

Equiprobable Approximation to Bivariate Lognormal Returns

This handout starts by defining a convenient notation to represent two rate-of-return shocks that are distributed according to a multivariate lognormal that allows for nonzero covariances. It next presents a computationally simple method for constructing a numerical approximation to that joint distribution. The final section uses both the numerical approximation, and more accurate (but enormously slower) standard numerical integration tools to assess the accuracy of the Campbell-Vicera analytical approximation to the solution to the optimal portfolio choice problem described in [Portfolio-Multi-CRRA](#).

1 Statistical Theory

Consider a set of two normally distributed risks

$$\begin{aligned}\theta_{2,t+1} &\sim \mathcal{N}(-0.5\Sigma^2, \Sigma^2) \\ \theta_{1,t+1} &\sim \mathcal{N}(-0.5(x\Sigma)^2, (x\Sigma)^2)\end{aligned}$$

which are statistically independent ($\theta_{1,t+1} \perp \theta_{2,t+1}$) even though the *scale* of shock 1's *standard deviation* is determined by the proportionality factor x multiplied by the value of shock 2's standard deviation, Σ . (This permits us to increase or decrease the size of both risks by changing Σ and to adjust the relative size of risk 1 vs risk 2 by adjusting x .)

The θ are interpreted as the logarithms of level variables Θ , and the means of the log variables have been chosen such that the expectations of the levels are independent of the size of the risk (cf. [MathFacts](#)). That is, defining $\Theta_i = \mathbb{E}_t[e^{\theta_{i,t+1}}]$ for $i \in \{1, 2\}$:

$$\begin{aligned}\Theta_i &\equiv \mathbb{E}_t[\Theta_{i,t+1}] = 1 \quad \forall i \\ \log \mathbb{E}_t[\Theta_{i,t+1}] &= 0 \quad \forall i.\end{aligned}\tag{1}$$

From the first shock we can construct a log rate-of-return variable that can be represented equivalently in either of two ways:

$$\begin{aligned}\mathbf{r}_{1,t+1} &= \overbrace{\mathbf{r}_1 + 0.5(x\Sigma)^2}^{\equiv \mathbf{r}_1} + \theta_{1,t+1} \\ &= \underbrace{\mathbf{r}_1 - 0.5(x\Sigma)^2}_{=\mathbf{r}_1} + \underbrace{0.5(x\Sigma)^2 + \theta_{1,t+1}}_{\equiv \theta_{1,t+1}}\end{aligned}\tag{2}$$

where $\theta_{1,t+1} \sim \mathcal{N}(0, (x\Sigma)^2)$; the notation for θ is motivated by the fact that the addition of the extra term cancels the nonzero mean of the original θ . Then

$$\begin{aligned}\mathbf{R}_{1,t+1} &\equiv e^{\mathbf{r}_{1,t+1}} \\ \mathbf{R}_1 &\equiv \mathbb{E}_t[\mathbf{R}_{1,t+1}] = e^{\mathbf{r}_1 + (x\Sigma)^2/2} \\ \log \mathbb{E}_t[\mathbf{R}_{1,t+1}] &= \mathbf{r}_1 + (x\Sigma)^2/2 = \mathbf{r}_1\end{aligned}\tag{3}$$

where note to avoid confusion that $\mathbf{r}_1 \neq \mathbb{E}_t[\mathbf{r}_{1,t+1}]$ while $\log \mathbb{E}_t[\mathbf{R}_{1,t+1}] \neq \mathbf{r}_1 = \mathbb{E}_t[\mathbf{r}_{1,t+1}]$.

Using $\mathbb{E}_t[\theta_{2,t+1}] = -0.5\Sigma^2$ and $\mathbb{E}_t[(\omega/x)\theta_{1,t+1}] = -0.5((\omega/x)\Sigma)^2$, we can by analogy define a second return

$$\begin{aligned}\mathbf{r}_{2,t+1} &\equiv \mathbf{r}_2 + \zeta + (\omega/x)\theta_{1,t+1} + \theta_{2,t+1} \\ &= \mathbf{r}_2 + \zeta + (\omega/x)(\theta_{1,t+1} - 0.5(x\Sigma)^2) + (\theta_{2,t+1} - 0.5\Sigma^2) \\ &= \underbrace{\mathbf{r}_2 + \zeta - 0.5(\omega/x)(x\Sigma)^2 - 0.5\Sigma^2}_{\mathbf{r}_2} + (\omega/x)\theta_{t+1,1} + \theta_{2,t+1}\end{aligned}\tag{4}$$

for some constants ω and ζ .

