Another Stochastic Volatility Model with Markov Switching

Peter Trubey Mahroo Ehsani

Abstract

In this paper, we briefly explain the theory behind a relationship between stock returns and trading volume. We then explain the stochastic volatility model and Markov switching addition, and motivate why we would use these in financial research. In OpenBUGS we build the SV model using AMD and Microsoft stock, and compare results. We then attempt the MSSV model, and note shortcomings of our implementation. We postulate why the MSSV model failed to work. Finally, we suggest steps we can take fix the MSSV model, to better understand this relationship.

1 Introduction

In this paper we attempt to model a relationship between stock returns, and traded share volume. This relationship is well studied in finance literature, and there are several reasons why this relation is important. In his paper[3], Karpoff cites three. First, it provides insight into the structure of financial markets. There are models that predict price-volume relationsships that depend on the rate of information flow to the market, and how the information is disseminated through the market. Second, the price-volume relation is important for event studies that use a combination of price and volume data from which to draw inferences. The construction of tests and validity of the inferences drawn depends on the joint distribution of price changes and volume. Third, price-volume relations has significant implications for the research into futures markets. Price variability affects the volume of trades in futures contracts. The time to delivery of a futures contract affects the volume of trading, and through this affect, possibly also the variablity of price. This circular path may have some effect on the residual nature of volatility in a stock.

1.1 Mixture of Distributions Hypothesis

The mixture of distributions hypothesis states that stock returns and trading volumes both depend on the same latent information flow variable[3]. Both models we present today fall under this hypothesis, theorizing that the flow of information into the market can be modelled as the volatility of the market. This expresses itself in both stock price changes, and changes in share volume.

In the stochastic volatility model, the variance of the stochastic process is itself randomly distributed. This is in contrast to other homoscedastic models, where the variance of the model is a constant term. The use of SV allows us to model the *fat tails* of the distribution of returns more effectively.

The Markov switching model, also known as the regime switching model, is another oft-used method of representing the changing volatility. In this model, we include an indicator variable indicating which regime the model is in at a given point, and allow the model to move between regimes using probabilities encased in a switching matrix. The regime governs the level of volatility in the market at a given time.

1.2 Trade Modeling

Establishing a relationship with trading behavior requires us to make some assumptions about the traders. In these models we assume the existence of a group called *noise traders*, or traders who make their trades in absence of information. Hereafter we will refer to them as the uninformed trading component. We also assume the existence of *informed traders*, who make trades based on the flow of information. In these models, we assume the informed trading component activity is proportional to the flow of information in the market, which we previously referred to as volatility.

The model resulting from combination of MDH, and this trade modeling relationship is called Modified Mixture Model(MMM). It is established using a time series to model the flow of information into the market.[4] The remainder of this paper is organized as follows: Section 2 is data preparation of the data set. Section 3 presents two models of stochastic vlatility and Markov-switching stochastic volatilty. Section 4 shows an empirical application on the return and trading volume for the data set. Finally, section 5 concludes.

2 Data Preparation

In this research we attempt to implement a modified Stochastic Volatility (SV) model of Mahieu and Bauer[4]. We further attempt to implement the Markov Switching SV model specified in So, et al.[7]. The data set we use comes from The Center for Research In Security Prices[2]. The raw data set includes daily stock prices and corressponding share volume from August 26, 2008 to Dec 31, 2012 for two companies of Advanced Micro Devices Inc.(AMD) and Microsoft Corp.(MSFT)

