Non-Gaussian Bridge Sampling with an Application*

Jin-Chuan Duan† and Changhao Zhang‡

(First Draft: September 18, 2015; This Draft: October 21, 2015)

Abstract

This paper provides a new bridge sampler that can efficiently generate sample paths, subject to some endpoint condition, for non-Gaussian dynamic models. This bridge sampler uses a companion pseudo-Gaussian bridge as the proposal and sequentially re-simulates sample paths via a sequence of tempered importance weights in a way bearing resemblance to the density-tempered sequential Monte Carlo method used in the Bayesian statistics literature. This bridge sampler is further accelerated by employing a novel idea of k-fold duplicating a base set of sample paths followed by support boosting. We implement this bridge sampler on a GARCH model estimated to the S&P 500 index series, and our implementation covers both parametric and non-parametric conditional distributions. Our performance study reveals that this new bridge sampler is far superior to either the simple-rejection method when it is applicable or other alternative samplers designed for paths with a fixed endpoint. Computing SRISK of the NYU-Stern Volatility Institute is then used to demonstrate the method's real-life applicability.

Keywords: sequential Monte Carlo, density tempering, Metropolis-Hastings, GARCH, systemic risk, infill estimation.

*This paper has benefited from the comments made by the participants of the seminar at UC-Berkeley and the Oberwolfach Workshop.

†Duan is with the National University of Singapore (Department of Finance, Risk Management Institute, and Department of Economics). E-mail: bizdjc@nus.edu.sg.

‡Zhang is a doctoral student in the Department of Finance, National University of Singapore. E-mail: changhao@nus.edu.sg.

1 Introduction

We propose a new bridge sampling method that can efficiently generating sample paths, subject to some endpoint condition, for any non-Gaussian dynamic model. A stochastic bridge is a random sample path, X_t, that starts from some known value initially, say, x_0, but always lands at a fixed value at some future time point, say, time T. Were the distribution of X_t (for any $0 < t < T$) conditional on x_0, X_T and other initial information variables at time 0 analytically known, sampling such a stochastic bridge would be rather straightforward. Other than Gaussian systems, general dynamic models are not expected to be analytically tractable. As a concrete example, the popular discrete-time GARCH(1,1) model's initial condition can be completely characterized by x_0 and σ_1, but the distribution of X_t conditional on x_0, σ_1 and X_T has no known tractable analytical solution. In this paper, we offer a novel sequential Monte Carlo (SMC) sampling scheme to effectively generate GARCH bridges or other general non-Gaussian bridges (discrete or continuous-time). In addition, we expand this new bridge sampling technique to allows for the endpoint condition to be a set of points as opposed to be just a singleton.

Bridge sampling is often needed for dynamic scenario analyses, particularly stress testing where one must generate many sample paths for some stochastic process of interest in order to assess the consequences of a prescribed stress scenario. Supervisor-led stress testing is becoming the norm for banks, with scenarios being prescribed by regulators and aggregated by banks in bottom-up fashion. Banks also conduct internal stress testing to determine economic capital and augment various value-at-risk measures. Supervisors themselves conduct top-down stress tests alongside banks.[1]

The well-known SRISK, a systemic risk measure championed by Nobel laureate, Professor Robert Engle, is a concrete example of dynamic scenario analysis. SRISK proposed by Acharya, *et al* (2012), uses the methodology in Brownlees and Engle (2012), and is released and maintained by the Volatility Institute at Stern School of Business, New York University. SRISK measures a financial institution's need for additional equity capital when the general stock market falls by a factor of at least 40% over a period of six months (126 trading days). SRISK assumes that a broad-based stock market index value follows the GJR-GARCH process of Glosten, *et al* (1993) on a daily frequency. Because the conditional distribution under the stress scenario has no tractable analytical solution, the SRIK methodology relies on a simple-rejection technique; that is, to simulate the GJR-GARCH model forward for six months and only keep those paths satisfying the endpoint condition. Since the endpoint condition is a set with a positive probability, simple-rejection sampling works theoretically but is rather inefficient computationally. At the average volatility level of the S&P500 index, the acceptance rate is about 0.244%, meaning that only 24.4 out of every

[1] For an overview of supervisory stress tests and regulatory perspectives, readers are referred to the BCBS Report on principles for sound stress testing practices and supervision (2009), Borio, *et al* (2014), Moretti, *et al* (2008), Cihak (2007) and Sorge (2004). Jurisdiction-specific reports are available in the EBA Report on EU-wide stress test (2014), Fed Report on SCAP (2009), Hirtle and Lehnert (2014), Borio and Drehmann (2009), Boss, *et al* (2007), Drehmann (2005) and Virolainen (2004).

10,000 generated paths will be kept.[2] As a comparison, our bridge sampler is able to reduce the computing time by a factor of over 10 times while having correct distributional properties.

At the heart of our bridge sampling scheme is the application of density-tempered SMC to sequentially modify a set of sample paths originated from a pseudo-Gaussian proposal to eventually arrive at a set of target bridge paths. The simulated sample paths are brought through a sequence of density-tempered intermediate targets with reweighting, resampling and support boosting (via the Metropolis-Hastings move). The density-tempered SMC algorithm can be likened to running many independent Markov chains simultaneously while feeding the chains with stationarily distributed random paths represented by the SMC sample for a particular intermediate target. The Metropolis-Hastings move is therefore used only for the purpose of support boosting as opposed to achieving convergence to the stationary distribution. The density-tempered SMC approach allows us to exercise control over the effective sample size at different stages and to balance between efficiency and richness of path variety. We further accelerate this bridge sampler by employing a novel idea of k-fold duplicating a base set of sample paths followed by support boosting. Our bridge-sampling algorithm is easy to implement, and delivers exact draws from the target bridge. This bridge-sampling approach is inspired by the density-tempered SMC method of Del Moral, *et al* (2006) and Duan and Fulop (2015), which are designed for Bayesian statistical analysis.

We use the ARCH/GARCH model, pioneered by Engle (1982) and Bollerslev (1986), to concretely demonstrate how our proposed bridge-sampling technique works for non-Gaussian stochastic processes. Specifically, we adopt the GJR-GARCH(1,1) version by Glosten, *et al* (1993) because it is this GARCH model underlying SRISK. Our bridge-sampling method is obviously applicable to any properly discretized version of continuous-time stochastic volatility models because they resemble the GARCH model. Intuitively, our bridge sampler is expected to work well for general non-Gaussian stochastic processes, and thus adds to the literature by providing an efficient and accurate means to simulate stochastic processes subject to an endpoint condition.

In the context of stochastic differential equations, simulating diffusion bridges has long been the interest of many researchers. By the theory of h-transforms available in, say, Rogers and Williams (2000), the diffusion bridge will share the local characteristics with the original diffusion, meaning that the diffusion term remains unchanged. The drift term of the diffusion bridge needs to be adjusted in a way to ensure that the diffusion bridge meets its fixed endpoint, but the adjustment term tends to be intractable. Nevertheless as shown in Papaspiliopoulos and Roberts (2012), an importance sampling scheme by using the linear diffusion bridge, whose dynamic is analytically solvable using the h-transform theory, can serve as a proposal. This result is fairly intuitive, because the likelihood ratio of the target diffusion bridge over the proposal of a linear diffusion bridge can be evaluated up to a proportional constant, which need not be known for the importance sampling purpose. A more sophisticated diffusion bridge sampling scheme has recently been proposed by Bladt and Sorensen (2014) where the proposal bridge is constructed by taking the path that forward and reverse sampling of the target diffusion intersect within the time span

[2]This acceptance rate is obtained with the GJR-GARCH model with conditional t distribution estimated to the daily time series of the S&P500 index values from 2 January 1990 to 18 September 2015.

of the bridge. In contrast to the literature on diffusion bridges, our bridge sampler is also based on importance sampling but can handle both continuous-time and discrete-time stochastic processes even with non-parametric conditional distributions. Moreover, the density-tempered SMC feature really distinguishes our bridge sampler. It is quite clear that the importance sampling schemes in the diffusion bridge literature can benefit from adopting the density-tempered SMC in their algorithmic designs.

Our bridge sampler can be fairly easily generalized to handle multidimensional bridges. More interestingly perhaps is the situation where the paths of a multivariate stochastic process are generated with an endpoint condition only applicable to one or a subset of variables. As Foglia (2009) notes, stress test scenarios should be coherent. In the typical example of a macroeconomic stress test involving a severe shock to GDP, one needs to generate coherent shocks to unemployment, FX, asset prices, as well as financial variables. Along the same vein, there is increasing emphasis on the integration/interaction of market risk and credit risk – conditional on a severe market shock, how would credit factors respond? Authors who have highlighted this issue include Hartmann (2010), Drehmann et al (2010), Bellini (2013), and Alessandri and Drehmann (2010). SRISK, an application examined in this paper, fits nicely in this category. SRISK in essence hinges on measuring how a financial institution's equity value would respond to a severe stock market downturn. The results presented in this paper show that our bridge sampler performs well vis-a-vis the simple rejection method employed by SRISK in a bivariate setting.

Our bridge sampler can also contribute to the literature on infill statistics in the context of estimating diffusion models. If analysts are only privy to discretely sampled observations, a "discretization bias" inevitably occurs when the Euler-Maruyama approximation is used in estimation. Pedersen (1995) proposed to reduce this bias by simulating extra points between two adjacent observations to bring the discretized version closer to the target model. These "infill" points are of course unobservable, and hence need to be integrated out in some way. The Perdersen (1995) solution is based on unconditional forward simulation to the last infill point and averages the conditional density of the endpoint evaluated at the last infill point. This solution, albeit creative, is a rather crude way of conducting a bias correction. The literature has expanded since to include the use of the Gaussian bridge as in Durham and Gallant (2002) and some importance sampler as in Elerian, *et al* (2001). Further reading on this literature can be found in Kleppe, *et al* (2014).

All in all, we add to the growing literature on density-tempered SMC with a unique methodological development with many potential applications. While our focus and the example are geared towards finance/economics, this new bridge sampler can certainly be applied to a wider range of problems in other disciplines.

2 A non-Gaussian bridge sampler based on density-tempered SMC

The non-Gaussian bridge sampler proposed here borrows the idea of density-tempered SMC in the Bayesian statistics literature such as Del Moral, *et al* (2006) and Duan and Fulop (2015) where density-tempering is used to bridge the far apart prior and posterior distributions. Here we choose bridge-sampling the GARCH model to describe the method and demonstrate it in action. The GARCH model exhibits an attractive nonlinear dynamic through volatility, and has been widely applied to model financial time series. We focus on the GJR-GARCH version by Glosten, *et al* (1993) because of its popularity and ability to capture asymmetric volatility responses. But our bridge sampler is generically applicable to a wide range of non-Gaussian stochastic processes.

