

Utah State University

DigitalCommons@USU

All Graduate Plan B and other Reports

Graduate Studies

5-2019

Predictive Distributions via Filtered Historical Simulation for Financial Risk Management

Tyson Clark
Utah State University

Follow this and additional works at: <https://digitalcommons.usu.edu/gradreports>



Part of the [Finance and Financial Management Commons](#), [Portfolio and Security Analysis Commons](#), [Probability Commons](#), and the [Statistical Models Commons](#)

Recommended Citation

Clark, Tyson, "Predictive Distributions via Filtered Historical Simulation for Financial Risk Management" (2019). *All Graduate Plan B and other Reports*. 1368.
<https://digitalcommons.usu.edu/gradreports/1368>

This Creative Project is brought to you for free and open access by the Graduate Studies at DigitalCommons@USU. It has been accepted for inclusion in All Graduate Plan B and other Reports by an authorized administrator of DigitalCommons@USU. For more information, please contact digitalcommons@usu.edu.



Predictive Distributions via Filtered Historical Simulation for Financial Risk Management

Tyson Clark *Utah State University*

Filtered historical simulation with an underlying GARCH process can be used as a valuable tool in VaR analysis, as it derives risk estimates that are sensitive to the distributional properties of the historical data of the produced predictive density. I examine the applications to risk analysis that filtered historical simulation can provide, as well as an interpretation of the predictive density as a poor man's Bayesian posterior distribution. The predictive density allows us to make associated probabilistic statements regarding the results for VaR analysis, giving greater measurement of risk and the ability to maintain the optimal level of risk per tolerance in portfolios to remain compliant with regulations.

Keywords: Filtered Historical Simulation, Predictive Density, Value at Risk, VaR, GARCH

Introduction

Investors often fail to account for risks included in their portfolios when things are going well. This was a major issue prior to the 2008 global financial crisis, and the crisis reminded many of the importance of evaluating risk. This was most apparent for investment banks that experienced large losses as a result of some risky assets in their portfolios for which they didn't properly account. Because of the importance of banks in our economy, there are regulations in place that require banks to stress-test their portfolios to evaluate the risk exposure.

Since financial data is not i.i.d. normal, it can be difficult to account for the periods of volatility when something such as the 2008 financial crisis occurs. In this paper, I show that a symmetric GARCH model when bootstrapped in a filtered historical simulation allows for this volatility clustering pattern and will provide more accurate VaR metrics for banks that are seeking to remain compliant with the 10-day, 1% significance VaR levels set by regulations established in Basel II.

In this paper I will show a model that banks have begun to use to meet Basel II requirements. Basel II requires banks using internal VaR models to evaluate their risk at the 99% significance level. It also requires banks to measure potential losses for a 10-day risk horizon. Filtered historical simulation is a semi-parametric approach that combines Monte Carlo simulation based on volatility clustering patterns with the empirical non-normal return distributions from historical data. This

overcomes many issues of basic historical simulation by allowing it to find an h-day VaR where h is more than just a few days. As it is a non-parametric model, it will not account for the volatility currently prevailing in the market (Alexander (2009)). Filtered historical simulation becomes this semi-parametric model by combining an estimated GARCH model with the i.i.d. bootstrap. This allows the model to account for value correlations without restricting them over time.

The GARCH process underlying this filtered historical simulation started as a less complex autoregressive conditional heteroskedasticity (ARCH) model invented by Robert Engle (Engle (1995)). This model, as well as the GARCH model I will be using to further discuss filtered historical simulation are explained in detail by Enders (2008). The model was made to account for volatility clustering patterns as is often seen with financial data that will experience market shocks such as bubbles, where a model assuming constant variance is inappropriate. For comparison to the later shown GARCH model, the simplest form of the ARCH model is the ARCH(1) model denoted by:

$$y_t = x_t' \beta + \epsilon_t$$

where:

$$\epsilon_t = u_t \sqrt{\alpha_0 + \alpha_1 \epsilon_{t-1}^2}$$

and $u_t \sim N(0,1)$. This ϵ_t term is what is allowing for the conditional heteroskedasticity.