2.1 Daily Returns

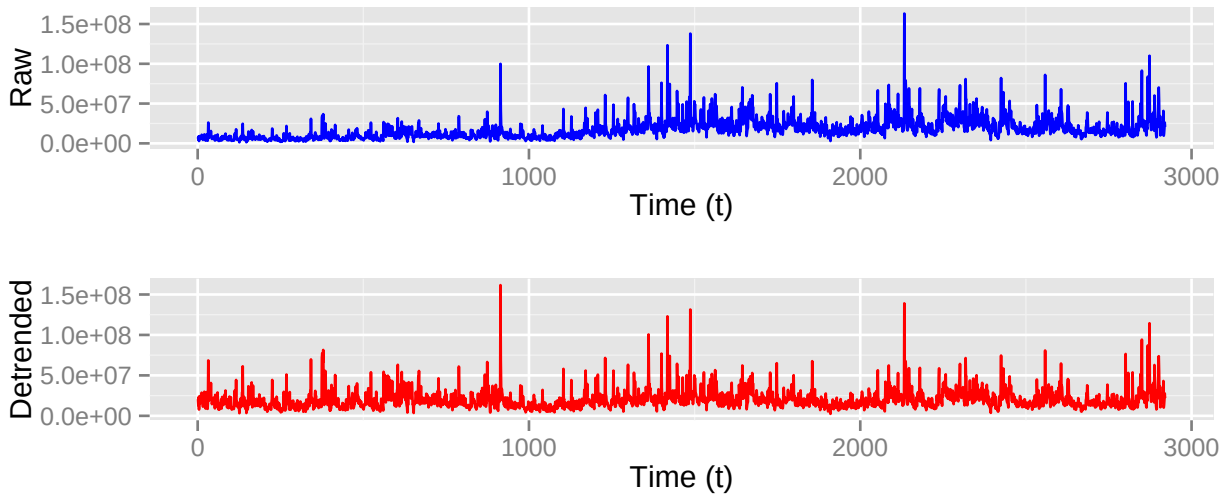
In the previous section, we mention the daily return of a stock. This value is equivalent to the percent return of a hypothetical investment made in the stock on the previous day, calculated from closing price to closing price. In this equation, y_t is daily return on day t, P_t is the closing share price of the stock on day t.

$$y_t = 100(\frac{P_t}{P_{t-1}})$$

2.2 Daily Share Volume

We call v_t the share volume traded on day t. The preparation for this term is a bit more involved. These models call for a detrended share volume, where long term trends are removed from the data and we are only concerned with day to day shifts. The literature was mute on how to go about this detrending, so we were forced to make a guess. What seems logical is to scale a local mean to the overall mean, in order to remove long-duration trends. To accomplish that, we

Figure 1: Raw and Detrended Share Volume Series' for AMD stock



calculate a smoothed trend line through the use of localized weighted scatterplot smoothing, or *lowess*. We then scale the observations of volume based on the ratio of the overall mean and the lowess value. See figure 1 to compare the raw and detrended series.

$$\pi(t) = lowess(V_t)$$
$$v_t = c(\frac{V_t(\frac{1}{n}\sum_{i=1}^{n} V_i)}{\pi(t)})$$

Where V_t is the share volume at time t, and v_t is the de-trended share volume at time t. We multiply by a scaling constant c to bring the volumes down to a level that can be efficiently modeled using a Poisson distribution. This scalar is dependent on the average share volume of the stock, and is calculated as a byproduct of scaling the detrended share volume to a target mean. We then round the detrended scaled volume measurements to be integers, so that they can be modeled by a Poisson distribution. These time series, y_t being the daily returns and v_t being the scaled share volume for $t = 1, \ldots, n$ constitute the inputs to the models we will be using.

3 Models

3.1 Stochastic Volatility - SV Model

The Stochastic Volatility Model for stock returns establishes a relationship between stock returns, and a latent variable *volatility*, which we model as e^{h_t}. We also assume that trading volume is related to the same latent variable. The code used to implement the time series is adapted from Meyer and Yu[6]. The model can be expressed as:

$$y_t|h_t \sim N(0, e^{-h_t})$$
$$v_t|h_t \sim Pois(c(m_0 + m_1 e^{h_t}))$$
$$h_t = \alpha + \phi(h_{t-1}) + \eta_t \qquad \eta \sim N(0, \tau_\eta^2)$$

Described here is the empirical volatility model of Mahieu[4], in terms of precision. In this paper, we make a slight modification to this model, which as a consequence of a non-centering autoregressive process on the log-volatility, assumes a generally increasing share volatility with respect to time. We prefer to not use such an assumption, so we propose using a centering autoregressive process, which can be expressed as:

$$h_t = \alpha + \phi(h_{t-1} - \alpha) + \eta_t$$

holding all other portions of the model the same. This has the effect of assuming stationarity for the volatility series. In this model, we are assuming y_t follows a normal distribution centered at 0, with a variance at time t equal to the volatility at time t. We further assume that the share volume traded follows a sum of two poisson distributions, representing informed and non-informed trading respectively. The non-infomed trading will follow a poisson distribution with $\lambda = m_0$. The informed trading component of share volume follows a poisson distribution using constant m_1, proportional to the latent volatility variable e^{h_t}. Note that in this model formulation we are multiplying by the scalar constant c. In the data preparation, we had already absorbed that c into the share volume, so we can regard it as already being multiplied into $m0$ and $m1$.

