{gathered} \min \sum_{i=1}^n p_i \log p_i, \\ \text{subject to} \\ \sum_{i=1}^n p_i h\left(\hat{\varepsilon}_i, \hat{\theta}_\tau\right) \boldsymbol{x}_i = 0, \\ \sum_{i=1}^n p_i = 1,\ p_i \geq 0, i = 1, \ldots, n. \end{gathered} \tag{A.8}$$

The first-order conditions provide the weights:

$$p_i^*(\theta) = \frac{\exp\left\{\boldsymbol{\lambda}(\theta)' h\left(\hat{\varepsilon}_i, \hat{\theta}_\tau\right) \boldsymbol{x}_i\right\}}{\sum_{j=1}^n \exp\left\{\boldsymbol{\lambda}(\theta)' h\left(\hat{\varepsilon}_i, \hat{\theta}_\tau\right) \boldsymbol{x}_j\right\}}, i = 1, \ldots, n, \tag{A.9}$$

[20]Endogeneity of $\boldsymbol{x}$ can be easily taken into account by replacing the orthogonality between $h\left(\hat{\varepsilon}_i, \hat{\theta}_\tau\right)$ and $\boldsymbol{x}$ with orthogonality between $h\left(\hat{\varepsilon}_i, \hat{\theta}_\tau\right)$ and instruments $\boldsymbol{z}$.

where the Lagrange multipliers, $\boldsymbol{\lambda}(\theta)$, solve:

$$\boldsymbol{\lambda}(\theta) = \arg\min_{\boldsymbol{\varphi}} n^{-1} \sum_{i=1}^{n} \exp\left\{\boldsymbol{\varphi}' \left[h\left(\hat{\varepsilon}_i, \hat{\theta}_\tau\right) \boldsymbol{x}_i\right]\right\}. \tag{A.10}$$

In turn, the posterior in (A.7) can be analyzed using standard Markov Chain Monte Carlo (MCMC) methods, such as the Metropolis-Hastings algorithm. In practice, to explore the posterior we use first- and second-order derivative information from the log posterior using the algorithm of Girolami & Calderhead (2011) described in Appendix B.

Appendix B. Markov Chain Monte Carlo

A standard MCMC resampling algorithm uses local information about both the gradient and the Hessian of the log-posterior conditional of $\boldsymbol{\theta}$ at the existing draw. Following this, a Metropolis test is used to determine if the generated candidate should be accepted. High acceptance rates can be achieved by using smaller transitions in the parameter space. Unfortunately, the smaller the transition, the longer it takes to traverse the parameter space. We recommend against this standard MCMC algorithm as the Girolami & Calderhead (2011) algorithm moves appreciably faster. Even if smaller transitions are used, in high dimensions, the typical random walk underlying traditional MCMC becomes inefficient, exhibiting both poor mixing of the chain and highly correlated samples.

To overcome these known issues with MCMC Girolami & Calderhead's (2011) algorithm adapts the Metropolis adjusted Langevin algorithm (MALA) and Hamiltonian Monte Carlo (HMC) into an overarching geometric framework. It has been found that the Girolami & Calderhead (2011) algorithm performs vastly superior relative to the standard Metropolis-Hastings algorithm with autocorrelations that are much smaller across the chain. The gains in efficiency are due to the fact that the geometric framework uses explicit structure of the target density to traverse the parameter space.

To more adequately explain the algorithm of Girolami & Calderhead (2011) let $\mathcal{L}(\boldsymbol{\theta}) = \log p(\boldsymbol{\theta}|\boldsymbol{X})$ denote the log posterior of $\boldsymbol{\theta}$ for simplicity. Moreover, define

$$\boldsymbol{G}(\boldsymbol{\theta}) = \text{est.cov}\frac{\partial}{\partial\boldsymbol{\theta}}\log p(\boldsymbol{X}|\boldsymbol{\theta}) \tag{B.1}$$

as the empirical counterpart of

$$\boldsymbol{G}_o(\boldsymbol{\theta}) = -E_{\mathcal{Y}|\boldsymbol{\theta}}\frac{\partial^2}{\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}'}\log p(\boldsymbol{X}|\boldsymbol{\theta}). \tag{B.2}$$