This ARCH model has some notable weaknesses. It assumes that positive and negative shocks have the same effects on volatility because it depends on the square of the previous shocks. α_t^2 must remain in certain intervals which limits the ability of an ARCH model to allow for excess kurtosis. It also is likely to overpredict the volatility because the model responds slowly to large isolated shocks to the series.

Due to some of these limitations, the generalized autoregressive conditional heteroskedasticity (GARCH) model was introduced by Bollerslev (1986). The GARCH model extends Engle's ARCH model by allowing the conditional variance to be an autoregressive moving average (ARMA) process. This allows for both autoregressive and moving average components in the heteroskedastic

variance. The basic GARCH(1,1) model can be denoted as:

$$\sigma_t^2 = \omega + \alpha e_{t-1}^2 + \beta \sigma_{t-1}^2$$

We can see that if this $\beta \sigma_{t-1}^2$ term in the equation is dropped out, we are left with our ARCH model from before.

There is much literature surrounding GARCH models and there exists an alphabet soup of extensions to the original GARCH model. Because of the leverage effect, there are many models that beat a GARCH(1,1) model ([Hansen and Lunde \(2003\)](#)). The leverage effect means the impact of bad news will have more of a volatility shock than good news. The simple GARCH(1,1) model doesn't take this into account, however it is still remarkable how well the model does at predicting volatility patterns, and is the model I will use throughout this paper. I do include some models of both the S&P 500 and IBM using the GJR-GARCH which includes an additional variable to account for this leverage effect, however in my sample data there is no significant difference between the simple GARCH(1,1) and a GJR-GARCH(1,1,1) model, showing the power of the GARCH(1,1) model.

It is easy to see here in the introduction to filtered historical simulation that all of these models and computational techniques are all from the classical frequentist perspective, but there is a very Bayesian feel to the overall method, especially concerning the results that are produced. Through the statistical bootstrap of the GARCH process, we produce a predictive distribution that shows a distribution of the VaR for an h-day horizon. As these are all classical techniques, filtered historical simulation is rarely, if ever, looked at from a Bayesian perspective.

VaR Overview and Models

Since the 1990s, almost all financial institutions have used some sort of VaR measure. [Alexander \(2009\)](#) outlines some important advantages of a VaR measure, including: VaR calculates an amount that could be lost with some predetermined probability, VaR measures the risk factors as well as their sensitivities, its scalability, and others. Risk management is important for banks because they need to know what their level of risk tolerance is to determine if certain risks will be held or hedged away. Before this can be determined, banks need a way to measure this risk.

As mentioned previously, VaR is a very important metric for banks to consider. As such, many different methods of calculating VaR have appeared. I will highlight three methods mentioned by [Alexander \(2009\)](#) in her book 'Market Risk Analysis: Volume IV': normal linear VaR, historical simulation, and Monte Carlo simulation. Again, as stated previously, filtered historical simulation uses aspects of historical simulation and Monte Carlo simulation that allow it to be a semi-parametric model, making it the optimal model for banks to calculate their h-day VaR models.

I will only briefly mention normal linear VaR, as it is only appropriate when a portfolio P&L is a linear function of its risk factors or asset returns. As such, it is not ideal for almost any financial asset.

Historical simulation is the method that is the easiest to conceptualize. Most, if not all, investors have looked at previous return data when looking at their future investment decisions. In fact, I would assume that most novice investors believe this is the best method of estimating future returns. Historical simulation's biggest issue is that it assumes that all possible future variations have occurred in the past. Looking back at each individual financial crisis, we can see that these models will vastly under-predict the possibility of some new market failure because of some issue that has not previously occurred.

The third method for modeling VaR is Monte Carlo simulation. Monte Carlo is useful since it can be applied to non-linear portfolios. Monte Carlo doesn't look back at historical data; however, and may not be the most useful tool for looking at risk estimates based off of historical volatility clustering patterns.

Filtered Historical Simulation

Filtered historical simulation was first introduced by [Barone-Adesi, Giannopoulos and Vosper \(1999\)](#). It is a great method for determining VaR for portfolios, as it uses both nonlinear econometric models and historical returns to build the predictive density that the portfolio could take for multiple days looking forward. Rather than most models, which suffer the problem of under-estimating the risk of extreme outcomes, filtered historical simulation risk estimated are derived directly from the tails of the distribution.