We consider volatility to be an indicator of information flow. As explained previously, theory tells us that both stock price movement, and share volume change as a result of information flow. Through the informed trading predictor m_1, this model attempts to see how much trading is occuring as a result of information flow, and how much occurs in absence of it.

3.2 Markov Switching Stochastic Volatility - MSSV Model

The Markov Switching Model allows us to change regimes between periods of high volatility and low volatility. In the Markov switching model, we introduce another parameter S_t, which gives the state at time t, and a transition matrix ξ which gives the probabilities of switching regimes dependent on the previous state. The BUGS code for the Markov Switching component is adapted from Martinez-Beneito, et al.[5]. The model changes the formulation of h_t as follows:

$$\begin{aligned} h_t &= \alpha_t + \phi(h_{t-1} - \alpha_{t-1}) + \eta_t \\ \alpha_t &= \alpha_0 + \alpha_1 S_t \end{aligned}$$

In this model, we assume α_1 to represent the additional volatility corresponding to periods of high volatility, whereas α_0 represents the baseline log-volatility inherent to the stock. The state variable S_t is a single-order Markov Switching Variable, to allow for α_1 to only appear during periods of high volatility. It moves using a transition matrix ξ which gives the probability for moving from any state to any other state. This model is a specific case of the model outlined in So, et al.[7] where we use a centered time series for the log-volatility, and we use a binary transition matrix to only consider 2 possible states.

$$\xi = \begin{bmatrix} \xi_{0,0} & \xi_{0,1} \\ \xi_{1,0} & \xi_{1,1} \end{bmatrix} \qquad S_t = \begin{cases} 0 & \text{w.p. } \xi_{(S_{t-1},0)} \\ 1 & \text{w.p. } \xi_{(S_{t-1},1)} \end{cases}$$

In this model, we assume the state parameter $S_t \in \{0, 1\}$ with $S_t = 1$ corresponding to a period of high volatility, and $S_t = 0$ a period of low volatility. For that reason, we are restricting α_1 to

be positive. ϕ refers to the residual effects of past volatility on the present volatility, and we also assume that to be positive.

4 Analysis

Both models were ran for both stocks. After initial results, a burn-in period of $20,000$ iterations was selected. Models were then ran for an additional $60,000$ iterations to yield the resulting distributions.

4.1 Stochastic Volatility Model

We start with the stochastic voltility model, In this model, we track $\theta = (\alpha, \phi, m_0, m_1, \tau_\eta)$.

We know from prior knowledge that AMD is a low share price, highly variable stock. It should stand to reason that it would have a relatively high volatility compared to a much more entrenched stock, Microsoft. In looking at the model summary in table 1, the effect of volatility is what we see. AMD has a comparitively high baseline volatility parameter $\alpha \approx 2.37$, and a comparitively low precision for η at $\tau_\eta \approx 2.6$. We also take note that the residual effect of the previous day, ϕ is high near 0.8, but not as high as the microsoft. When we observe the time series (figure 2) share volume and daily returns for AMD stock over time, and compare to the log volatility plot, we see large rises in share volume generally correspond with large rises in share volume. This result is logical.

On the other hand, Microsoft is a solid and comparatively stable stock. In table 2, we see the baseline volatility on Microsoft's stock using the SV Model is much lower, with a posterior mode of around .79. We can see the posterior distribution of this parameter in figure 4. The Precision component τ_η is much stronger for Microsoft, meaning the variance of the log-volatility time series is much lower than it is on AMD's stock. This coincides with our previous inference that Microsoft is a much less volatile stock. The non-informed trading components for both stocks are similar, however the informed trading components change in size significantly.

We see the effect of this in the log volatility time series on the Microsoft Stock. In figure 3, we see the log-volatility is overall much lower than it was with the AMD stock, along with changes in daily returns and share volume.

