Advanced Macro: The Log-Normal Distribution

Eric Sims
University of Notre Dame

Spring 2020

1 Introduction

Many of the papers in the CSV literature make use of the log-normal distribution. This note describes some properties of the distribution that come in handy.

2 A Few Things about Densities and Distributions

I'm not very good with probability and statistics. And integrals always kind of scare me. It's important to remember that integrals are just sums over continuous variables (rather than discrete, where we have straight old summation operators).

2.1 Continuous Random Variables

Let x be some random variable. For now, suppose it is continuous over $(-\infty, \infty)$. Let the *density* (or probability density function, pdf) be given by $f(x)$. Let the *cumulative distribution function* (or cdf, or what we'll often just call the distribution) be $F(x)$. Provided the distribution is differentiable, we have:

$$f(x) = F'(x) \tag{1}$$

In other words, the density is the first derivative of the distribution. The distribution measures the probability that the random variable is less than or equal to some cutoff value, $\bar{x}$. This is given by:

$$F(\bar{x}) = \int_{-\infty}^{\bar{x}} f(x)dx \tag{2}$$

Although it's not quite right in the continuous case (it is in the discrete case), we can think of the density evaluated at a point, $\bar{x}$, as giving the probability that the random variable equals that realization: $f(\bar{x}) = P(x = \bar{x})$. This isn't quite right in the continuous case because the probability of any single realization is zero. Hopefully this will be clearer in the discrete case. So the distribution is effectively just the sum (remember, an integral is a sum) of the probabilities that x takes on possible values up to $\bar{x}$. Since x must take on *some* value, we have:

$$\int_{-\infty}^{\infty} f(x)dx = 1 \tag{3}$$

This means that:

$$1 - F(\bar{x}) = \int_{\bar{x}}^{\infty} f(x)dx \tag{4}$$

The unconditional expectation of the random variable is:

$$\mathbb{E}(x) = \int_{-\infty}^{\infty} xf(x)dx \tag{5}$$

Again, this isn't quite right in the continuous case, but this is essentially the probability-weighted sum of all possible realizations of the random variable x. The *partial expectation* of the random variable is defined as:

$$g(\bar{x}) = \int_{-\infty}^{\bar{x}} xf(x)dx = \mathbb{E}[x|x \leq \bar{x}]P(x \leq \bar{x}) = \mathbb{E}[x|x \leq \bar{x}]F(\bar{x}) \tag{6}$$

Hence, we can write the conditional expectation as the partial expectation divided by the probability that x is in a region, which is given by the distribution function:

$$\mathbb{E}[x|x \leq \bar{x}] = g(x)F(\bar{x})^{-1} \tag{7}$$

2.2 Discrete Random Variables

I think it's easier to think about the world in discrete terms. Let's consider a particularly simple discrete distribution – uniform over the interval $[1, 10]$, or $x \sim U(1, 10)$. The probability mass function (technically we should use the term mass not density for the discrete case, but it is functionally the same thing) is just the inverse of the number of potential realizations. Letting there by n possible realizations, we'd have:

$$f(x) = \frac{1}{n} \tag{8}$$

The unconditional expectation is just the probability weighted sum of potential outcomes. For this particular example:

$$\mathbb{E}[x] = \sum_{x=1}^{n} xf(x) = \sum_{x=1}^{10} \frac{x}{10} = \frac{1}{10}(1 + 2 + 3 + 4 + 5 + 6 + 7 + 8 + 9 + 10) = 5.5 \tag{9}$$

The cumulative distribution is just the probability of $x \leq \bar{x}$. So, for example, we'd have for this distribution:

$$F(5) = \frac{1}{10} + \frac{1}{10} + \frac{1}{10} + \frac{1}{10} + \frac{1}{10} = \frac{1}{2} \tag{10}$$

In other words, there is a 50 percent change of drawing 5 or less from a $U(1, 10)$ discrete distribution.

The partial expectation of $x \leq 5$ is:

$$g(5) = \sum_{x=1}^{5} \frac{x}{10} = \frac{1}{10}(1 + 2 + 3 + 4 + 5) = \frac{3}{2} \tag{11}$$

The conditional expectation effectively re-weights the probabilities – if you condition on knowing $x \geq 5$, there is a 20 percent chance of each realization, not a 10 percent chance. So:

$$\mathbb{E}[x|x \leq 5] = \sum_{x=1}^{5} \frac{x}{5} = \frac{1}{5}(1 + 2 + 3 + 4 + 5) = 3 \tag{12}$$

As above in the continuous, we can take (10)-(12) to realize the conditional expectation is just the partial expectation divided by the distribution function:

$$\mathbb{E}[x|x \leq 5] = \frac{g(5)}{F(5)} = \frac{3}{2} \times \frac{2}{1} = 3 \tag{13}$$

3 The Log-Normal

Let ω be a random variable. We assume that:

$$\ln \omega \sim N(\mu, \sigma^2) \tag{14}$$

Note that the support for ω must be $(0, \infty)$, since you can't take the log of something negative. $N(\cdot)$ is the normal distribution, μ is the mean, and σ^2 is the variance.