In filtered historical simulation, price series are not forced to conform to a probability distribu-

tion, but the data are allowed to speak for themselves. Historical data is still considered. This is an important aspect to filtered historical simulation since a basic Monte Carlo simulation only draws from a theoretical distribution, which smooths the empirical distribution and may underestimate the risk of catastrophe. It generates thousands of scenarios for the mean and variance of each risk factor in a multi-period horizon. This allows filtered historical simulation to provide a more accurate depiction of the tail structure in future prices. These features make filtered historical simulation very easy to implement in risk analysis and stress-testing.

Filtered historical simulation works by using a parametric model of return volatility, such as a GARCH model, to simulate log returns over some predefined risk horizon. Using the estimated GARCH model

$$\hat{\sigma}_{t+1}^2 = \hat{\omega} + \hat{\alpha}r_t^2 + \hat{\beta}\hat{\sigma}_t^2$$

, filtered historical simulation will assume that GARCH draws from the standardized empirical distribution, so therefore the standardized innovations are $\epsilon_t = r_t/\hat{\sigma}_t$. r_t represents the historical daily log return and $\hat{\sigma}_t^2$ represents the GARCH daily standard deviation.

The simulation process begins by setting initial conditions $\tilde{\sigma}_0$ and \tilde{r}_0 . \tilde{r}_0 will be set equal to the log return from the previous day, but there are at least two different options for setting $\tilde{\sigma}_0$.

The first of the two options is setting $\tilde{\sigma}_0$ equal to the last estimated daily conditional volatility from the GARCH model. The second option is to set $\tilde{\sigma}_0$ equal to the long run unconditional volatility:

$$\hat{\omega}/(1 - \hat{\alpha} - \hat{\beta})$$

. In my model, I will use the latter. This will give a return on the forecasted day 1 of $r_1 = \epsilon_1\hat{\sigma}_1$. The value ϵ_t will be drawn independently via the i.i.d. bootstrap. I can now simply iterate through:

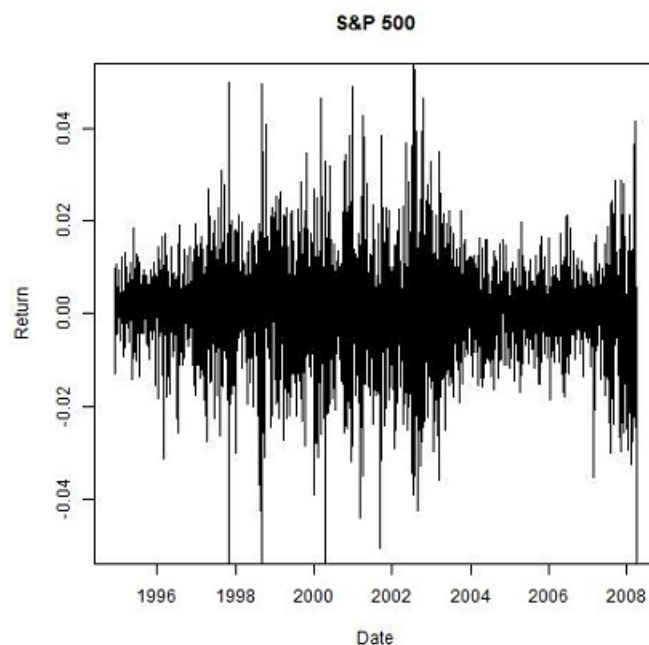
$$\hat{\sigma}_{t+1}^2 = \hat{\omega} + \hat{\alpha}r_t^2 + \hat{\beta}\hat{\sigma}_t^2$$

. Since I have used log returns, the h-day return is simply $\sum_t^h r_t$.

Repeating this through simulation will produce a predictive density that can be used for VaR estimation. The h-day return will be displayed in the distribution to show the probability distribution of the h-day returns. To demonstrate this, I have coded filtered historical simulation models

with corresponding bootstraps of the underlying GARCH processes for both the S&P 500 from December 1, 1994 to March 31, 2008 and IBM from the beginning of 1999 to the end of 2003. The S&P data was chosen to show the volatility clustering patterns even for the overall market, and the extreme shocks shown in 2008 that occurred because of the financial crisis. IBM is a relatively 'safe' stock that has not had any major issues in the past, but will work well for my purposes in showing the effectiveness of filtered historical simulation.