Table 1: SV Model on AMD Stock

	Mean	SD	MCσ	$Z_{.025}$	Median	$Z_{.975}$
α	2.365	0.1058	0.001646	2.154	2.366	2.57
m_0	3.276	0.1475	0.003315	2.974	3.28	3.554
m_1	0.1093	0.01197	0.0002647	0.08744	0.1089	0.1341
ϕ	0.789	0.03316	0.0007996	0.7198	0.7906	0.8496
τ_η	2.593	0.448	0.01141	1.857	2.543	3.605

Table 2: SV Model on Microsoft Stock

	Mean	SD	MCσ	$Z_{.025}$	Median	$Z_{.975}$
α	0.7919	0.1499	0.002459	0.4928	0.7928	1.089
m_0	3.245	0.1024	0.002182	3.044	3.244	3.446
m_1	0.2208	0.03174	6.98E-04	0.1613	0.2197	0.2851
ϕ	0.9056	0.02971	8.79E-04	0.8397	0.9084	0.9555
τ_η	6.328	2.132	0.07033	3.294	5.96	11.91

4.2 MSSV Model

In the Markov Switching Stochastic Volatility Model, we additionally track the Additional baseline volatility parameter, along with the diagonal components of the switching matrix ξ. Thus, $\theta = (\alpha_0, \alpha_1, m_0, m_1, \phi, \xi_{0,0}, \xi_{1,1}, \tau_\eta)$.

In the model summary for AMD's stock in table 3 we see similar estimates for m_0, m_1, τ_η as we saw in the SV model indicating our estimates for the informed and non-informed ratios of the marketplace are consistent. However, we see α_0 is significantly larger than α was in the SV model. α_1 is also positive, meaning for periods where $S_t = 1$, α_t will be even larger. This is contrary to the goals of the MSSV model, as we were hoping to isolate the high volatility portions of the series into the $\alpha_1 S_t$ component. Adding the Markov Switching effect instead increased the baseline volatility. Interestingly, the residual volatility from the previous iteration did not feature so strongly in this model, with $\phi \approx .58$.

There are some reasons we might have seen this. We may have specified too strong priors on some components yielding this result. On the other hand, our method of detrending and scale used may not be consistent with the authors of the model. We need to do additional work to make the model work consistently and with results we can interpret. Unfortunately, The Markov Chains for the MSSV Model on Microsoft Stock did not converge, and therefore we cannot make inferences about it. There may be reasons as to why the chains converged for AMD and not for Microsoft. We are unable to explain that at this time.

Table 3: MSSV Model on AMD Stock

	Mean	SD	MCσ	$Z_{.025}$	Median	$Z_{.975}$
α_0	2.837	0.09796	0.002393	2.643	2.837	3.032
α_1	1.12	0.3276	0.009397	0.603	1.076	1.898
m_0	3.288	0.1479	0.003931	2.984	3.292	3.561
m_1	0.1093	0.01205	2.95E-04	0.08767	0.1086	0.1348
ϕ	0.5778	0.0605	0.001733	0.4509	0.5799	0.6922
$\xi_{0,0}$	0.9837	0.009524	2.56E-04	0.9611	0.9856	0.995
$\xi_{1,1}$	0.9641	0.02414	7.20E-04	0.911	0.9689	0.9892
τ_η	2.484	0.4357	0.01126	1.768	2.439	3.466

5 Conclusion

The stochastic volatility model we adapted from Mahieu and Bauer[4] proved to be more stable for both models, and yield more useful results. These results were consistent with our reasoning, showing the more stable stock to be more stable. This tautological reasoning at least tells us the model is consistent.

On the other hand, the Markov Switching Stochastic Volatility model did not converge for the more stable stock. For the less stable stock, it increased the apparent baseline volatility, rather than isolating the periods of high volatility into the high volatility parameter. Indeed, in looking at the plot of S_t relative to share volume, and daily returns, we see S_t often tends to be going in the opposite direction from how we would expect it to move. We have no explanation for this behavior at this time.

Still, with the SV model there are various inferences we can draw. We can say that the structure of stock returns and share volume are definitely different for the two stocks sampled. The companies we chose both belong to the tech sector, so that may account for the similar trading structure regarding informed and uninformed trading. The difference in ϕ is of particular interest, because according to it the volatility in Microsoft's stock is longer lived than that of AMD. There may be reasons as to why this is that could be of interest to us.