Let the logarithmic asset price at time t be denoted by X_t. Assume that the asset return (continuously compounded) follows the GJR-GARCH model and mean is allowed to be autoregressive of first order:

$$
\begin{aligned}
X_t &= \mu + \rho X_{t-1} + \epsilon_t && (1)\\
\sigma_t^2 &= \omega + (\alpha + \gamma I_{t-1})\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 && (2)\\
I_{t-1} &= \begin{cases} 0 & \text{if } \epsilon_{t-1} \geq 0 \\ 1 & \text{if } \epsilon_{t-1} < 0 \end{cases} && (3)\\
\epsilon_t \mid \mathcal{F}_{t-1} &\sim D(0, \sigma_t^2) && (4)
\end{aligned}
$$

where ϵ_t, conditional on the information set $\mathcal{F}_{t-1}$ at time $t-1$, follows some distribution with mean 0 and variance σ_t^2. Denote by $f(X_t \mid X_{t-1}, \sigma_t)$ the conditional density function, which can be either parametric or non-parametric.

We wish to generate random paths, $\mathbf{X} = (x_0, X_1, X_2, \cdots, X_T)$, subject to the condition that $X_T \in A$ where A is a set. If A is a set with a small probability of occurrence, say, some level below x_0, a straightforward simulation-rejection method may be applied albeit inefficient. When A is a singleton, say, at exactly some level below x_0, bridge sampling seems to be the only sensible way to generate a sample, because a singleton set has a zero probability of occurrence. Our proposed bridge sampler can deal with both situations.

2.1 Initialization by a pseudo-Gaussian bridge

Given the starting value x_0 at time 0, one first samples the endpoint $X_T \in A$ with an easy-to-sample distribution, $g(X_T \mid x_0, \sigma_1, A)$. When A is a singleton, this distribution trivially equals to 1 at the only permissible endpoint. We will discuss in a later section how to choose $g(X_T \mid x_0, \sigma_1, A)$ when A is an interval.

Once X_T is in place, a pseudo-Gaussian bridge sampler is applied to generate $\{X_1, X_2, ..., X_{T-1}\}$ sequentially, knowing that it is a sample from a wrong but easy-to-sample bridge. This pseudo-Gaussian bridge utilizes the conditional mean and variance implied by the Gaussian bridge. However, motivated by the fact that the target GARCH model has heavy tails, the bridge samples are

generated with t distribution. The pseudo-Gaussian bridge also takes into account the conditional volatility implied by the target GARCH model as one moves along a sample path, i.e., σ_t, so that it more closely resembles the target model. In other words, advance the system from $t = 1$ to $T-1$ one time point at a time using the following system:

$$\begin{aligned} X_t &= \alpha_t(X_{t-1}) + \beta_t X_T + u_t & (5)\\ \alpha_t(X_{t-1}) &= \mu + \rho X_{t-1} - \beta_t \left(\frac{1-\rho^{T-t+1}}{1-\rho}\mu + \rho^{T-t+1} X_{t-1} \right) & (6)\\ \beta_t &= \rho^{T-t} \frac{1-\rho^2}{1-\rho^{2(T-t+1)}} & (7)\\ \zeta_t^2 &= k_b \frac{1-\rho^2-\beta_t^2(1-\rho^{2(T-t+1)})}{1-\rho^2}\sigma_t^2 & (8)\\ u_t &\sim \zeta_t \sqrt{\frac{\nu_b-2}{\nu_b}} t(0,\nu_b) \;\; i.i.d. & (9) \end{aligned}$$

where $t(0,\nu_b)$ is t distribution with ν_b degrees of freedom and $\nu_b > 2$ to ensure a finite variance. Note that the terms $\alpha_t(X_{t-1})$ and β_t come from the standard Gaussian bridge results. For the variance term (8), we made the design choice to introduce a multiplier k_b to increase the span of the Gaussian bridge draws because the target GARCH has a known stochastic mixing feature that increases the distribution's tail thickness. Also, instead of scaling up a constant volatility as in standard Gaussian bridge, σ_t^2 is used to track the volatility of the target model and is updated according to equation (2) recursively.

If $\nu_b = \infty$, the pseudo-Gaussian bridge becomes a sequence of standard Gaussian bridge draws based on changing volatilities. In the limiting case of $\rho = 1$, in other words no mean reversion, the above set of equations can be simplified to:

$$\begin{aligned} X_t &= \frac{T-t}{T-t+1} X_{t-1} + \frac{1}{T-t+1} X_T + u_t & (10)\\ \zeta_t^2 &= k_b \frac{T-t}{T-t+1}\sigma_t^2 & (11)\\ u_t &\sim \zeta_t \sqrt{\frac{\nu_b-2}{\nu_b}} t(0,\nu_b) \;\; i.i.d. & (12) \end{aligned}$$

Each simulated path is essentially a (high-dimensional) point and should be associated with an importance weight. The log-density of the path under the pseudo-Gaussian bridge sampler is:

$$\begin{aligned} \ln \mathcal{L}_{PGbridge}(\mathbf{X}) &= (T-1)\ln \frac{\Gamma(\frac{\nu_b+1}{2})}{\sqrt{\pi(\nu_b-2)}\Gamma(\frac{\nu_b}{2})} - \sum_{t=1}^{T-1} \left[\ln \zeta_t + \frac{\nu_b+1}{2} \ln\left(1 + \frac{u_t^2}{\zeta_t^2(\nu_b-2)}\right)\right] \\ &\quad + \ln g(X_T \mid x_0, \sigma_1, A). \qquad (13) \end{aligned}$$

Note that the first term on the right-hand side of the above equation is a norming constant and need not be evaluated for the algorithm, because converting importance weights into probabilities will naturally remove the norming constant.

The log-density of this path under the true GJR-GARCH bridge with the conditional density function, $f(X_t \mid X_{t-1}, \sigma_t)$, is:

$$\ln \mathcal{L}_{garch}(\mathbf{X}) = \sum_{t=1}^{T} \ln f(X_t \mid X_{t-1}, \sigma_t) - \ln Prob(X_T \in A \mid X_0, \sigma_1). \tag{14}$$

The simulated path is naturally associated with the importance weight, $\frac{\mathcal{L}_{garch}(\mathbf{X})}{\mathcal{L}_{PGbridge}(\mathbf{X})}$, and for which $Prob(X_T \in A \mid X_0, \sigma_1)$ also becomes an irrelevant constant because it only depends on x_0, σ_1, A and the parameters of the target GARCH model which are fixed for all simulated paths. Normalizing importance weights into probabilities for different simulated paths cancels out this term.

The sample generated by the above scheme is legitimate when being coupled with correct importance weights, but the sample would have a very poor numerical quality because the importance weights are highly uneven across different simulated paths, translating into a very low effective sample size (ESS[3]). Thus, the sample must undergo further refinements, which will be explained in detail later.

2.2 Tempering importance weights sequentially

Due to the high dimensionality of $\mathbf{X}$ and its non-Gaussian dynamic structure, the sample of simulated paths from the above section 2.1 generally exhibit a low ESS. A straightforward application of importance weights is clearly not advisable.[4] As stated earlier, our solution borrows the idea of density-tempered SMC such as Del Moral, *et al* (2006) and Duan and Fulop (2015) and treats the pseudo-Gaussian bridge like the prior distribution and the target GARCH bridge as the posterior distribution.

Consider an intermediate target where the density is given by[5]:

$$f_\delta(\mathbf{X}) \propto \mathcal{L}_{PGbridge}(\mathbf{X}) \left(\frac{\mathcal{L}_{garch}(\mathbf{X})}{\mathcal{L}_{PGbridge}(\mathbf{X})} \right)^\delta \tag{15}$$

One starts with $\delta = 0$, which represents the distribution under the pseudo-Gaussian bridge. Through a sequence of intermediate targets $0 < \delta_1 < \delta_2 < ... < 1$, one can arrive at $\delta = 1$, which is exactly the intended GARCH bridge. Clearly, the sequence of intermediate targets are artificial, but nevertheless completely valid. Similar to the density-tempered SMC methodology, one begins with an easy-to-sample distribution, i.e., the pseudo-Gaussian bridge ($\delta = 0$), and traverses through a sequence of densities to the ultimate target ($\delta = 1$), i.e. the GARCH bridge with its endpoint falling in a pre-defined set.

[3]Unless otherwise stated, ESS in this paper takes on the standard definition of $1/\sum_i \omega_i^2$, where ω_i are the importance weights converted into probabilities.

[4]Our experience suggests that unless the bridge is short, such as a few time points, the pseudo-Gaussian bridge proposal is bound to be a poor path for its target GARCH bridge.

[5]The proportional relationship "$\propto$" is to accommodate the fact that the right-hand side may not be integrated to 1 and is thus off by a constant multiplier. In the case of $\delta = 0$ or 1, it is a density function by definition.

Moving from δ_i to a new intermediate target represented by δ_{i+1} can be implemented by re-weighting the paths by the ratio of the two weights.

$$W(\delta_i, \delta_{i+1}, \mathbf{X}) = \frac{f_{\delta_{i+1}}(\mathbf{X})}{f_{\delta_i}(\mathbf{X})} \propto \frac{\mathcal{L}_{PGbridge}(\mathbf{X}) \left(\frac{\mathcal{L}_{garch}(\mathbf{X})}{\mathcal{L}_{PGbridge}(\mathbf{X})} \right)^{\delta_{i+1}}}{\mathcal{L}_{PGbridge}(\mathbf{X}) \left(\frac{\mathcal{L}_{garch}(\mathbf{X})}{\mathcal{L}_{PGbridge}(\mathbf{X})} \right)^{\delta_i}} = \left(\frac{\mathcal{L}_{garch}(\mathbf{X})}{\mathcal{L}_{PGbridge}(\mathbf{X})} \right)^{\delta_{i+1} - \delta_i} \tag{16}$$

Consider the case where one moves from $\delta = 0$ to δ_1. The density for the pseudo-Gaussian bridge is $f_0(\mathbf{X}) = \mathcal{L}_{PGbridge}(\mathbf{X})$, and this target is also the basis for our initial sample. In moving to δ_1, we target an abstract intermediate, $f_{\delta_1}(\mathbf{X})$, and couple it with the importance weight of $\left(\frac{\mathcal{L}_{garch}(\mathbf{X})}{\mathcal{L}_{PGbridge}(\mathbf{X})} \right)^{\delta_1}$. Likewise, this can be generally applied to the move from δ_i to δ_{i+1} with incremental importance weight as in equation (16). Having draws from an intermediate target, one can progress to the next in the sequence.

Similar to Del Moral, *et al* (2006) and Duan and Fulop (2015), we adaptively choose the tempering sequence $\{\delta_i\}$ to ensure sufficient path diversity. Moving from δ_i, we pick δ_{i+1} so that the ESS derived from (16) stays at least as large as some preset level, say, $b_I\%$ of the sample size N.[6] Unless otherwise stated, we set $b_I = 80$ in our empirical implementation later. Note that in this manner, we directly control the variability of the importance weights and thereby avoid sample impoverishment.

We conduct resampling with the assigned importance weight, which is a common step in the SMC literature. Without resampling, weights get accumulated sequentially and will also cause sample impoverishment. Resampling and thus reassignment of equal weights to all simulated paths are essential to the success of our bridge sampler.