When looking at the volatility of the S&P 500 from January 2, 1995 to March 31, 2008 below, it is easy to notice the volatility clustering patterns that make GARCH an excellent candidate for determining the volatility in the underlying. With just a simple volatility calculation, these patterns will not be taken into account, and the VaR will be underestimated, potentially leading to greater than expected losses in a portfolio.



To calculate the VaR through filtered historical simulation, first I need to estimate a GARCH model to account for volatility clustering patterns as seen above. I used the optimizer 'rugarch' from R to optimize the GARCH parameters $\hat{\omega}$, $\hat{\alpha}$, and $\hat{\beta}$. These are the estimated parameters that were used in the i.i.d. bootstrap process for filtered historical simulation. The full S&P 500 and IBM code can be found in the linked Github Repository ([Clark, Tyson \(2019\)](#)). As is shown by [McDonald \(2009\)](#), there are jumps in the return data for IBM surrounding earnings announcements

that the GARCH model has difficulty explaining, so those 5 particular dates have been omitted from the optimization process.

After the optimizer has found the values of $\hat{\omega}$, $\hat{\alpha}$, and $\hat{\beta}$, they can be used in the i.i.d. bootstrap process of filtered historical simulation. I will first find the scaled volatility adjusted VaR using the GARCH process to compare my results in filtered historical simulation.

Quantile	Unadjusted	GARCH
0.10%	4.84%	7.47%
1%	2.83%	4.19%
5%	1.78%	2.69%
10%	1.27%	2.10%

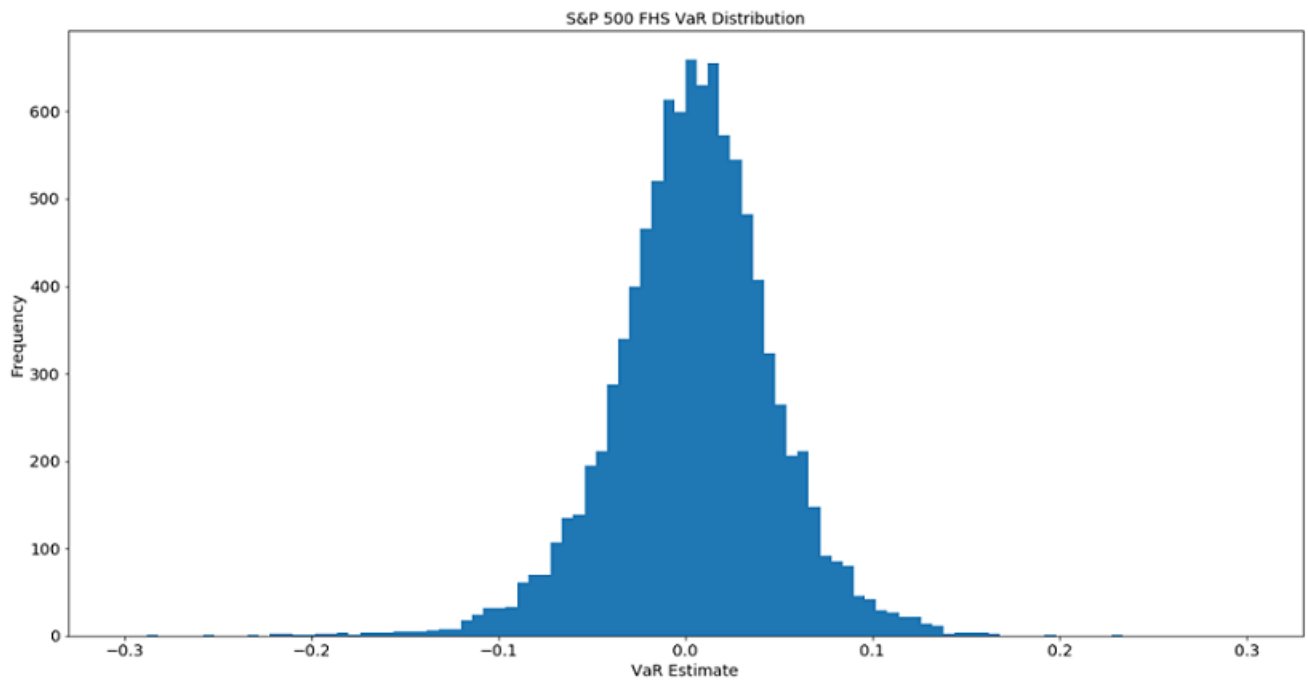
The table above shows the 1-day VaR for the S&P 500 on March 31, 2008, the last day of my sample. We can see here that unadjusted volatility will severely underestimate a VaR. Using GARCH to correctly adjust for volatility clustering shows a much higher VaR in each respective quantile. This is a very important finding, especially for banks that have a limit on VaR they are allowed to have in their respective portfolios. A simple volatility model will allow banks to take on substantial risk, which can be catastrophic in the event of a tail event, as we saw within this sample in 2008.

