Notes on risk-neutral distributions

All equation references are to `BCM disasters ALT Jan 05 09 MC.pdf`.

Suppose, as in equation (10), that

$$\log \frac{c_t}{c_{t-1}} = \mu_c - \sigma_c^2/2 - \omega_c \left(e^{\theta_c + \delta_c^2/2} - 1 \right) + \sigma_c \varepsilon_{c,t} + z_{c,t}$$

where $\varepsilon_{c,t} \sim N(0,1)$ and $z_{c,t}$ is an independent random variable that captures the effects of disasters. **At this point I deviate from the setup currently in the paper.** *Roughly speaking*, we can think of $z_{c,t}$ as a random variable that is zero with probability $1-\omega_c$, and takes some Normally distributed value (the size of the disaster to log consumption) with probability ω_c. To be precise, it is convenient to define $z_{c,t}$ as follows.

$$z_{c,t} \sim \begin{cases} 0 & \text{w.p. } e^{-\omega_c} \\ N(\theta_c, \delta_c^2) & \text{w.p. } \omega_c e^{-\omega_c} \\ \vdots & \vdots \\ N(n\theta_c, n\delta_c^2) & \text{w.p. } \frac{\omega_c^n}{n!} e^{-\omega_c} \\ \vdots & \vdots \end{cases} \tag{1}$$

The number, n, of disasters that take place in a given period follows a Poisson distribution with parameter ω_c ("$Po(\omega_c)$"). Conditional on n disasters, the disaster sizes are each independent $N(\theta_c, \delta_c^2)$ random variables, so the sum of the disaster sizes is distributed $N(n\theta_c, n\delta_c^2)$.

Since the disaster probability, ω_c, is small, the probability of no disasters occurring in a given period $e^{-\omega_c} \approx 1 - \omega_c$, the probability of one disaster occurring $\omega_c e^{-\omega_c} \approx \omega_c$ and the probability of more than one disaster occurring is less than ω_c^2, hence extremely small. Thus definition (1) is *almost* equivalent to assuming that there is no disaster with probability $1 - \omega_c$, and a single disaster—with size distribution $N(\theta_c, \delta_c^2)$—with probability ω_c; but this apparently simpler assumption would greatly complicate the mathematics.

The probability distribution function of the Normal shock size and disaster shock size is

$$p_{\varepsilon,z}(x,y) = \underbrace{\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2}}_{N(0,1)} \cdot \sum_{n=0}^{\infty} \underbrace{\frac{\omega_c^n}{n!} e^{-\omega_c}}_{Po(\omega_c)} \underbrace{\frac{1}{\sqrt{2\pi n \delta_c^2}} e^{-\frac{1}{2} \frac{(y-n\theta_c)^2}{n\delta_c^2}}}_{N(n\theta_c, n\delta_c^2)}$$

Roughly speaking, this represents the probability (density) that state (x, y) occurs: $\varepsilon_{c,t} = x$ and $z_{c,t} = y$. From now on, I will drop the subscripts ε and z and denote the real-world probability (distribution function) as $p(x, y)$.

Motivated by the finite-state logic, we can compute the risk-adjusted (or "risk-neutral") probability distribution function via

$$p^*(x, y) = p(x, y)M(x, y)R_f \tag{2}$$

where $M(x, y)$ is the value taken by the stochastic discount factor in state (x, y). With power utility, $M = e^{-\rho-\alpha \log C_t/C_{t-1}}$, so

$$M(x, y) = \exp\left\{-\rho - \alpha\left[\mu_c - \frac{1}{2}\sigma_c^2 - \omega_c\left(e^{\theta_c+\frac{1}{2}\delta_c^2} - 1\right)\right] - \alpha\sigma_c x - \alpha y\right\}$$

The normalization factor R_f in (2) simply ensures that p^* integrates to 1, so in the calculations that follow, we can ignore constants of proportionality and simply keep track of the "shape" of distributions. Using (2), we have

$$\begin{aligned} p^*(x, y) &\propto \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2 - \alpha\sigma_c x} \cdot \sum_{n=0}^{\infty} \frac{\omega_c^n}{n!} \frac{1}{\sqrt{2\pi n\delta_c^2}} e^{-\frac{1}{2}\frac{(y-n\theta_c)^2}{n\delta_c^2}} e^{-\alpha y} \\ &\vdots \\ &\propto \underbrace{\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(x+\alpha\sigma_c)^2}}_{N(-\alpha\sigma_c, 1)} \cdot \sum_{n=0}^{\infty} \underbrace{\frac{\left(\omega_c e^{\frac{1}{2}\alpha^2\delta_c^2 - \alpha\theta_c}\right)^n}{n!}}_{Po\left(\omega_c e^{\frac{1}{2}\alpha^2\delta_c^2 - \alpha\theta_c}\right)} \underbrace{\frac{1}{\sqrt{2\pi n\delta_c^2}} e^{-\frac{1}{2}\frac{\left[y - n\left(\theta_c - \alpha\delta_c^2\right)\right]^2}{n\delta_c^2}}}_{N(n(\theta_c - \alpha\delta_c^2), n\delta_c^2)} \end{aligned} \tag{3}$$

Under the risk-adjusted distribution, $\varepsilon_{c,t}$ is still Normally distributed, but its mean shifts from 0 down to $-\alpha\sigma_c$; the rate of disaster arrivals increases from ω_c to $\omega_c^* = \omega_c e^{\frac{1}{2}\alpha^2\delta_c^2 - \alpha\theta_c}$ (which is greater than ω_c because if disasters are bad news on average, then $\theta_c < 0$); and the jump size distribution is Normal with mean shifted from θ_c down to $\theta_c^* = \theta_c - \alpha\delta_c^2$. As one would expect, increasing risk aversion magnifies the distinction between the risk-adjusted and real-world probability distributions.

As in the current version, you can look at things like the jump security, whose price is

$$\frac{1}{R_f}\mathbb{P}^*\left(\text{at least one jump}\right) = \frac{1}{R_f}\left(1 - e^{-\omega_c^*}\right)$$

and whose excess return (really, log excess-return-ratio...) is therefore

$$\log\left(\frac{\mathbb{P}\left(\text{at least one jump}\right)}{1 - e^{-\omega_c^*}}\right) = \log\left(\frac{1 - e^{-\omega_c}}{1 - e^{-\omega_c^*}}\right) \approx \log(\omega_c/\omega_c^*) = \alpha\theta_c - \frac{1}{2}\alpha^2\delta_c^2$$

which is negative because, of course, the security is a hedge.

You can also compute the returns on a "size security"—I did this but didn't get particularly neat expressions. I don't think the calculations that are in the paper at the moment are right: eg, equations (19)–(22) (and also for a different reason (18) as noted in my email) are incomplete, because you can't just price a security conditionally, you have to specify what it pays off in each state of the world. (Thus, roughly speaking, equation (19) is missing a factor that accounts for the probability that the security pays off.)

I think that we can also now calculate option prices using (3)—hopefully getting the same results as before...