6 Future Work

As the MSSV model did not act in a way in accordance with our goals of the model, our course of action should be to amend that. There are several steps to take in order to accomplish this. The first of which would be to review our choice of priors, and potentially set stronger priors to reflect our beliefs of how the model *should* act. Second, Abanto[1] implemented a multi-move sampler for h_t in order to limit autocorrelation of draws within the chain. Autocorrelation was a large hurdle in fitting this model, as all the parameters in our chains were seeing significant autocorrellation well beyond a lag of 100 iterations. Even use of the "over-fit" function in OpenBUGS had little noticeable effect. Implementing the multimove sampler should result in reducing these autocorrelations, and speed convergence.

References

[1] Carlos A. Abanto-Valle, Helio S. Migon, and Hedibert F. Lopes. Bayesian modeling of financial returns: A relationship between volatility and trading volume. *Applied Stochastic Models in Business and Industry*, 26(2):172–193, 2010.

[2] Center for Research in Security Prices CRSP. Graduate school of business, university of chicago.

[3] Jonathan M. Karpoff. The relation between price changes and trading volume: A survey. *The Journal of Financial and Quantitative Analysis*, 22(1):pp. 109–126, 1987.

[4] Ronald Mahieu and Rob Bauer. A bayesian analysis of stock return volatility and trading volume. *Applied Financial Economics*, 8(6):671–687, 1998.

[5] Miguel A. Martnez-Beneito, David Conesa, Antonio Lopez-Qulez, and Aurora Lopez-Maside. Bayesian markov switching models for the early detection of influenza epidemics. *Statistics in Medicine*, 27(22):4455–4468, 2008.

[6] Renate Meyer and Jun Yu. Bugs for a bayesian analysis of stochastic volatility models, 2000.

[7] Mike K P So, K Lam, and W K Li. A stochastic volatility model with markov switching. *Journal of Business & Economic Statistics*, 16(2):244–53, 1998.