This simple GARCH volatility estimate can now be bootstrapped to allow for an h-day VaR to be calculated. This filtered historical simulation process will allow the GARCH volatility to revert to the long-run mean.

Quantile	GARCH_scaled	FHS
0.10%	23.61%	19.90%
1%	13.24%	10.96%
5%	8.52%	6.68%
10%	6.66%	4.72%

As is shown in this table, the adjusted GARCH model overestimates risk when compared

to filtered historical simulation. This is because the square root scaling of a daily VaR model estimated through GARCH doesn't allow for mean reversion, and will subsequently overestimate the VaR when taken on a day with a higher than usual variance. This is why it is important to use the statistical bootstrap, and why filtered historical simulation is better than a simple square root adjustment to a single day GARCH estimation. This will hold in almost all of the 10,000 simulations carried out. We can see this distribution in the produced predictive density. My simulated predictive density is shown below.



I also want to point out the comparison of the FHS percentile returns with the VaR estimates that are unadjusted for volatility clustering. The table is shown below.

Quantile	Unadjusted	FHS
0.10%	4.84%	19.90%
1%	2.83%	10.96%
5%	1.78%	6.68%
10%	1.27%	4.72%

When returns are unadjusted for volatility clustering patterns, risk estimates will severely underperform when compared to a filtered historical simulation VaR model. Banks should use filtered historical simulation to account for these risks, and when they use these models to remain compliant with Basel II standards, they will be much less likely to become insolvent as we saw with many banks during the 2008 financial crisis.

This method of filtered historical simulation has been an incredible innovation in risk estimation. In fact, it has been extended to apply to more complex VaR models for portfolios with changing weights that may even include derivative securities ([Barone-Adesi, Giannopoulos and Vosper \(1999\)](#)). The same method I have applied is used in this paper, by simulating asset volatilities to depend on the most recent portfolio returns. The combination of GARCH modeling and historical returns creates this same semi-parametric model to find the VaR, and is shown to work quite remarkably.

Predictive Density

This produced predictive density helps us see the distribution of returns through a bootstrap of 10,000 simulations. This distribution can be quite useful when a bank or investment manager is looking at the VaR for his or her respective portfolio. The density is computed using frequentist methods, but can be looked at as a poor-man's Bayesian posterior distribution when there is an uninformed prior, as is most likely the case with a VaR model ([Murphy \(2012\)](#)). In fact, the way filtered historical simulation is modeled, it is very similar to a Bayesian model. Without getting too involved in the Bayesian perspective, Bayes' rule can be defined as:

$$P(\theta|x) \propto P(x|\theta)P(\theta)$$

, where $P(\theta)$ is the prior applied to the Bayesian model. It can be seen that when this prior is uninformed, it will drop out, leaving just $P(x|\theta)$, which will be informed by the model just as filtered historical simulation.

It is shown here that the actual VaR model is being computed using frequentist methods, but this produced density is the frequentist approach to a Bayesian posterior. I would like to quote Terry Speed about how 'We don't need to be Bayesian; we just need to be approximately so. We

don't need theory to tell us our method works; we just need to simulate and see' (Speed (2011)). This Bayesian interpretation can be important. In 'Computer Age Statistical Inference', Efron and Hastie provide some comparisons between frequentist methods and Bayesian methods (Efron and Hastie (2016)). An important note they make is that a Bayesian places a large bet on his or her prior being correct, or at least not harmful, but a frequentist will take a more defensive approach, that by statistical methods he or she can produce a useful distribution, regardless of the true μ .