2.3 Boosting the empirical support

With repeated reweighting and resampling, the empirical support of the simulated paths would gradually shrink, i.e., fewer and fewer distinct paths in the sample of size N. Periodically boosting the empirical support, i.e., having more distinct paths, is thus a must for the sampling scheme to succeed. Support boosting can be accomplished by running the Metropolis-Hastings (MH) move at each δ step. Note that this MH move is not for achieving convergence as in the typical MCMC run. Rather, it is designed to increase the diversity of sample paths and thus the representativeness of the sample through the use of the correct MH kernel to maintain the intermediate target distribution. The MH move is also an essential component of the SMC methodology.

The MH move requires a proposal, and the quality of the proposal determines the acceptance rate of the proposed replacement. In contrast to the SMC in the Bayesian estimation where the

[6] When dealing a sample of equal or unequal weights, the correct ESS calculation should first collapse non-distinct paths with the compensating weight adjustment. Otherwise, the ESS is likely overstated. However, collapsing non-distinct paths require computing time, and leaving it as an inflated figure in the intermediate steps can still serve the control purpose well.

high-dimensional parameters do not have an intrinsic internal structure, the GARCH path with an ending condition naturally imposes a highly restrictive relationship among points along a path. For our problem, the obvious proposal is again the pseudo-Gaussian bridge sampler used in the initialization discussed above. However, instead of sampling the entire path, we sample and replace segments with random starting and ending times, which are shorter and therefore increase chances of being accepted.

We sample starting and ending times randomly with $1 \le t_s \le t_e \le T-1$ and attach them to the current path to be a point in an augmented space. Once t_s and t_e are in place, we sample a pseudo-Gaussian bridge to replace the segment running from t_s to t_e inclusive, in a way similar to the initialization sampling where the bridge is naturally anchored at time $t_s - 1$ and $t_e + 1$. When A is not a singleton set, we need to resample $X_T \in A$ with a reasonable probability of, say, 50% so that X_T gets rejuvenated. Let Z be an independent Bernoulli random variable determining whether the endpoint is resampled. When $Z = 1$, t_e is always set equal to $T-1$ so that the replaced segment is drawn consistently with the resampled endpoint. We also find that rejuvenating the segment near the starting point beneficial. When $Z = 0$, we set $t_s = 0$ with a probability of say 50%. In other words, let Y be another independent Bernoulli random variable, and set $t_s = 0$ if $Y = 1$.

The current path augmented by Z, starting time and ending time is denoted by $(\mathbf{X}, Z, t_s, t_e)$, whereas the new proposal is $(\mathbf{X}^*, Z^*, t_s^*, t_e^*)$. The log-density of the new path given the current path $\mathbf{X}$ while conditioning on Z, t_s and t_e is:

$$\begin{aligned} & \ln \mathcal{L}_{MH}(\mathbf{X}^* \mid \mathbf{X}; Z, t_s, t_e) \\ = & \begin{cases} \sum_{t=t_s}^{t_e} p_t(\mathbf{X}^*, t_e) & \text{if } A \text{ is a singleton set} \\ 1_{\{Z=1\}} \left[\sum_{t=t_s}^{T-1} p_t(\mathbf{X}^*, T-1) + g(X_T \mid X_{t_s-1}, \sigma_{t_s}(\mathbf{X}), A) \right] & \\ \quad +1_{\{Z=0\}} \sum_{t=t_s}^{t_e} p_t(\mathbf{X}^*, t_e) & \text{otherwise} \end{cases} \end{aligned} \tag{17}$$

where $g(X_T \mid X_{t_s-1}, \sigma_{t_s}(\mathbf{X}), A)$ is the density of the same endpoint sampler as in the initialization stage which will be described in a later section, and the density of the pseudo-Gaussian bridge sampler for $t_s \le t \le T-1$ is:

$$p_t(\mathbf{X}^*, t_e) = \ln \frac{\Gamma(\frac{\nu_b+1}{2})}{\sqrt{\pi(\nu_b-2)}\Gamma(\frac{\nu_b}{2})} - \ln \zeta_t(\mathbf{X}^*, t_e) - \frac{\nu_b+1}{2} \ln \left(1 + \frac{u_t^2(\mathbf{X}^*, t_e)}{\zeta_t^2(\mathbf{X}^*, t_e)(\nu_b-2)} \right) \tag{18}$$

where

$$\begin{aligned} u_t(\mathbf{X}^*, t_e) &= X_t^* - \alpha_{t,t_e}(X_{t-1}) - \beta_{t,t_e} X_{t_e+1} \\ \zeta_t^2(\mathbf{X}^*, t_e) &= k_b \frac{1-\rho^2-\beta_{t,t_e}^2(1-\rho^{2(t_e-t+2)})}{1-\rho^2} \sigma_t^2(\mathbf{X}^*) \\ \alpha_{t,t_e}(X_{t-1}) &= \mu + \rho X_{t-1} - \beta_{t,t_e} \left(\frac{1-\rho^{t_e-t+2}}{1-\rho} \mu + \rho^{t_e-t+2} X_{t-1} \right) \\ \beta_{t,t_e} &= \rho^{t_e-t+1} \frac{1-\rho^2}{1-\rho^{2(t_e-t+2)}}. \end{aligned}$$

Again, in the limiting case of $\rho = 1$, the above simplifies to:

$$\begin{aligned} u_t(\mathbf{X}^*, t_e) &= X_t^* - \frac{t_e - t + 1}{t_e - t + 2} X_{t-1} - \frac{1}{t_e - t + 2} X_{t_e+1} \\ \zeta_t^2(\mathbf{X}^*, t_e) &= k_b \frac{t_e - t + 1}{t_e - t + 2} \sigma_t^2(\mathbf{X}^*). \end{aligned}$$

Worth noting is the fact that $\sigma_t(\mathbf{X}^*)$, the GARCH volatility deduced from the new path, will be different from $\sigma_t(\mathbf{X})$ at time $(t_s + 1)$ and all the way to the end of the path even though $\mathbf{X}^*$ and $\mathbf{X}$ share the identical segment after t_e.

The MH acceptance probability conditioning on Z, t_s and t_e is:

$$a(\mathbf{X} \to \mathbf{X}^* \mid Z, t_s, t_e) = \min\left(1, \frac{f_\delta(\mathbf{X}^*)}{f_\delta(\mathbf{X})} \frac{\mathcal{L}_{MH}(\mathbf{X} \mid \mathbf{X}^*; Z, t_s, t_e)}{\mathcal{L}_{MH}(\mathbf{X}^* \mid \mathbf{X}; Z, t_s, t_e)}\right). \tag{19}$$

We now elaborate on how the MH moves essentially provide new draws from the intermediate target distribution. By the standard result, the Markov transition kernel conditional on (Z, t_s, t_e) becomes

$$\begin{aligned} \mathcal{K}(\mathbf{X} \to \mathbf{X}^* \mid Z, t_s, t_e) &= a(\mathbf{X} \to \mathbf{X}^* \mid Z, t_s, t_e) \mathcal{L}_{MH}(\mathbf{X}^* \mid \mathbf{X}; Z, t_s, t_e) d\mathbf{X}^* \\ &\quad + B(\mathbf{X}; Z, t_s, t_e) \delta_{\mathbf{X}}(d\mathbf{X}^*) \end{aligned} \tag{20}$$

where $\delta_{\mathbf{X}}(d\mathbf{X}^*)$ is a Dirac measure that puts unit mass at $\mathbf{X}$, and $B(\mathbf{X}; Z, t_s, t_e) = 1 - \int_{\mathbf{Y} \in \mathbf{S}} a(\mathbf{X} \to \mathbf{Y} \mid Z, t_s, t_e) \mathcal{L}_{MH}(\mathbf{Y} \mid \mathbf{X}; Z, t_s, t_e) d\mathbf{Y}$ with $\mathbf{S} = \{x_0\} \times \mathbf{R}^{T-1} \times A$ denoting the set of all possible paths. Note that $B(\mathbf{X}; Z, t_s, t_e)$ captures the probability of the Markov chain staying at $\mathbf{X}$ after one transition. Again by the standard result of the MH algorithm, $\mathcal{K}(\mathbf{X} \to \mathbf{X}^* \mid Z, t_s, t_e) f_\delta(\mathbf{X}) d\mathbf{X}$ is reversible, and hence $f_\delta(\mathbf{X})$ is a stationary density of this Markov transition kernel for any (Z, t_s, t_e).

Our sampling scheme of advancing $(\mathbf{X}, Z, t_s, t_e)$ to $(\mathbf{X}^*, Z^*, t_s^*, t_e^*)$ is conducted in the augmented space, and thus has the following Markov transition kernel: $\mathcal{K}(\mathbf{X} \to \mathbf{X}^* \mid Z, t_s, t_e) Prob\{Z^*\}$ $Prob\{t_s^*, t_e^*\}$. It follows that

$$\begin{aligned} & \sum_{Z, t_s, t_e} \int_{\mathbf{X} \in \mathbf{S}} Prob\{Z\} Prob\{t_s, t_e\} \mathcal{K}(\mathbf{X} \to \mathbf{X}^* \mid Z, t_s, t_e) f_\delta(\mathbf{X}) d\mathbf{X} Prob\{Z^*\} Prob\{t_s^*, t_e^*\} \\ = & \left(\sum_{Z, t_s, t_e} Prob\{Z\} Prob\{t_s, t_e\}\right) \left(\int_{\mathbf{X} \in \mathbf{S}} \mathcal{K}(\mathbf{X} \to \mathbf{X}^* \mid Z, t_s, t_e) f_\delta(\mathbf{X}) d\mathbf{X}\right) Prob\{Z^*\} Prob\{t_s^*, t_e^*\} \\ = & Prob\{Z^*\} Prob\{t_s^*, t_e^*\} f_\delta(\mathbf{X}^*) d\mathbf{X}^*. \end{aligned} \tag{21}$$

because the first term in the second equality of the above equals 1, and the second term becomes $f_\delta(\mathbf{X}^*) d\mathbf{X}^*$ due to the fact that $f_\delta(\mathbf{X})$ is a stationary density of $\mathcal{K}(\mathbf{X} \to \mathbf{X}^* \mid Z, t_s, t_e)$. Therefore, equation (21) implies that $Prob\{Z\} Prob\{t_s, t_e\} f_\delta(\mathbf{X}) d\mathbf{X}$ is a stationary measure for the Markov transition kernel in the augmented space. Obviously, this Markov transition kernel is also irreducible

and aperiodic by design, and hence convergence to the stationary distribution takes place with repeated MH moves.