Since $(\omega/x)\theta_{1,t+1}$ is the only component of $\mathbf{r}_{2,t+1}$ that is correlated with $\mathbf{r}_{1,t+1}$,

$$\begin{aligned}\text{cov}(\mathbf{r}_{1,t+1}, \mathbf{r}_{2,t+1}) &= \text{cov}(\mathbf{r}_{1,t+1}, (\omega/x)\mathbf{r}_{1,t+1}) \\ &= (\omega/x) \underbrace{\text{cov}(\mathbf{r}_{1,t+1}, \mathbf{r}_{1,t+1})}_{=x^2\Sigma^2} \\ &= \omega x \Sigma^2.\end{aligned}$$

Thus, the parameter ω controls the covariance between the risky returns. If we set $\omega = 0$ then $\mathbf{r}_{1,t+1} \perp \mathbf{r}_{2,t+1}$ (the returns are independent).

Next we want to find the value of ζ such that the expected level of the return is unaffected by Σ (so that we will be able to explore independently the distinct effects of the components of each shock and their covariance):

$$\mathbf{R}_2 \equiv \mathbb{E}_t[\mathbf{R}_{2,t+1}] = e^{\mathbf{r}_2}\tag{5}$$

regardless of the values of x and Σ . Using (4), we therefore need:

$$\begin{aligned}\mathbb{E}_t[e^{\zeta + (\omega/x)\theta_{1,t+1} + \theta_{2,t+1}}] &= 1 \\ \log \mathbb{E}_t[e^{\zeta + (\omega/x)\theta_{1,t+1} + \theta_{2,t+1}}] &= 0.\end{aligned}\tag{6}$$

Using standard facts about lognormals (cf. **MathFacts**), and for convenience defining $\hat{\omega} = (\Sigma/x)\omega$, we have

$$\begin{aligned}0 &= \zeta - 0.5\hat{\omega}x^2 - 0.5\Sigma^2 + 0.5\hat{\omega}^2x^2 + 0.5\Sigma^2 \\ &= \zeta - 0.5x^2\hat{\omega}(1 - \hat{\omega}) \\ \zeta &= 0.5(\hat{\omega} - \hat{\omega}^2)x^2 = 0.5(\omega x - \omega^2)\Sigma^2\end{aligned}\tag{7}$$

which means that we can rewrite (4) directly as

$$\begin{aligned}\mathbf{r}_{2,t+1} &= \mathbf{r}_2 + 0.5(\omega x - \omega^2)\Sigma^2 + (\omega/x)\theta_{1,t+1} + \theta_{2,t+1} \\ &= \mathbf{r}_2 + 0.5(\omega x - \omega^2)\Sigma^2 - 0.5(\omega/x)(x\Sigma)^2 - 0.5\Sigma^2 + (\omega/x)\theta_{t+1,1} + \theta_{2,t+1} \\ &= \mathbf{r}_2 + 0.5\omega x \Sigma^2 - 0.5\omega x \Sigma^2 - 0.5\omega^2 \Sigma^2 - 0.5\Sigma^2 + (\omega/x)\theta_{t+1,1} + \theta_{2,t+1} \\ &= \underbrace{\mathbf{r}_2 - (1 + \omega^2)0.5\Sigma^2}_{\mathbf{r}_2} + (\omega/x)\theta_{t+1,1} + \theta_{2,t+1}.\end{aligned}$$

Hence, from the independent mean-one lognormally distributed shocks Θ_1 and Θ_2 , we have constructed a pair of jointly lognormally distributed shocks whose covariance is controlled by the parameter ω , whose relative and absolute variances are controlled by the parameters x and Σ , and whose means are $\mathbf{r}_1$ and $\mathbf{r}_2$.

To sum up, the process can be described in either of two ways:

$$\begin{pmatrix} \mathbf{r}_{1,t+1} \\ \mathbf{r}_{2,t+1} \end{pmatrix} \sim \begin{pmatrix} \mathcal{N}(\mathbf{r}_1, x^2 \Sigma^2) \\ \mathcal{N}(\mathbf{r}_2, \Sigma^2) \end{pmatrix} = \begin{pmatrix} \mathcal{N}(\mathbf{r}_1 - 0.5(x\Sigma)^2, x^2 \Sigma^2) \\ \mathcal{N}(\mathbf{r}_2 - (1 + \omega^2)0.5\Sigma^2, \Sigma^2) \end{pmatrix}$$

with covariance matrix

$$\begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix} \equiv \begin{pmatrix} (x\Sigma)^2 & x\omega\Sigma^2 \\ x\omega\Sigma^2 & (\omega^2 + 1)\Sigma^2 \end{pmatrix} \quad (8)$$

where a final useful result that follows from (8) and (1) is

$$\begin{aligned} \text{corr}(\mathbf{r}_{1,t+1}, \mathbf{r}_{2,t+1}) &\equiv \text{cov}(\mathbf{r}_{1,t+1}, \mathbf{r}_{2,t+1}) / \sigma_1 \sigma_2 \\ &= \omega / (1 + \omega^2)^{0.5}. \end{aligned} \quad (9)$$

