

Variance Ratio Results

August 6, 2016

To compute our variance ratios, we begin with daily price data and compute (1) daily returns; (2) five-day returns (for non-overlapping five day intervals); (3) ten-day returns (for non-overlapping ten day intervals); (4) twenty-day returns (for non-overlapping twenty day intervals); and (5) fifty-day returns (for non-overlapping fifty day intervals). We log returns for all of these periods. Pre-crisis, banks that are active throughout this period (2002-2007) have approximately 1500 daily return observations, 300 five-day return observations, 150 ten-day return observations, and so forth.

We then compute the variance of each of these samples before and after the crisis. We compute the variance ratio using the variance of log returns for each interval relative to the variance of our daily returns.

Our goal is to ascertain whether (1) there is evidence of positive auto-correlation in the pre-crisis period and (2) there is more positive autocorrelation in the pre-crisis period relative to the post-crisis period (which would be an argument in favor of the market misunderstanding of risk hypothesis). If there is more positive serial correlation in the pre-crisis period, we would expect that variance ratios have decreased in the aftermath of the Great Recession for each of our intervals; for example, $\frac{ret_{5_{pre}}}{ret_{1_{pre}}} > \frac{ret_{5_{post}}}{ret_{1_{post}}}$.

We present results first for our big-6 banks (Table I) and then for the other large US financial institutions (Table II) separately in the pages that follow.

I'm not sure exactly how to interpret these very simple comparisons, but overall it looks like for most of our ratios (except ten-day log returns), on average the big-6 banks had higher variance ratios in the pre-crisis relative to the post-crisis period. In many cases, these are rather minute differences as the variance ratios look fairly similar in both periods.

We then do this same analysis for other large domestic financial institutions. There are 44 banks for which we have returns data in both the pre- and post-crisis period. For all of our variance ratios, we see that the majority (and in the case of some ratios, the large majority) of financial institutions have higher variance ratios in the pre-crisis period.

However, just comparing pre- and post-crisis variance ratios does not tell us much about whether there actually is any statistically significant autocorrelation in the pre- or post-crisis period. From a simple glance at our results, it seems that there is more likely to be evidence of *negative* autocorrelation than positive autocorrelation in our samples. In order to test for whether there is significant autocorrelation, we can use the following test statistic (from the Campbell-Ho-MacKinlay textbook):

The asymptotic distribution of $\hat{VR}(q)$ is

$$\sqrt{nq}(\hat{VR}(q) - 1) \sim N(0, 2(q - 1))$$

where n is the number of observations in each of our samples (i.e. 300 for five-day VR) and q is the number of days in each interval (i.e. 5 for five-day VR). Note that in this formulation, the VR is always centered around 1 because the test-statistic calls for dividing by the number of days in each of our intervals (5, 10, 20, and so forth).

We test for statistical significance of autocorrelation for our Big-6 banks. We do not repeat this exercise for the other large US financial institutions in this pass, though we certainly can in the future. This is slightly more complicated because there is often missing data either pre- or post-crisis for these institutions and we will need to use the exact number of observations we have for each bank in our sample for these computations.

In Table III, we present our same variance ratio table for the Big-6 but now with demarkations of statistical significance for our VR. There is no evidence of any positive serial correlation at any period for any of the

Big-6 banks. In fact, to the extent that there is any serial correlation that is statistically significant, it is negative serial correlation. Note that especially when we get up to the twenty and fifty-day windows, we lose a lot of power because we have so few observations. Lo and MacKinlay have a way to implement this test statistic for overlapping periods which we could try to do, though I doubt we'll get very different results.

Campbell also discusses (in his textbook and his course) how results for autocorrelations for individual stocks are related to the time horizon of the return measurements. Our results are broadly consistent with Nagel (2013) who documents small negative autocorrelations for individual stocks in daily, weekly, and even monthly data.