A sample of independently simulated paths, $\{\mathbf{X}_i; i = 1, 2, \cdots, N\}$, that rightly represents the intermediate target density, $f_\delta(\mathbf{X})$, can be viewed as an independent sample generated from the stationary distribution (after marginalization to focus on $\mathbf{X}$) corresponding to the Markov transition kernel. A one-step MH move applied on all N paths will naturally produce a new sample of independent paths that also represents the same target distribution. In this sense, the MH move is used to refine the sample rather than one of many moves needed to achieve Markov Chain Monte Carlo convergence to the stationary distribution. Excessive application of MH moves is therefore unnecessary after the empirical support has been satisfactorily boosted. To achieve adequate boosting at the intermediate steps, we find it effective to apply a rule of thumb requiring each path to accept on average a fixed number of MH moves, say, twenty moves. Having multiple accepted moves is essential, because every accepted MH move simply replaces a random segment instead of the whole path.

2.4 Sample endpoints that fall in the target set

When the endpoint is not fixed, there are two places in the algorithm that the endpoint needs to be sampled – the initialization step and the MH move in the support boosting stage. We now describe how to sample endpoints falling in an intended set. Without loss of generality, we can focus on an interval, i.e., $A = \{a < X_T \leq b\}$.

The endpoint will be sampled with a truncated t distribution with ν_e degrees of freedom.[7] The mean and variance of the truncated t distribution should also reflect those implied by the target process; that is, the mean and variance used in the t sampler must suitably reflect the the length of sample path, and the mean return and the initial volatility of the target GARCH process. Let $\Psi(\cdot; \nu_e)$ be the cumulative distribution function for t random variable with ν_e degrees of freedom. The truncated t sampler used here has the following log-density for $X_T \in A$:

$$\begin{aligned} & \ln g(X_T \mid X_t, \sigma_{t+1}, A) \\ = \; & \ln \frac{\Gamma(\frac{\nu_e+1}{2})}{\sqrt{\pi(\nu_e-2)}\Gamma(\frac{\nu_e}{2})} - \ln \eta(t,T) - \frac{\nu_e+1}{2} \ln\left(1 + \frac{[X_T - \lambda(t,T)]^2}{\eta^2(t,T)(\nu_e-2)}\right) \\ & - \ln\left[\Psi\left(\frac{b-\lambda(t,T)}{\eta(t,T)}\sqrt{\frac{\nu_e}{\nu_e-2}}; \nu_e\right) - \Psi\left(\frac{a-\lambda(t,T)}{\eta(t,T)}\sqrt{\frac{\nu_e}{\nu_e-2}}; \nu_e\right)\right] \end{aligned} \tag{22}$$

where

$$\begin{aligned} \lambda(t,T) &= \rho^{T-t} x_t + \frac{1-\rho^{T-t}}{1-\rho}\mu \\ \eta^2(t,T) &= k_e \frac{1-\rho^{2(T-t)}}{1-\rho^2}\sigma_1^2. \end{aligned}$$

[7] The truncated Cauchy distribution would have been a direct alternative, since it is the limiting case of the t distribution. However, it leads to substantially lower MH acceptance rates, possibly due to the greater dissimilarity between the two distributions.

Note that the mean function, $\lambda(t,T)$, reflects the multiperiod conditional mean of the target process. To ensure that the simulated endpoint can reach a wide enough range to cover heavy tails and/or larger volatility caused by the target GARCH process, $k_e \geq 1$ is introduced into the variance function by design. In our implementation later, we apply $k_e = 1.96$.

Our algorithm for executing the above truncated t sampler comprises: (1) generating points using the truncated standardised Cauchy distribution based on inversion of its cumulative distribution function, (2) applying the standard rejection simulation technique to obtain draws from the standardised t distribution, and (3) applying a scale factor of $\eta(t,T)$ and shift factor of $\lambda(t,T)$ to the standardised t draw, with corresponding adjustments to the density function, to obtain draws for X_T. The acceptance probability for rejection sampling in step (2) is:

$$a_{t,h}(x) = \frac{1}{R}\frac{t(x;\nu_e,A')}{h(x;A')} \tag{23}$$

where $t(x;\nu_e,A')$ denotes the standardised-t distribution with degrees of freedom ν_e, truncated to the standardised set A', $h(x;A')$ denotes the standardised Cauchy distribution truncated to A', and R is the scaling factor to facilitate rejection sampling by ensuring that $t(x;\nu_e,A')$ is bounded by $Rh(x;A')$. Analytically, the local optima of $\frac{t(x;\nu_e,A')}{h(x;A')}$ occur at -1, 0 and 1 and so if any of these points lie within A', we set R accordingly to the maximum. If these points do not lie in A', then we can improve the tightness by evaluating $\frac{t(x;\nu_e,A')}{h(x;A')}$ at $|\sup A'| < -1$ or $|\inf A'| > 1$, whichever the case may be.

2.5 k-fold duplication followed by support boosting

By iteratively progressing through a series of intermediate targets, through reweighting, resampling and boosting the sample, we finally arrive at a representative sample that reflects the target GARCH bridge. Indeed, with the suitable selection of parameters, the sample provides an accurate empirical representation of the target bridge. While the algorithm as-is already provides computational efficiency vis-a-vis, say, the simple rejection sampling (i.e., simulate-discard) algorithm, we can leverage on the support boosting technique further to increase efficiency by another order of magnitude.

Suppose that we require N paths in our sample. It is not necessary for all N paths to be carried through the intermediate targets. We can simply traverse the density-tempered bridge with a small representative sample. Once at the final target distribution, "duplicate-and-boost" rounds may be conducted to rapidly multiply the sample size whilst maintaining sample diversity. This involves duplicating the representative sample by k folds, i.e., increasing the sample such that there will be k identical sets of paths. Support-boosting can then be conducted in the same fashion as described in section 2.3 to restore distinctiveness of paths. As already shown, the Metropolis-Hastings moves provide valid proposals to replace the duplicated paths.

In our implementation later, the target is 50,000 paths. Starting with a base sample of 2000 paths, we traverse the density-tempered bridge to reach the final target distribution. Two rounds

of duplicate-and-boost are conducted with the number of paths duplicated by 5 folds at each round for a total duplication factor of $5^2 = 25$ that multiplies 2,000 paths to 50,000.

Our second enhancement involves improving the pseudo-Gaussian bridge sampler by utilizing the empirical distribution of paths. When $\delta = 0$, the pseudo-Gaussian bridge sampler performs well because that is exactly the initial (intermediate) target distribution. However, when δ increases and approaches 1, its performance degrades (i.e., the MH acceptance rate decreases). This is because the target approaches that of the true GARCH bridge, and more proposals are required of the pseudo-Gaussian sampler to obtain GARCH-like paths.

Observe equations (5) to (9) and note that the pseudo-Gaussian bridge essentially provides the location and scale for X_t given X_{t-1} and X_T. We improve $\alpha_t(X_{t-1})$ and β_t by updating them with a linear regression of X_t on X_{t-1} and X_T, using the empirical distribution which has been generated via the density-tempered bridge and is currently at the final target. Bridge sampling then uses these regression-updated coefficients as the basis of the location and scale for the new path proposals.

Note that a balance needs to be struck as the regression-updated bridge sampler need not be superior to the pseudo-Gaussian bridge sampler. First, the regression coefficients are subject to sampling errors when the SMC sample used to compute these coefficients is not large enough. Second, we found the pseudo-Gaussian bridge may have higher acceptance rates when support-boosting is conducted over the intermediate stages of the density tempering sequence. For the results reported in this paper, we apply the pseudo-Gaussian bridge sampler for all intermediate density targets and use the regression-updated bridge sampler for the duplicate-and-boost rounds where $\delta = 1$.

2.6 Summary of the non-Gaussian bridge sampler

To sample N paths with the endpoint in set A, the key steps are:

1. At $\delta_0 = 0$, initialize the base sample by sampling a fraction of N paths, say N_b, with (1) simulating $X_T \in A$ using $g(X_T; 0, x_0, \sigma_1)$ as in equation (22) when A is not a singleton set, and (2) running a pseudo-Gaussian bridge sampler as in equations (5)-(9).

2. Determine the next intermediate target by selecting δ_{i+1} such that the ESS computed with incremental importance weights, $W(\delta_i, \delta_{i+1}, \mathbf{X})$ in (16), does not fall below $b_I\%$ of N_b, and $b_I = 80$ is used.

3. At the selected δ_{i+1}, resample N_b paths according to $W(\delta_i, \delta_{i+1}, \mathbf{X})$, and then reassign equal weights to all resampled paths.

4. Perform the Metropolis-Hastings updating moves to replace random segments of each path as described in Section 2.3 to boost the empirical support.

5. Repeat steps 2-4 until $\delta = 1$.

6. Duplicate-and-boost N_b until N sample paths are obtained, as per section 2.5; that is, at each round of duplication, conduct regression to update the bridge sampler's location and scale, duplicate the paths by k folds, and then boost support as per section 2.3.

3 A performance study

3.1 Performance of the density-tempered SMC bridge sampler subject to a endpoint set constraint

Our density-tempered SMC bridge sampler can be compared with a simple rejection sampler when the endpoint is required to fall in a set, i.e., A, with a non-negligible probability. The simple rejection sampler is, for example, used in the SRIK implementation. Rejection sampling is conducted by simulating GARCH paths forward and discarding those where $X_T \notin A$. We use this simple rejection sampler as the benchmark method, knowing that it works straightforwardly even if it may be inefficient.

The target GJR-GARCH model is defined by the parameters estimated to the daily S&P500 Index series from the beginning of January 1990 to the end of May 2015. The conditional distribution is assumed to be t with ν degrees of freedom, and is motivated by the common empirical finding of conditional heavy tails of the S&P500 Index return. We set $\rho = 1$ per usual practice in modeling asset prices, and the parameter estimates by the maximum likelihood principle are listed in Table 1. Although we use a conditional t distribution in this study, our bridge sampler also works for a non-parametric conditional distribution, which will be shown later.

Based on the adopted conditional distribution, our target GARCH process has the following conditional log-density:

$$\ln f(X_t \mid X_{t-1}, \sigma_t) = \ln \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\pi(\nu-2)}\Gamma(\frac{\nu}{2})} - \ln \sigma_t - \frac{\nu+1}{2} \ln \left[1 + \frac{(X_t - X_{t-1} - \mu)^2}{\sigma_t^2(\nu-2)} \right]. \tag{24}$$

The degrees of freedom, ν_b and ν_e, of the two t-distributed samplers used in the algorithm (i.e., the pseudo-Gaussian bridge and the endpoint sampler) need not equal ν, the degrees of freedom for the target model's conditional t distribution. In the performance analysis, we set $\nu_b = \nu - 1$ and $\nu_e = \nu - 4$, giving the two samplers heavier tails in order to accommodate the increased tail thickness due to stochastic mixing effect of the GARCH dynamics.

The GARCH path to be bridge-sampled is 6 months long (126 trading days). The baseline case for the set of endpoints, i.e., A, is at least 40% percentage below the starting asset value. For the logarithmic index value, the baseline case becomes $A = (-\infty, x_0 + \ln 0.6]$. The choice of this baseline case is motivated by the SRISK measure of Acharya, *et al* (2010) and Brownlees and Engle (2012). The baseline case also sets the initial conditional volatility to the stationary volatility computed with the estimated parameters. The full list of parameters for the baseline case is provided in Table 1.