The Langevin diffusion, which is a stochastic differential equation and is commonly used to explain Brownian motion, is given by:

$$d\boldsymbol{\theta}(t) = \frac{1}{2}\tilde{\nabla}_{\boldsymbol{\theta}}\mathcal{L}\{\boldsymbol{\theta}(t)\}dt + d\boldsymbol{B}(t) \tag{B.3}$$

where

$$\tilde{\nabla}_{\boldsymbol{\theta}}\mathcal{L}\{\boldsymbol{\theta}(t)\} = -\boldsymbol{G}^{-1}\{\boldsymbol{\theta}(t)\}\cdot\bigtriangledown_{\boldsymbol{\theta}}\mathcal{L}\{\boldsymbol{\theta}(t)\} \tag{B.4}$$

is the so called "natural gradient" of the Riemann manifold generated by the log posterior. The elements of the Brownian motion are

$$\begin{aligned}\boldsymbol{G}^{-1}\{\boldsymbol{\theta}(t)\}d\boldsymbol{B}_i(t) =& |\boldsymbol{G}\{\boldsymbol{\theta}(t)\}|^{-1/2}\sum_{j=1}^{K_\beta}\frac{\partial}{\partial\boldsymbol{\theta}}\left[\boldsymbol{G}^{-1}\{\boldsymbol{\theta}(t)\}_{ij}|\boldsymbol{G}\{\boldsymbol{\theta}(t)\}|^{1/2}\right]dt \\ &+\left[\sqrt{\boldsymbol{G}\{\boldsymbol{\theta}(t)\}}d\boldsymbol{B}(t)\right]_i.\end{aligned} \tag{B.5}$$

The lack of differentiability, due to the use of generalized residuals stemming from the check function, is not concerning given that it occurs only at a single, known, point.

To implement the resampling scheme we use the discrete form of the stochastic differential equation as follows:

$$\begin{aligned}\tilde{\boldsymbol{\theta}}_i =& \boldsymbol{\theta}_i^o + \frac{\epsilon^2}{2}\left\{\boldsymbol{G}^{-1}\left(\boldsymbol{\theta}^o\right)\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}^o\right)\right\}_i - \epsilon^2\sum_{j=1}^{K_\theta}\left\{\boldsymbol{G}^{-1}\left(\boldsymbol{\theta}^o\right)\frac{\partial\boldsymbol{G}\left(\boldsymbol{\theta}^o\right)}{\partial\boldsymbol{\theta}_j}\boldsymbol{G}^{-1}\left(\boldsymbol{\theta}^o\right)\right\}_{ij} \\ &+ \frac{\epsilon^2}{2}\sum_{j=1}^{K_\theta}\left\{\boldsymbol{G}^{-1}\left(\boldsymbol{\theta}^o\right)\right\}_{ij}\operatorname{tr}\left\{\boldsymbol{G}^{-1}\left(\boldsymbol{\theta}^o\right)\frac{\partial\boldsymbol{G}\left(\boldsymbol{\theta}^o\right)}{\partial\boldsymbol{\theta}_j}\right\} + \left\{\epsilon\sqrt{\boldsymbol{G}^{-1}\left(\boldsymbol{\theta}^o\right)}\boldsymbol{\xi}^o\right\}_i \\ =& \boldsymbol{\mu}\left(\boldsymbol{\theta}^o,\epsilon\right)_i + \left\{\epsilon\sqrt{\boldsymbol{G}^{-1}\left(\boldsymbol{\theta}^o\right)}\boldsymbol{\xi}^o\right\}_i\end{aligned} \tag{B.6}$$

where $\boldsymbol{\beta}^o$ is the current draw. The proposal density is

$$q\left(\tilde{\boldsymbol{\theta}}|\boldsymbol{\theta}^o\right) = N_{K_\theta}\left(\tilde{\boldsymbol{\theta}}, \epsilon^2\boldsymbol{G}^{-1}\left(\boldsymbol{\theta}^o\right)\right) \tag{B.7}$$

and convergence to the invariant distribution is ensured by using the standard form of the corresponding Metropolis-Hastings probability

$$\min\left\{1, \frac{p\left(\tilde{\boldsymbol{\theta}}|\cdot,\mathcal{Y}\right)q\left(\boldsymbol{\theta}^o|\tilde{\boldsymbol{\theta}}\right)}{p\left(\boldsymbol{\theta}^o|\cdot,\mathcal{Y}\right)q\left(\tilde{\boldsymbol{\theta}}|\boldsymbol{\theta}^o\right)}\right\}. \tag{B.8}$$

We select the hyperparameter ϵ during the burn-in phase so that 20-30% of candidates are, eventually, accepted.