Figure 2: Log Volatility h_t from the SV model on AMD Stock

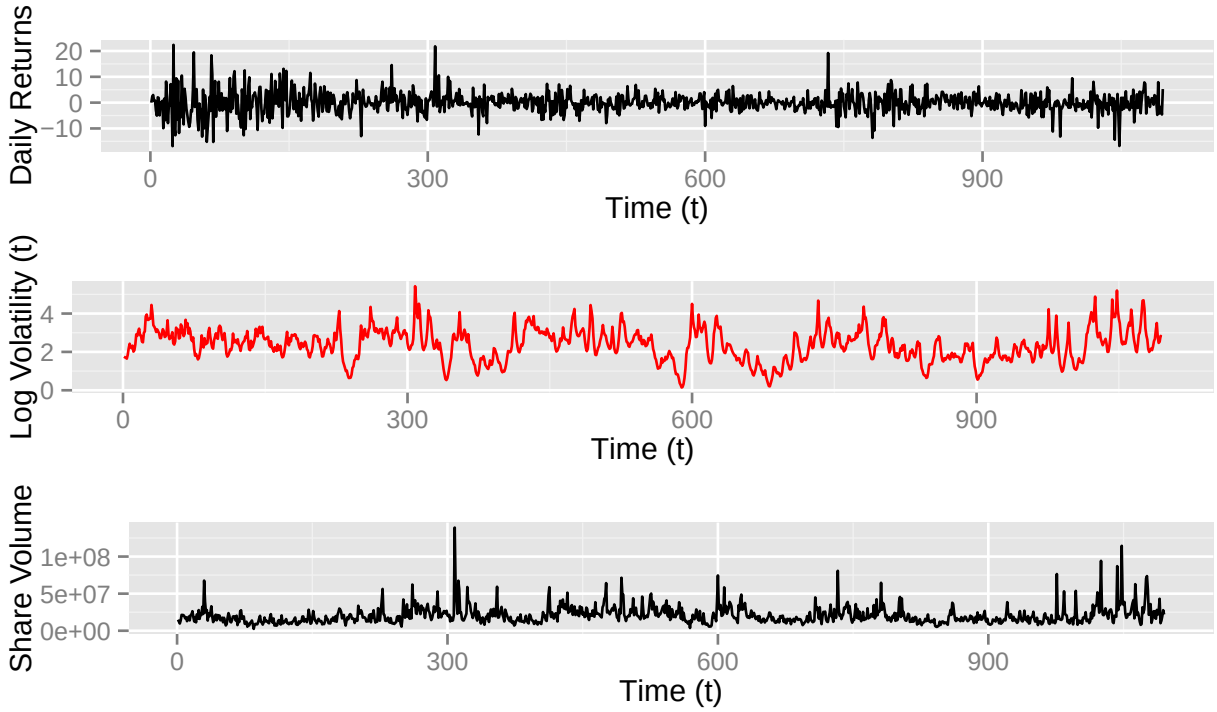


Figure 3: Log Volatility of SV Model on Microsoft Stock

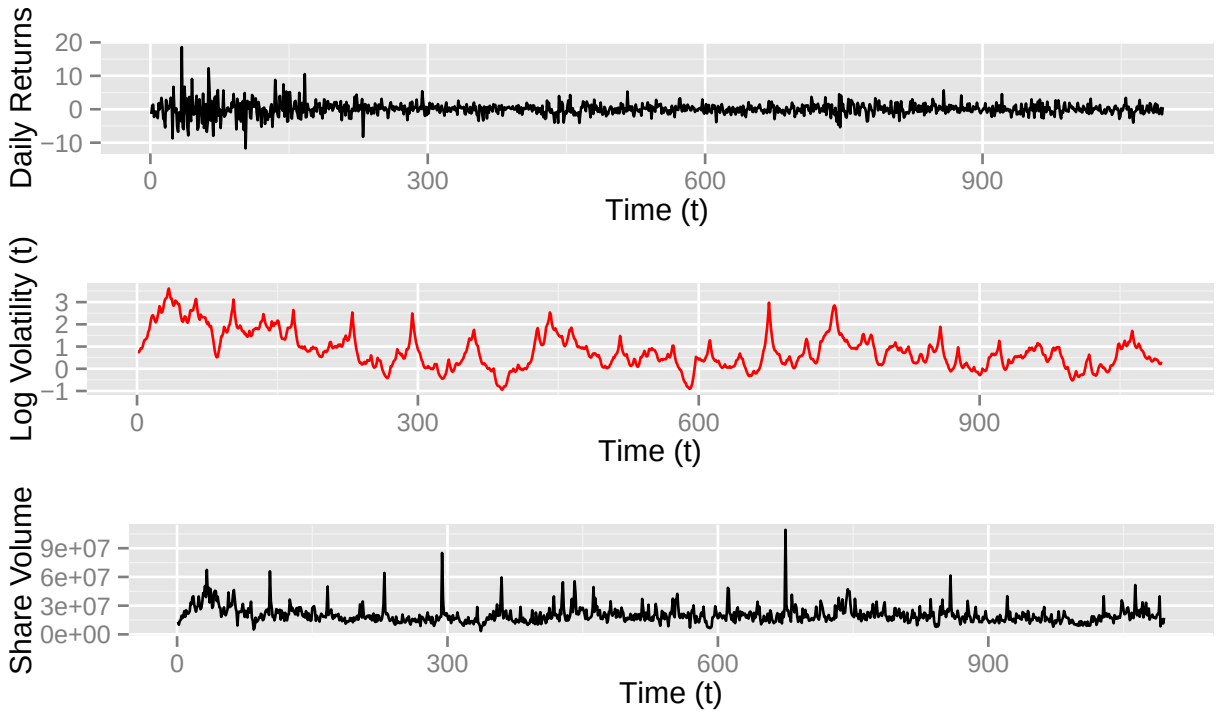


Figure 4: Posterior Distributions for SV Model: α