Table 1: Performance Study Parameters

Bridge		
x_0	$\ln(1000)$	Initial price
A	$(-\infty, \ln(600)]$	Endpoint set
T	126 trading days	Time horizon
GJR-GARCH		
μ	4.04×10^{-4}	Drift
ρ	1	AR term
ω	1.16×10^{-6}	Constant
α	3.85×10^{-7}	ARCH term
β	0.918	GARCH term
γ	0.140	Leverage term
ν	7.69	Degrees of freedom for t distributed innovations
σ_{ann}	0.159	Implied annualized volatility (252 days)
DTSMC		
N_b	2000	Representative base sample size
ν_e	3.69	Degrees of freedom for the t-distributed endpoint sampler
k_e	1.96	Endpoint sampler multiplier
ν_b	6.69	Degrees of freedom for the t distribution in the bridge sampler
k_b	1.44	Bridge sampler multiplier
$P(Z = 1)$	0.5	Probability of proposing new right segment
$P(Y = 0)$	0.25	Probability of proposing new mid segment
$P(Y = 1)$	0.25	Probability of proposing new left segment
b_I	80	Target ESS for new delta
MH Moves	20	Avg moves req per path at intermediate targets
# of rounds	2	Number of duplicate-and-boost rounds conducted
k	5	k-fold multiplier for duplicate-and-boost
MH Moves	10	Average numbers of moves for duplicate-and-boost

For the baseline, we show in Figure 1 the QQ plots[8] of the two corresponding empirical distributions (bridge sampling vs simple rejection sampling) at three time points along the path – X_{41}, X_{83} and X_{126}. In addition, we show QQ plots of the path average, minimum and maximum values. The plots provided are based on 50,000 paths. Overall, the results lend credibility to our bridge-sampling scheme, matching draws from the simple rejection method accurately. More importantly, the bridge sampler (coded in C++) is much more efficient with the running time less than 1/10 of the time needed for simple rejection sampling (also coded in C++) for the baseline case. Since SMC is naturally parallelizable just like the simple rejection method, their relative performance should remains fairly constant when one chooses to perform parallel processing.

Our scheme is also robust, and continues to deliver high-quality samples when the shock is varied between 0% and 80% below the starting value of the path. We also alter the initial conditional volatility to various levels between the minimum and maximum conditional volatilities in the time series. Lastly, we assess the results when the time horizon is shortened and extended to 63 and 252 trading days, respectively. The results for various cases are summarized in Table 2. The QQ plots do not suggest any inconsistency in the distribution of density-tempered SMC paths as compared to rejection sampling, and are skipped to conserve space. Across the spectrum of test cases, our bridge sampler dominates the simple rejection method except for cases when the probability of landing in the endpoint set is high; such as a shock of 0% or more or when the initial volatility is at its historical maximum. However, those cases should not call for bridge sampling in the first place.

3.2 Fixed endpoint sampler and alternatives

3.2.1 Quality of the density-tempered SMC bridge sampler under a fixed endpoint

The performance comparison reported earlier shows that the density-tempered SMC bridge sampler is accurate and delivers superior computational efficiency when the endpoint is confined to a set with a low probability. Logically, the accuracy should remain true when the paths are required to end at a fixed value. However, simple rejection sampling can no longer serve as the benchmark because the probability for a path to hit a fixed endpoint is zero.

In order to ascertain the numerical quality of the density-tempered SMC bridge sampler which uses 20 accepted MH moves on average for all intermediate targets, we use the sample produced by the same bridge sampler as the benchmark except that the number of accepted MH moves on average is set to a higher value of 100, with confidence that the empirical support would have been boosted to a very high level of satisfaction.

Also, the fixed endpoint will be set equal to the starting point, i.e., $X_T = x_0$, which is the case for assessing alternative samplers later. We reduce the shock level for a reason. The alternative samplers to be discussed later are deployed in infill estimation where extreme moves over a data period are far less likely to occur under the circumstance of their applications. We compare the

[8]QQ plots shown in this paper, unless otherwise stated, are based on comparing 100 points corresponding to the percentile levels 0.5% to 99.5% in increments of 1%.

Figure 1: QQ Plots of Density-Tempered SMC Bridge Sampling vs Simple Rejection Sampling (conditional t)

Comparison of two samples of 50,000 GJR-GARCH paths generated by two sampling methods, and each path covers six months (126 trading days). All paths share the same starting value and end at 40% or lower than the starting value. The initial volatility is set equal to the stationary volatility. Simple rejection sampling takes 2858 seconds to complete whereas density-tempered SMC bridge sampling needs 230 seconds.

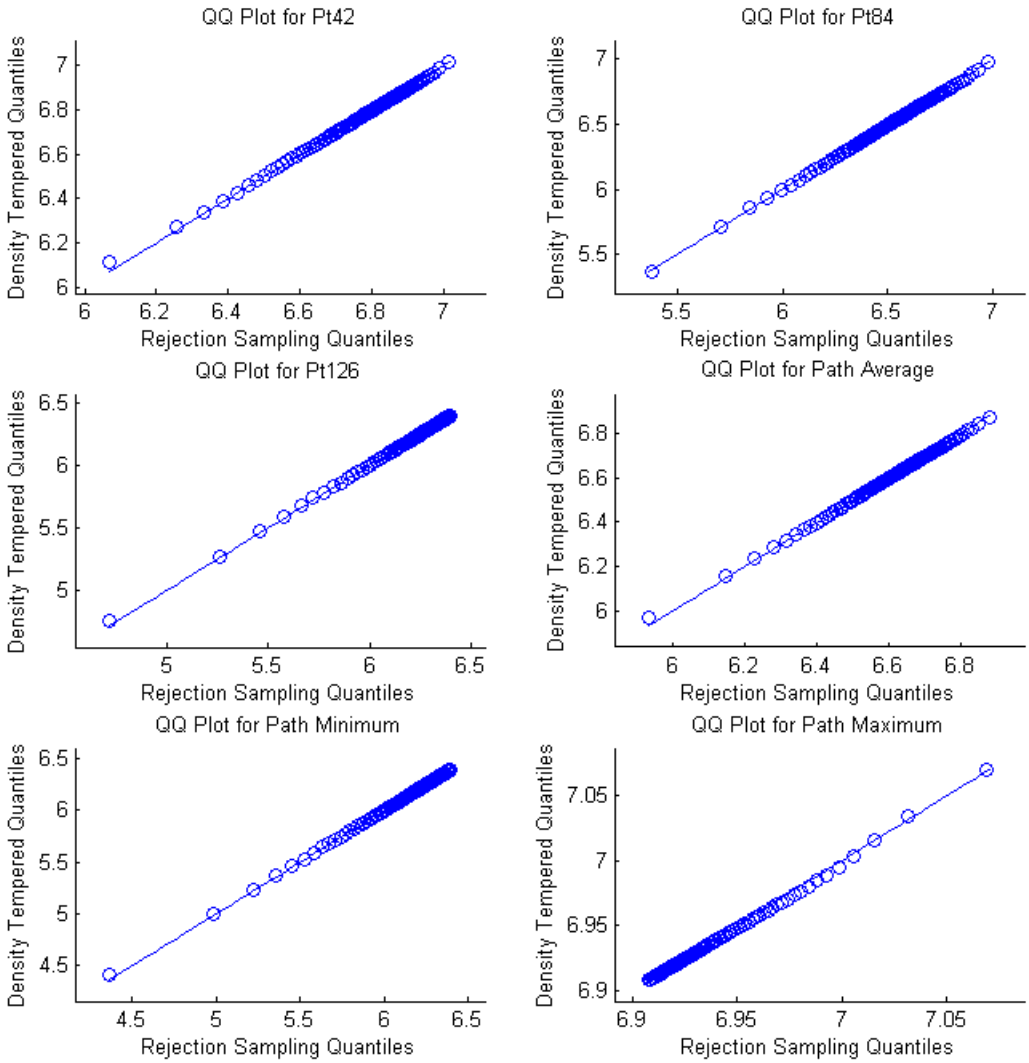

Table 2: Efficiency of Simple Rejection Sampling vs Density-Tempered SMC Bridge Sampling (conditional t)

Comparison of the computing time of the rejection sampling vs the density-tempered SMC bridge sampler. Density-tempered SMC provides efficiency gains under the majority of the test cases. The computing time is based on execution on a desktop computer with Intel i5-4460 @ 3.20 GHz, DDR3-1600MHz CL9 (4 GB x 2), and dedicated graphics.

	Rejection Sampling	**DTSMC Bridge**
	(in seconds, unless otherwise stated)	
Baseline (Figure 1)		
Shock 40%, Volatility 15.9%, T=126	2858	230
Varying Shock Level		
Shock of at least 80%	$>$ 46 h	125
Shock of at least 60%	$>$ 5 h	108
Shock of at least 20%	320	85
Shock of at least 0%	27	75
Varying Initial Volatility (annualized)		
Historical minimum 6.5%	3.0 h	282
Half of unconditional volatility 8.0%	2.4 h	277
Twice of unconditional volatility 31.8%	573	193
Historical maximum 89.0%	80	188
Varying Time Horizon		
1 quarter horizon (T=63)	1.6 h	102
1 year horizon (T=252)	2427	612

sample of 10,000 GJR-GARCH bridge paths generated in the way as before against this benchmark sample of the same size but obtained at a higher number of MH moves. The QQ plots of these two samples as in Figure 2 reveal that the two empirical distributions are essentially identical, suggesting that the density-tempered SMC bridge sampler is already accurate at 20 accepted MH moves on average.

3.2.2 Alternative samplers and their quality

The first fixed-endpoint sampler in comparison is the Pedersen sampler (1995) used in infill estimation. Here, we adopt the idea and apply it to simulate the GARCH bridge. The method can be summarized as follows: simulate the GARCH process forward per usual up to time $T-1$ ignoring the influence of the endpoint on the path. Because the path must end at x_T, the density of the last step from time $T-1$ to T solely determines the importance weight of a path, which is $f(x_T \mid X_{T-1}, \sigma_T)$.[9] The importance weights for such simulated paths will be highly uneven, simply because X_{T-1} is generated without factoring in the fact that the next point in the path must end at x_T.

The second alternative sampler being considered is the Gaussian bridge sampler (GBS) which was used in Durham and Gallant (2002) as an infill correction for discretization bias. Here, we apply it to the GJR-GARCH model. We use equations (5)-(12), except that we set $k_b = 1$ and u_t to a Gaussian random variable. In other words, we apply a standard GBS, except that we use the conditional volatility as prescribed by the GARCH dynamics, the left pin being $t-1$ and right pin T. Naturally, each sampled path must be assigned an importance weight which equals the ratio of its density under the GARCH bridge to its density under the Gaussian bridge. As we shall see, GBS fares poorly and a few paths dominate the entire sample. Note that the initialization sampler in our density-tempered SMC bridge sampler is in effect a more sophisticated version of GBS. It may serve as a starting point of an ultimate sampler, but needs to go through a sequence of revisions in order to arrive at a quality sample for a GARCH bridge.