2 Computational Theory

To reduce clutter, define $\tilde{\mathbf{R}}_i = \mathbf{R}_{i,t+1}$ and interpret $\mathbb{E}$ as $\mathbb{E}_t[\cdot]$, so that we can write the expectation of some function h that depends on the realization of the return shocks as:

$$\mathbb{E}[h(\tilde{\mathbf{R}}_1, \tilde{\mathbf{R}}_2)] = \int_{\tilde{\mathbf{R}}_1}^{\tilde{\mathbf{R}}_1} \int_{\tilde{\mathbf{R}}_2}^{\tilde{\mathbf{R}}_2} h(\tilde{\mathbf{R}}_1, \tilde{\mathbf{R}}_2) dF(\tilde{\mathbf{R}}_1, \tilde{\mathbf{R}}_2) \quad (10)$$

where $F(\tilde{\mathbf{R}}_1, \tilde{\mathbf{R}}_2)$ is the cumulative distribution function for a multivariate lognormal with the covariance matrix defined by (8).¹ Standard numerical computation software can compute this double integral, but at such a slow speed as to be unusable for many purposes. Computation of the expectation can be massively speeded up by advance construction of a numerical approximation to $F(\tilde{\mathbf{R}}_1, \tilde{\mathbf{R}}_2)$.

Such approximations often take the approach of replacing the distribution function with a discretized approximation to it; appropriate weights $w[i, j]$ are attached to each of a finite set of points indexed by i and j , and the approximation to the integral is given by:

$$\mathbb{E}[h(\tilde{\mathbf{R}}_1, \tilde{\mathbf{R}}_2)] \approx \sum_{i=1}^n \sum_{j=1}^m h(\hat{\mathbf{R}}_1[i, j], \hat{\mathbf{R}}_2[i, j]) w[i, j] \quad (11)$$

where various methods are used for constructing the weights $w[i, j]$ and the nodes (corresponding to the $\{i, j\}$ pairs). The matrices $\hat{\mathbf{R}}_1$ and $\hat{\mathbf{R}}_2$ contain the conditional mean values of $\tilde{\mathbf{R}}_1$ and $\tilde{\mathbf{R}}_2$ associated with each of the regions.

Perhaps the most popular such method is Gauss-Hermite interpolation (see Judd (1998) for an exposition, or Kopecky and Suen (2010) for a recent candidate for a better choice). Here, we will pursue a particularly simple and intuitive alternative: Equiprobable discretization. In this method, $m = n$ and boundaries on the joint CDF are determined in such a way as to divide up the total probability mass into submasses of equal size (each of which therefore has a mass of n^{-2}). This is conceptually easier if we represent the underlying shocks as statistically independent, as with Θ_1 and Θ_2 above; in that case, each submass is a square region in the Θ_1 and Θ_2 grid. We then

¹Under the lognormal assumption, $\tilde{\mathbf{R}}_i = 0$ and $\tilde{\mathbf{R}}_i = \infty$ for both shocks $i \in \{1, 2\}$.




Figure 1 Portfolio Share By Method as a Function of ω

compute the average value of Θ_1 and Θ_2 *conditional* on their being located in each of the subdivisions of the range of the CDF. Since Θ_1 and Θ_2 are IID, the representation of the approximating summation is even simpler than (11):

$$\mathbb{E} \left[h \left(\tilde{\mathbf{R}}_1, \tilde{\mathbf{R}}_2 \right) \right] \approx n^{-2} \sum_{i=1}^n \sum_{j=1}^n h \left(\mathbf{R}_1(\hat{\Theta}_1[i], \hat{\Theta}_2[j]), \mathbf{R}_2(\hat{\Theta}_1[i], \hat{\Theta}_2[j]) \right)$$

where $\mathbf{R}_i(\Theta_1, \Theta_2)$ are the (linear) functions relating the return shocks to the IID shocks.

3 Computational Results

Figure 1 compares the computed optimal portfolio share for a numerical solution using the built-in numerical optimizer and maximization functions (the lowest, black, locus), the Campbell-Viceira solution (the highest, red locus) and an equiprobable approximation using 20 approximation points (green, middle) as well as the solution using the equiprobable approximation at an evenly-spaced grid of points (blue dots).

Careful examination indicates that the numerical approximation is quite close to the full numerical solution, while the CV approximation diverges substantially from the numerical answer. The tradeoff is that the equiprobable solution is about 2000 times slower than the CV approximation, while the direct solution is more than 100 times

slower than the equiprobable solution. Depending on the requirements of the problem being examined, these differences in efficiency can make a tremendous difference in the feasibility of a research project.²

References

- JUDD, KENNETH L. (1998): *Numerical Methods in Economics*. The MIT Press, Cambridge, Massachusetts.
- KOPECKY, KAREN A., AND RICHARD M.H. SUEN (2010): “Finite State Markov-Chain Approximations To Highly Persistent Processes,” *Review of Economic Dynamics*, 13(3), 701–714, <http://www.karenkopecky.net/RouwenhorstPaper.pdf>.

²Details can be found in the *Mathematica* notebook associated with this handout.