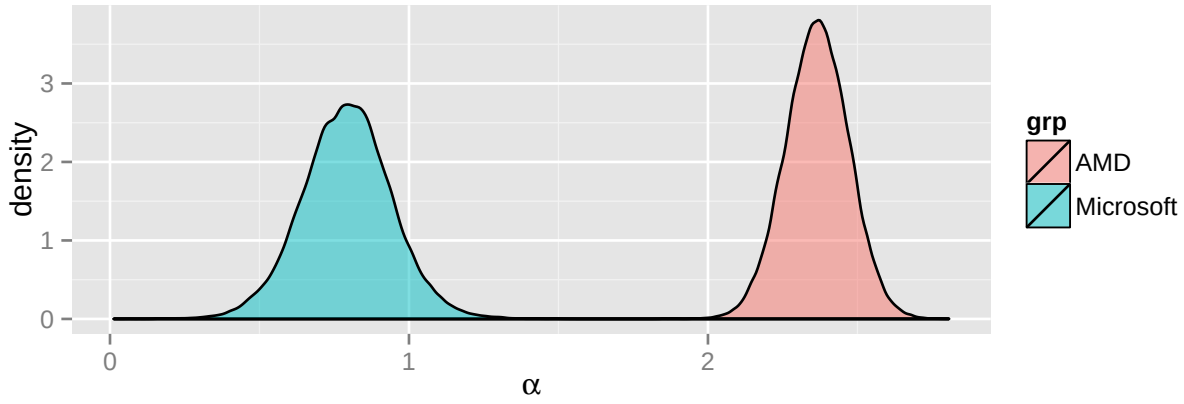


Figure 5: Posterior Distribution for SV Model: ϕ

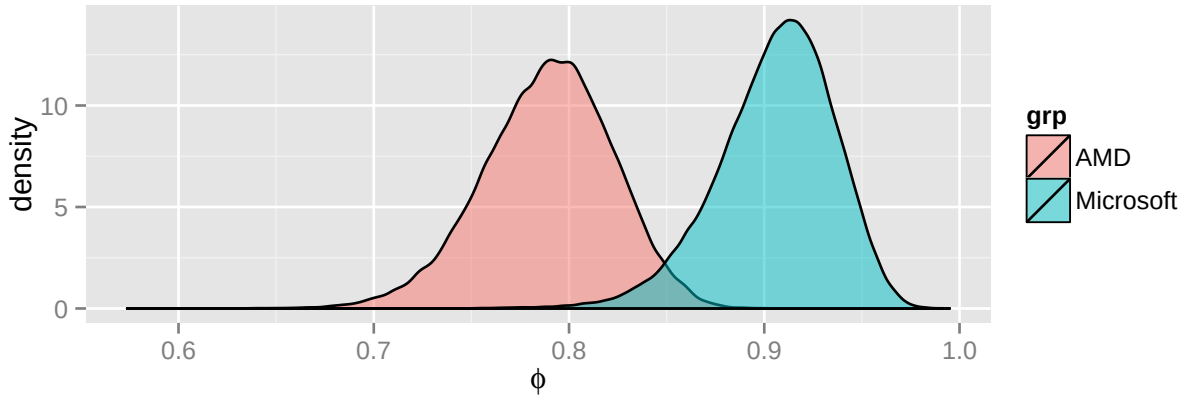


Figure 6: Posterior Distributions for SV Model: τ_η

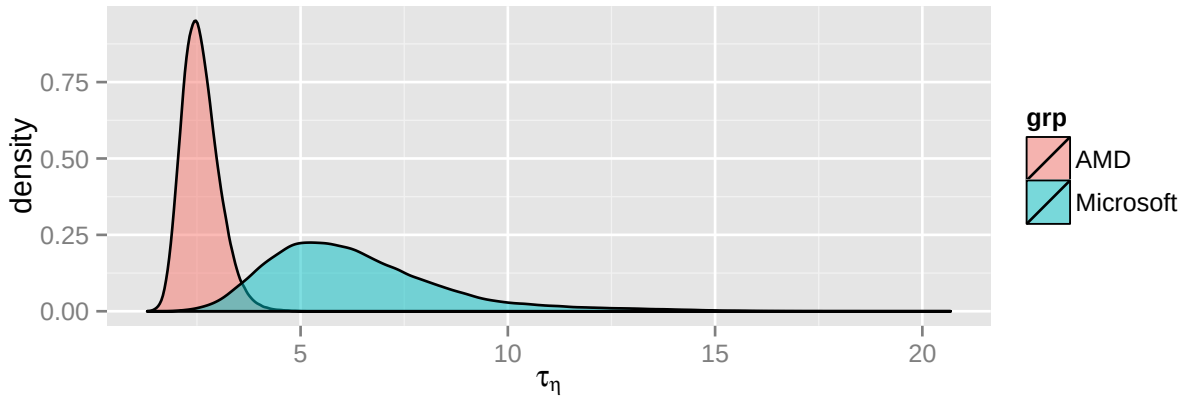


Figure 7: Posterior Distributions for SV Model: m_0, m_1