Table 3 provides the ESS based on 10,000 sample paths produced by the Pedersen and GBS samplers, respectively. The target model is still GJR-GARCH(1,1) but with both conditional normal and t distributions estimated to the S&P500 index daily series from 2 January 1990 to 18 September 2015.

Having established the density-tempered SMC bridge sampler as a quality method, we use it as the benchmark to assess the two alternative bridge samplers. The Pedersen sampler generally performs poorly, regardless of the conditional distribution. Even with benign shocks of -10% and 10%, the ESS is below 11%. GBS appears to fare better, in the range of 33.3% to 67.9%, but only when the conditional distribution is Gaussian. When a heavier-tailed distribution like t is involved,

[9] When A is a set containing more than one element, the Pedersen sampler needs to further simulates X_T using the conditional distribution at time $T-1$ under the GARCH model and factoring in the conditioning set, $X_T \in A$. The importance weight becomes $Prob(X_T \in A \mid X_{T-1}, \sigma_T)$, which arises from conditioning at time $T-1$.

Figure 2: QQ Plots for 20 vs 100 MH Moves under Density-Tempered SMC Bridge Sampling with a Fixed Endpoint

Comparison of two samples of 10,000 GJR-GARCH paths generated by the density-tempered SMC bridge sampler. All paths share the same starting value and end at the same point as the starting value. Paths produced for both 20 and 100 MH moves use the baseline parameter provided earlier.

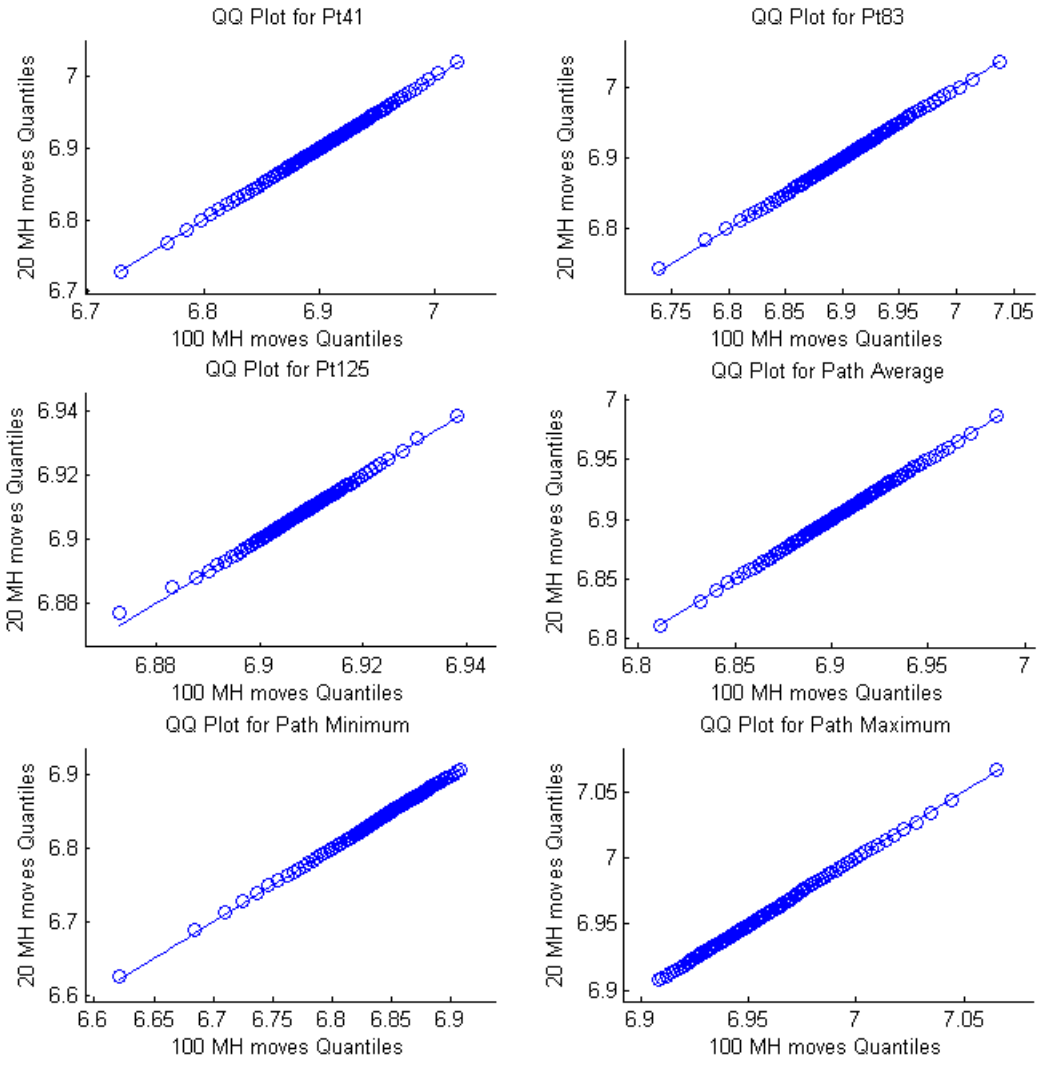

which is more reasonable for applications, the ESS falls drastically to between 3.38% and 5.03%, which is worse than the Pedersen sampler.

The overall low ESS implies that important weights are highly uneven, or stated differently, a few paths dominate the entire sample. This is illustrated in the QQ plots of Figure 3, where a sample of 100 paths is used to plot against the quantitles produced by the density-tempered SMC bridge sampler.[10] Many of the GARCH paths are essentially repeats of each other. Overall, our results suggest that the Pedersen and GBS samplers are inappropriate for simulating GARCH bridges, and by extension, are unlikely to work well for other non-Gaussian stochastic processes.

Table 3: Effective Sample Sizes (ESS) of Alternative Samplers

	Price Change		
	-10%	0%	10%
Conditional Gaussian			
Pedersen	4.57%	10.9%	10.4%
GBS	33.3%	67.9%	65.6%
Conditional *t*			
Pedersen	3.00%	7.71%	10.1%
GBS	3.38%	4.18%	5.03%

4 Application - Computing SRISK

Our sampler can be used in a wide range of applications, and here we demonstrate it with the particular example of SRISK computation. SRISK is a value critically hinges on the long-run marginal expected shortfall (LRMES), where LRMES of a financial firm is defined as its expected return conditional on a systemic event. We understand that V-Lab defines this event as a negative shock of at least 40% to the S&P500 index[11] over a period of 6 months (126 days). SRISK is basically LRMES translated to a dollar value of additional equity capital required through factoring in a financial institution's capital structure.

Computing LRMES requires simulation of market prices under stress. Our bridge sampler needs to accommodate several differences in the specific SRISK implementation. First, instead of using conditional t innovations, LRMES computation uses a non-parametric conditional distribution. The density-tempered SMC bridge sampler can accommodate non-parametric conditional distributions, and will be demonstrated shortly. Second, we need to simulate a bivariate system based on the

[10] We chose to show a sample size of 100 to illustrate the flat portions of the QQ plot, indicating that a few actual paths dominate the entire sample, leading to low ESS. This effect exists at 10,000 paths, but is harder to detect visually.

[11] The actual implementation of SRISK for US financial firms is based on the SPDR S&P500 ETF which tracks the S&P500 index very closely. We have used the latter to be consistent with the earlier sections of our paper.

Figure 3: QQ Plot of the Gaussian Bridge Sampler and Pedersen Samplers against the Density-Tempered SMC Bridge Sampler

Comparison of samples of 100 GJR-GARCH paths generated by the Gaussian bridge and Pedersen samplers against the target quantiles obtained by the density-tempered SMC bridge sampler. All paths share the same starting value and end at the same point as the starting value. Circles represent the Pedersen sampler while the crosses represent the Gaussian bridge sampler.

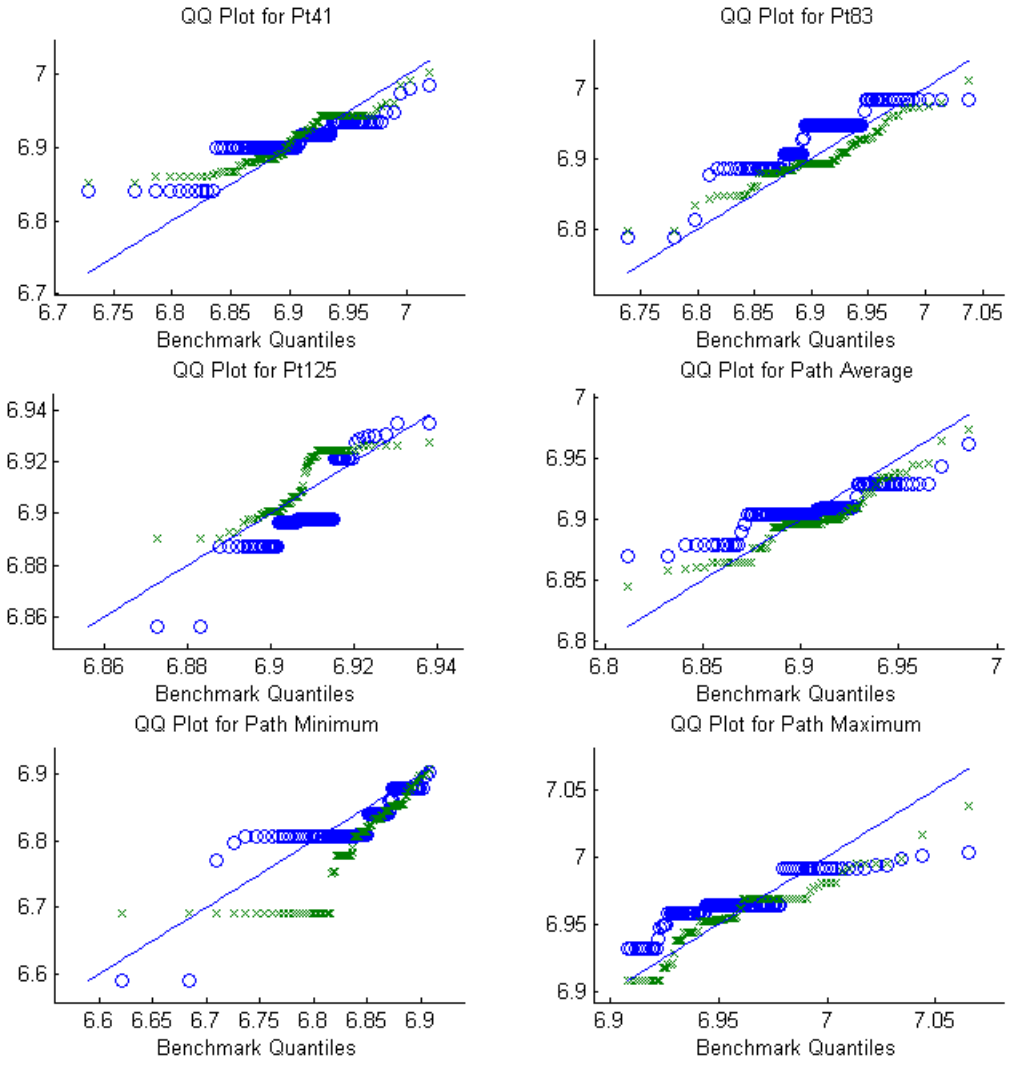

GARCH-DCC model of Engle (2002) while making one stochastic process subject to an endpoint condition. This is because LRMES is the expected return of a financial institution given a severe S&P500 index shock.