In the perspective of filtered historical simulation for VaR, this is an important point since there can be no prior belief regarding the distribution of returns if one wants to properly account for risk, especially in the tails. Though calculated using frequentist method, the distribution shown above is very close to a Bayesian posterior distribution.

When a risk manager or trader is thinking about risk, they can't think about these frequentist models. Though the models may do a good job of summarizing data, the true uncertainty must be thought of from a subjective sense. Filtered historical simulation produces a predictive density that allows traders and risk managers to do just that. Though filtered historical simulation is a frequentist method, rarely, if ever, is the produced distribution thought of from a Bayesian subjective sense, though that is exactly what the model is producing. This way of thinking about uncertainty incorporates the theory that prices reflect information, or tacit knowledge, rather than just being something that is calculated through a model (Hayek (1945)). This way of thinking will help risk managers and traders more properly account for risk in their models, and is a great way to look at filtered historical simulation.

Conclusion

It can be shown in the VaR estimates generated from filtered historical simulation compared to the estimates based off the unadjusted volatility that the probability of extreme loss is severely underestimated. This has been an issue with banks in the past, as many have modeled VaR using only historical returns, which assumes that all future shocks to their portfolios have occurred in the past. Just in the last couple of decades we can see that this is not a correct assumption, as new market failures send asset prices sharply downward, such as the financial crisis of 2008. Though outside of the scope of my current research, if bank managers had applied this semi-parametric

filtered historical simulation process to portfolios of mortgage-backed securities or collateralized debt obligations that were the recipient of extreme losses, there is a decent probability that the exposure to these assets would have been reduced and the severity of the crisis would have been reduced.

The method of VaR estimates produced in the paper can be generalized to many different assets or portfolios of assets to produce valuable distributions for risk managers of banks and investment portfolios alike. It has been shown that these methods can even be used to model VaR estimates of derivative securities ([Barone-Adesi, Giannopoulos and Vosper \(1999\)](#)). To maintain Basel II requirements, a 99% significance level must be used, but for agencies seeking AA ratings or above, they must use significance levels above 99.7%, which as seen in the tables above will be vastly underestimated if one doesn't account for volatility clustering patterns that assets clearly possess.

The produced predictive density allows risk managers and traders alike to think about the uncertainty in their portfolios from a subjective sense, rather than simply the result of a model. It allows them to take into account all of the information incorporated into markets and how it will be reflected in their VaR analysis. Filtered historical simulation is very valuable in this way, and has a wide variety of useful applications.

Bibliography

- Alexander, Carol. 2009. *Market Risk Analysis, Value at Risk Models*. Vol. 4 John Wiley & Sons.
- Barone-Adesi, Giovanni, Kostas Giannopoulos and Les Vosper. 1999. "VaR without correlations for portfolios of derivative securities." *Journal of Futures Markets* 19(5):583–602.
- Bollerslev, Tim. 1986. "Generalized autoregressive conditional heteroskedasticity." *Journal of econometrics* 31(3):307–327.
- Clark, Tyson. 2019. "Predictive Densities via Filtered Historical Simulation for Financial Risk Management Github Repository." <https://github.com/Totyso/Predictive-Densities-via-Filtered-Historical-Simulation-for-Financial-Risk-Management.git>.
- Efron, Bradley and Trevor Hastie. 2016. *Computer age statistical inference*. Vol. 5 Cambridge University Press.
- Enders, Walter. 2008. *Applied econometric time series*. John Wiley & Sons.
- Engle, Robert. 1995. *ARCH: selected readings*. Oxford University Press.
- Hansen, P Reinhard and Asger Lunde. 2003. Does anything beat a GARCH (1, 1)? A comparison based on test for superior predictive ability. In *2003 IEEE International Conference on Computational Intelligence for Financial Engineering, 2003. Proceedings*. IEEE pp. 301–307.
- Hayek, F.A. 1945. "The Use of Knowledge in Society." *The American Economic Review* 35(4):519–530.
- McDonald, Robert L. 2009. "Fundamentals of derivatives markets."
- Murphy, Kevin P. 2012. *Machine learning: a probabilistic perspective*. MIT press.
- Speed, Terry. 2011. "IMS Elections: time to vote."