4.1 Non-parametric innovations

We first demonstrate the ability of the density-tempered bridge sampler to handle non-parametric innovations. By "non-parametric", we are referring to the distribution of ϵ_t in equation (4). In other words, ϵ_t is no longer t-distributed, and is instead constructed from the empirical distribution of the standardized residuals. As per usual notation, let $\epsilon_t = \sigma_t z_t$, where z_t is the standardised innovation with mean 0 and variance 1.

Our benchmark is the rejection sampler, which we implement by extracting z_t from the time series of the S&P index using the estimated GJR-GARCH model where the estimation is conducted with the quasi-maximum likelihood principle requiring no commitment to a specific conditional distribution. Following the convention of setting $\mu = 0$ as per SRISK and estimating the remaining parameters, the sample of standardised residuals based on the S&P500 index series shows that it has a negative skewness of -0.4132 and excess kurtosis of 1.588, which could potentially pose a model fitting issue for a symmetric heavy-tailed distribution like t, which in a way motivates its use in SRISK.

For the rejection sampler, we consider the discrete empirical distribution, which is obtained by bootstrapping the standardized residuals. Whenever an innovation is required for the next price point, we randomly draw from these residuals with equal probabilities to obtain z_t. z_t is then scaled by the conditional volatility at that point to obtain the innovation $\epsilon_t = \sigma_t z_t$. The rejection sampler performs about four times as fast as the case of conditional t-distribution when the endpoint condition is for the final price drops by 40% or more. Aside from the differences in parameter estimates, we identify two key reasons. First, the error distribution is now negatively skewed which increases the probability of reaching the endpoint set A (0.386% vs 0.244%). Second, sampling the standardized residuals is faster because it simply picks from the pool of stored residuals instead of drawing from a t-distribution.

On the other hand, handling a non-parametric distribution is more onerous for the density-tempered SMC bridge sampler. Recall equation (14) where the log-density under the target GARCH model is computed. The importance weights are then formed as the ratio of the GARCH density to the sampler's density as per equation (16). In other words, we need to estimate the density function for the non-parametric distribution. We do so using kernel density estimation (KDE). Our implementation of KDE uses a Gaussian kernel with bandwidth selected as per Sheather and Jones (1991). We use the sample of standardized residuals available from 2 January 1990 to 18 September 2015 to pre-compute density values over a mesh of 2048 points spread equally over a domain of $[\min\{z_t\} - 5, \max\{z_t\} + 5]$, where $\{z_t\}$ are the stored standardized residuals.

The density-tempered SMC bridge sampler continues to perform decently in term of accuracy as compared to the simple rejection sampler, which is revealed in the QQ plots in Figure 4. Slight

differences exist for a few possible reasons. First, bootstrapping uses discrete points but the density-tempered SMC bridge sampler uses the continuous kernel density function. Consequently, the support of our standardised residuals spans a continuous domain rather than selecting from a set of 6481 discrete points. Second, as mentioned in the previous paragraph, we work with a wider domain to accommodate larger shocks which may be proposed by our bridge sampler.

Our density-tempered SMC bridge sampler also performs well in terms of computing time. For the baseline case studied in the earlier section, producing 50,000 paths takes 372 seconds, compared to 767 seconds taken by the rejection sampler, or roughly twice as fast. Naturally, we expect greater advantages in speed when conditional volatilities are low.

4.2 Simulation of a bivariate GARCH-DCC model

As per Brownlees and Engle (2012), both the processes for the S&P index and the financial institution are assumed to follow GJR-GARCH with dynamic conditional correlation. In other words, equations (1) to (4) apply to both processes, and μ is set to 0 for both the S&P500 index and the financial firm. The conditional distribution D in equation (4) is jointly non-parametric:

$$\begin{bmatrix}\epsilon_{mt}\\ \epsilon_{it}\end{bmatrix}\Bigg| F_{t-1} \sim D\left(\mathbf{0}, \begin{bmatrix}\sigma_{it}^2 & r_{it}\sigma_{it}\sigma_{mt}\\ r_{it}\sigma_{it}\sigma_{mt} & \sigma_{mt}^2\end{bmatrix}\right) \tag{25}$$

where the sub-indices i and m refer to the financial institution and the market, respectively. r_{it} refers to the dynamic conditional correlation between i and m at time t, which is updated according to a companion pseudo-correlation matrix Q_{it} to ensure positive semi-definiteness:

$$\text{Corr}\begin{bmatrix}\epsilon_{mt}\\ \epsilon_{it}\end{bmatrix} = \begin{bmatrix}1 & r_{it}\\ r_{it} & 1\end{bmatrix} = \text{diag}(Q_{it})^{-\frac{1}{2}} Q_{it} \text{diag}(Q_{it})^{-\frac{1}{2}} \tag{26}$$

$$Q_{it} = (1-\alpha_C-\beta_C)S_i + \alpha_C \begin{bmatrix}\epsilon_{mt-1}\\ \epsilon_{it-1}\end{bmatrix}\begin{bmatrix}\epsilon_{mt-1}\\ \epsilon_{it-1}\end{bmatrix}' + \beta_C Q_{it-1} \tag{27}$$

where α_C and β_C are the DCC coefficients and S_i is the unconditional correlation matrix of the financial institution and market index. Engle (2009) provides the details for estimation of this system. Table 4 shows the estimated parameters for four selected US financial institutions.

Next, consider a pair $\{z_{mt}, \xi_{it}\}$, where ξ_{it} are standardised, linearly orthogonal shocks of the financial institution to the market on day t and computed as:

$$\xi_{it} = \frac{(z_{it} - r_{it}z_{mt})}{\sqrt{1-r_{it}^2}} \tag{28}$$

To obtain LRMES, we first obtain 50,000 simulated paths of $\{X_{mt}\}$ that satisfy $X_{m,126} \leq X_{m,0} + \ln(0.6)$ via the sampler. The implied shocks ϵ_{mt} and standardized residuals $z_{mt} = \epsilon_{mt}/\sigma_{mt}$

Figure 4: QQ Plots of Density-Tempered SMC Bridge Sampling vs Simple Rejection Sampling (non-parametric conditional distribution)

Comparison of two samples of 50,000 GJR-GARCH paths generated by two sampling methods, the rejection sampler based on continuous empirical distribution and the density tempered sampler. All paths cover six months (126 trading days), have the same starting value, and end at 40% or lower than the starting value. The initial volatility is set equal to the stationary volatility.

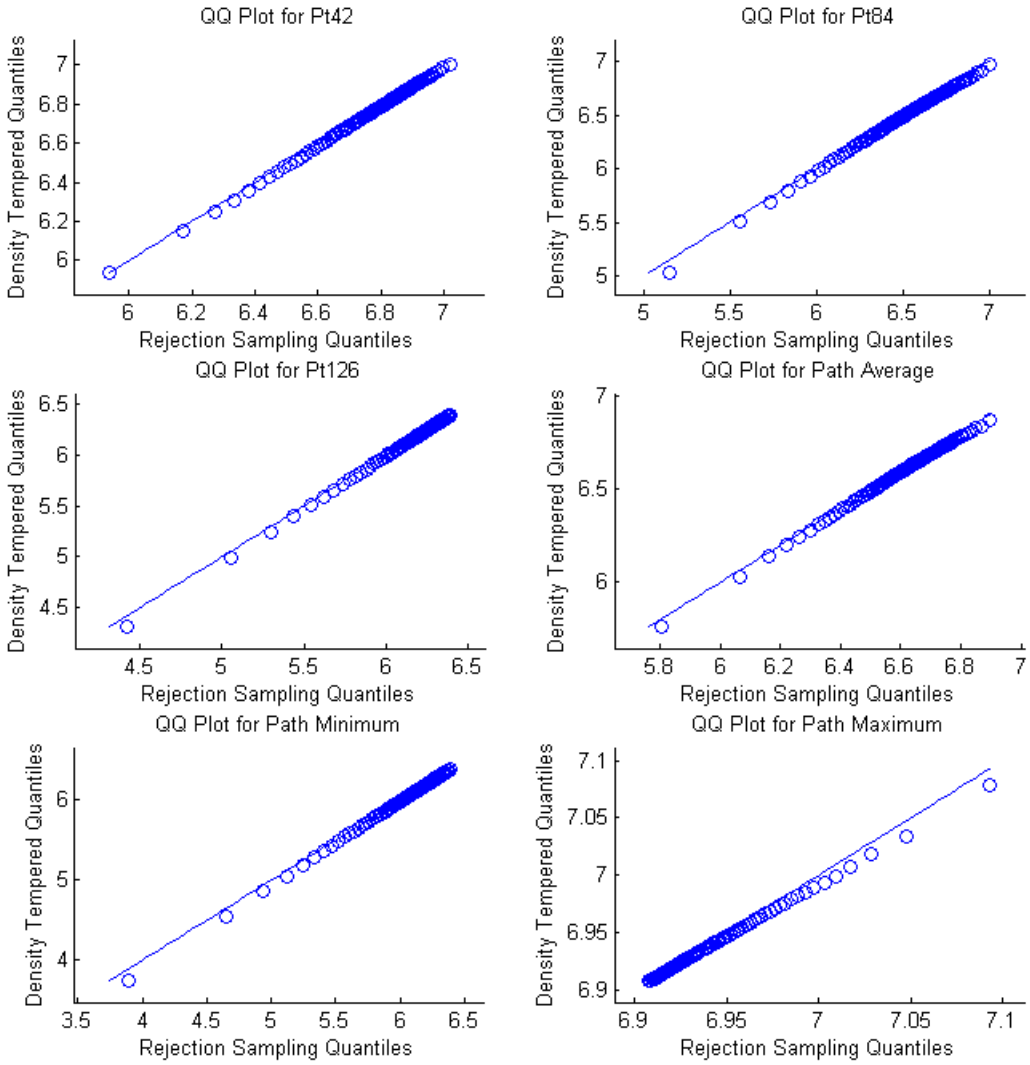

Table 4: GARCH-DCC Parameter Estimation for Four US Financial Firms

	S&P500	Citibank	BOA	JPM	WellsFargo
ω_i ($\times 10^{-6}$)	1.67	2.10	2.95	2.15	1.36
α_i	9.68×10^{-7}	0.0356	0.0388	0.0191	0.0200
β_i	0.913	0.936	0.930	0.941	0.940
γ_i	0.144	0.0542	0.0532	0.0783	0.0771
α_C		0.0295	0.0272	0.0225	0.0265
β_C		0.957	0.958	0.973	0.967
σ_{i1}	0.0145	0.0219	0.0228	0.0212	0.0219
r_{i1}		0.824	0.789	0.844	0.830

are inferred from each of these paths. For each simulated z_{mt}, a correlated z_{it} is generated according to:

$$z_{it} = \sqrt{1 - r_{it}^2}\xi_{it} + r_{it} z_{mt} \tag{29}$$

where ξ_{it} is a corresponding orthogonal shock component for the financial institution.

We consider two cases of LRMES, one assuming only linear correlations, and the other taking into account non-linear dependency. For the first case, since we assume only linear correlation in the shocks, r_{it} captures the correlation behaviour fully and we simply pick ξ_{it} randomly from the stored set of shocks whenever it is required. For the second case, since non-linear dependency is not explicitly modelled, we pick the corresponding ξ_{it} from the pairs of $\{z_{mt}, \xi_{it}\}$ stored earlier.

For the non-linear dependency case, there is a nuance. The rejection sampler draws discretely from the stored $\{z_{mt}, \xi_{it}\}$ and there is a one-to-one matching of ξ_{it} for each z_{mt}. On the other hand, the density-tempered sampler produces continuous $\{z_{mt}\}$, and will naturally generate $\{z_{mt}\}$ which need not be constrained to historical observations. Accordingly, given z_{mt}, we need to choose the corresponding ξ_{it} appropriately. This is done by randomly choosing between the two stored residuals (on the left and right) which are closest to z_{mt}, with probabilities of choosing either of them based on their distance to z_{mt}, assigning the closer one with a higher probability (for example, 1/4 vs 3/4 yield probabilities of 3/4 vs 1/4). Since we apply KDE in the density tempered sampler, there will be cases where z_{mt} exceeds the bounds of the stored set of standardized residuals. In these cases, we simply select the residual at the boundary. There are also cases of multiple identical observations of z_{mt} but different ξ_{it}, in particular $z_{mt} = 0$, to pick from. We select one of these observation randomly with equal probabilities. Once the corresponding ξ_{it} is determined, it is used to update the price path for the financial institution.

For computing LRMES, we continue to use the parameters as calibrated to the data sample from 2 January 1990 to 18 September 2015. But the conditional volatilities and dynamic correlation

take their values as of September 18, 2015.[12] We simulate the bivariate process and compute the LRMES for each financial institution according to:

$$LRMES_i = \frac{1}{N}\sum_{p=1}^{N}\left[1 - e^{\sum_{t=1}^{126}\epsilon_{it,p}}\right] \tag{30}$$

where i refers to the financial institution, t to the time point, and N to the total number of paths.

Table 5 below provides a comparison of the LRMES measures computed under the rejection sampler and density-tempered SMC bridge sampler. Overall, the numbers match well in terms of magnitude and are comparable. As we have alluded to earlier, the differences which exist are likely to be attributable to the fact that the density tempered sampler operates in a continuous domain while the rejection sampler uses discrete values.[13] Arguably, simulating with KDE even in terms of rejection sampling would yield a smoother estimate of LRMES, and thus a better estimate than that from bootstrapping the standardized residuals.

Table 5: LRMES estimates under the Rejection Sampler and DTSMC Sampler (LRMES in percentage points and standard errors in parentheses)

	Citibank	**BOA**	**JPM**	**WellsFargo**
Linear Dynamic Correlation				
Rejection Sampler	48.698	45.093	53.310	48.849
	(0.092)	(0.091)	(0.087)	(0.095)
Density-Tempered Sampler	49.129	45.430	53.777	49.145
	(0.094)	(0.090)	(0.087)	(0.095)
Non-Linear Dependency				
Rejection Sampler	48.477	44.161	50.905	48.650
	(0.096)	(0.097)	(0.089)	(0.097)
Density-Tempered Sampler	48.920	43.004	50.470	49.351
	(0.135)	(0.094)	(0.089)	(0.097)

5 Conclusion

We propose a new bridge sampling scheme for efficiently generating paths based on a general non-Gaussian stochastic process, particular the GARCH process. The endpoint of such paths can be either a set of points or a singleton. This bridge sampling scheme relies on SMC and specifically

[12] Although this is a point of relatively high volatility, the DTSMC is still faster and produces 50,000 paths in 349 seconds compared to 498 seconds taken by the rejection sampler.

[13] While both our samplers are comparable, they do not match the published LRMES from NYU-Stern V-Lab in terms of magnitude. However, their loose rankings are aligned. While we have attempted to follow the stated SRISK setup as far as possible, differences are still likely to exist.

density-tempered SMC to facilitate smooth transition across a sequence of intermediate targets so as to deliver exact draws from the target bridge. Careful control of the ESS is essential and is accomplished by reweighting, resampling and support boosting via the Metropolis-Hastings move at each stage of the sequence. As applied to generating a GARCH bridge defined by a stress scenario of an at least 40% negative shock to the S&P 500 index value, our bridge sampling scheme is accurate and delivers computational efficiency of over 10 times in comparison with simple rejection sampling for a baseline. When the path's endpoint is fixed, the density-tempered SMC bridge sampler is also far superior to the alternative bridge samplers. This bridge sampling method performs well under a range of parameter values and for both parametric and non-parametric conditional distributions.

The SRISK application shows that this new bridge sampling scheme makes stress scenario generation computationally efficient. The relevance of our bridge sampling scheme is its ability to efficiently generate paths of a non-Gaussian bridge under severe stress scenarios, and this has many implications. The Bank for International Settlements in its "Principles for Sound Stress Testing Practices and Supervision (2009)" notes that "most bank stress tests were not designed to capture the extreme market events that were experienced", and that "scenarios tended to reflect mild shocks". With this new bridge sampler, one can better cater to the needs for more onerous stress test scenarios with much reduced computational burden.

References

[1] Acharya, V., Engle, R., and Richardson, M. (2012). Capital shortfall: A new approach to ranking and regulating systemic risks. *American Economic Review* 102(3), 59-64.

[2] Alessandri, P., and Drehmann, M. (2010). An economic capital model integrating credit and interest rate risk in the banking book. *Journal of Banking and Finance* 34(4), 730-742.

[3] Andersen, L. (2008). Simple and efficient simulation of the Heston stochastic volatility model. *Journal of Computational Finance* 11(3), 42.

[4] Basel Committee on Banking Supervision. (2009). Principles for sound stress testing practices and supervision. Bank for International Settlements.

[5] Bellini, T. (2013). Integrated bank risk modeling: A bottom-up statistical framework. *European Journal of Operational Research* 230(2), 385-398.

[6] Bladt, M., and Sorensen, M. (2014). Simple simulation of diffusion bridges with application to likelihood inference for diffusions. *Bernoulli* 20(2), 645-675.

[7] Board of Governors of the Federal Reserve System. (2009). The Supervisory Capital Assessment Program: Design and Implementation.

[8] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31(3), 307-327.

[9] Borio, C.E., and Drehmann, M. (2009). Towards an operational framework for financial stability: 'fuzzy' measurement and its consequences.

[10] Borio, C., Drehmann, M., and Tsatsaronis, K. (2014). Stress-testing macro stress testing: does it live up to expectations?. *Journal of Financial Stability* 12, 3-15.

[11] Boss, M., Fenz, G., Krenn, G., Pann, J., Puhr, C., Scheiber, T., ... and Ubl, E. (2008). Stress tests for the Austrian FSAP update 2007: Methodology, scenarios and results. OeNB Financial Stability Report, 15, 68-92.

[12] Brownlees, C.T., and Engle, R.F. (2012). Volatility, correlation and tails for systemic risk measurement. Available at SSRN 1611229.

[13] Cihk, M. (2007). Introduction to applied stress testing. IMF Working Papers, 1-74.

[14] Del Moral, P., Doucet, A., and Jasra, A. (2006). Sequential monte carlo samplers. *Journal of the Royal Statistical Society Series B* 68(3), 411-436.

[15] Drehmann, M. (2005). A market based macro stress test for the corporate credit exposures of UK banks. In BCBS Seminar – Banking and Financial Stability: Workshop on Applied Banking Research.

[16] Drehmann, M., Sorensen, S., and Stringa, M. (2010). The integrated impact of credit and interest rate risk on banks: A dynamic framework and stress testing application. *Journal of Banking and Finance* 34(4), 713-729.

[17] Duan, J.-C., and Fulop, A. (2015). Density-Tempered Marginalized Sequential Monte Carlo Samplers. *Journal of Business and Economic Statistics* 33(2), 192-202.

[18] Durham, G.B., and Gallant, A.R. (2002). Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business and Economic Statistics* 20(3), 297-338.

[19] Elerian, O., Chib, S., and Shephard, N. (2001). Likelihood inference for discretely observed nonlinear diffusions. *Econometrica* 69(4), 959-993.

[20] Engle, R. (2002). Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics* 20, 339-350.

[21] Engle, R. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of UK inflation. *Econometrica* 50, 987-1008.

[22] Engle, R. (2009). Anticipating Correlations: A New Paradigm for Risk Management: A New Paradigm for Risk Management. Princeton University Press.

[23] European Banking Authority. (2014). Results of 2014 EU wide stress test.

[24] Foglia, A. (2008). Stress testing credit risk: a survey of authorities' approaches. Bank of Italy Occasional Paper (37).

[25] Glosten, L., Jagannathan, R., and Runkle, D. (1993). Relationship between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 1779-1801.

[26] Hartmann, P. (2010). Interaction of market and credit risk. *Journal of Banking and Finance* 34(4), 697-702.

[27] Hirtle, B., and Lehnert, A. (2014). Supervisory stress tests. FRB of New York Staff Report, (696).

[28] Kleppe, T. S., Yu, J., and Skaug, H. J. (2014). Maximum likelihood estimation of partially observed diffusion models. *Journal of Econometrics* 180(1), 73-80.

[29] Moretti, M., Stolz, S.M., and Swinburne, M. (2008). Stress Testing at the IMF. International Monetary Fund.

[30] Papaspiliopoulos, O., and Roberts, G. (2012). Importance sampling techniques for estimation of diffusion models. In Statistical Methods for Stochastic Differential Equations (M. Kessler, A. Lindner and M. Srensen, eds.). Monogr. Statist. Appl. Probab. 124 311340. Boca Raton, FL: CRC Press.

[31] Pedersen, A.R. (1995). A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations. *Scandinavian Journal of Statistics* 22(1), 55-71.

[32] Robert, C. (1995). Simulation of truncated normal variables. *Statistics and Computing* 5, 121-125.

[33] Rogers, L.C.G., and Williams, D. (2000). Diffusions, Markov Processes and Martingales: Volume 2, Ito Calculus. Cambridge University Press.

[34] Sorge, M. (2004). Stress-testing financial systems: an overview of current methodologies. Bank for International Settlements working paper, 165.

[35] Virolainen, K. (2004). Macro stress testing with a macroeconomic credit risk model for Finland. Bank of Finland discussion paper, 18.