Let $\Phi(\cdot)$ and $\phi(\cdot)$ be the cumulative distribution function and density function for a standard normal distribution (i.e. $N(0, 1)$). Then we have the cumulative distribution and density of the log-normal random variable satisfying:

$$F(\omega) = \Phi\left(\frac{\ln \omega - \mu}{\sigma}\right) \tag{15}$$

$$f(\omega) = F'(\omega) = \phi\left(\frac{\ln \omega - \mu}{\sigma}\right)\frac{1}{\omega\sigma} \tag{16}$$

Note that $\phi(\cdot) = \Phi'(\cdot)$; the multiplication by the inverse of $\omega\sigma$ is effectively the "derivative of the inside" part of the chain rule. Note that the log-normal density is given by:

$$f(\omega) = \frac{1}{\omega}\frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln \omega - \mu)^2}{2\sigma^2}\right) \tag{17}$$

As above, the cumulative distribution, $F(\cdot)$, is just the probability that $\omega \leq \bar{\omega}$, for some $\bar{\omega}$. That is:

$$F(\bar{\omega}) = \int_0^{\bar{\omega}} f(\omega) d\omega \tag{18}$$

The expected value satisfies:

$$E[\omega] = \int_0^{\infty} \omega f(\omega) d\omega \tag{19}$$

Again, in the discrete analog, this is just the weighted average of realizations, ω, times probabilities, $f(\omega)$. For this particular distribution, the mean works out to:

$$\mathbb{E}[\omega] = \exp\left(\mu + \frac{1}{\sigma^2}\right) \tag{20}$$

For most of the applications we deal with, we will need $\mathbb{E}[\omega] = 1$. This means that $\mu + \frac{1}{2}\sigma^2 = 0$, so we have:

$$\mu = -\frac{1}{2}\sigma^2 \tag{21}$$

Once again, we might be interested in the *partial* expectation – i.e. the expected value of ω conditional on being in some region, times the probability of being in that region:

$$g(\bar{\omega}) = \mathbb{E}[\omega | \omega \leq \bar{\omega}] P(\omega < \bar{\omega}) = \int_0^{\bar{\omega}} \omega f(\omega) d\omega \tag{22}$$

In words, this is the expected value of ω, conditional on $\omega \leq \bar{\omega}$, times the probability that $\omega \leq \bar{\omega}$. For the log-normal distribution where $\mathbb{E}[\omega] = 1$, this works out:

$$\int_0^{\bar{\omega}} \omega f(\omega) d\omega = \Phi\left(\frac{\ln \bar{\omega} - \mu - \sigma^2}{\sigma}\right) \tag{23}$$

Where again $\Phi(\cdot)$ is the cdf of a normal distribution. Similarly, we have:

$$\int_{\bar{\omega}}^{\infty} \omega f(\omega) d\omega = \Phi\left(\frac{\mu + \sigma^2 - \ln \bar{\omega}}{\sigma}\right) \tag{24}$$

3.1 Leibniz Rule and Differentiating wrt an Integral Bound

There will be some instances in this literature where we are interested in some function of a cutoff value, $\bar{\omega}$, where this cutoff value appears as an integral bound.

For example, suppose we are interested in differentiating the partial expectation with respect to $\bar{\omega}$:

$$g(\bar{\omega}) = \int_{\bar{\omega}}^{\infty} \omega f(\omega) d\omega \tag{25}$$

Most of us are not used to differentiating with respect to an integral bound – we kind of intuitively know that we can move a derivative through an integral (since an integral is a effectively a sum, and the derivative is a linear operator), but we get scared when we see the variable we are differentiating with respect to not inside the integral but rather as one of the bounds.

Well, we can use Leibniz rule. In the general form it looks scary. Suppose we have:

$$\int_{a(x)}^{b(x)} f(x,z)dz \tag{26}$$

The derivative of this sucker with respect to x is:

$$\frac{d}{dx}\left[\int_{a(x)}^{b(x)} f(x,z)dz\right] = f(x,b(x))b'(x) - f(x,a(x))a'(x) + \int_{a(x)}^{b(x)} f_x(x,z)dz \tag{27}$$

In words, this is the function evaluated at the upper bound, times the derivative of the upper bound with respect to x; minus the function evaluated at the lower bound, evaluated at the lower bound, times the derivative of the lower bound with respect to x; plus the integral of the derivative of the function inside the integral.

This all may look very confusing. But you have functionally probably used this rule many times in your life when the bounds of the integral are constants and you simply differentiate inside the integral.

So now let's go back to (25). What is $g'(\bar{\omega})$? Using the general formula, we have:

$$g'(\bar{\omega}) = \infty f(\infty)\frac{d\infty}{d\bar{\omega}} - \bar{\omega} f(\bar{\omega})\frac{d\bar{\omega}}{d\bar{\omega}} + \int_{\bar{\omega}}^{\infty} \big(\omega f'(\omega) + f(\omega)\big)\frac{d\omega}{d\bar{\omega}}d\omega \tag{28}$$

Okay, this looks messy. But it's not. Why? Because ω and $\bar{\omega}$ are different – one has nothing to do with the other in some sense. So $\frac{d\infty}{d\bar{\omega}} = 0$, $\frac{d\bar{\omega}}{d\bar{\omega}} = 1$, and $\frac{d\omega}{d\bar{\omega}} = 0$. But this means only the middle term is left!

$$g'(\bar{\omega}) = -\bar{\omega} f(\bar{\omega}) = -\bar{\omega} F'(\bar{\omega}) \tag{29}$$

Via similar logic, suppose we were interested with a different partial expectation, say:

$$h(\bar{\omega}) = \int_0^{\bar{\omega}} \omega f(\omega)d\omega \tag{30}$$

We'd then have:

$$h'(\bar{\omega}) = \bar{\omega} f(\bar{\omega})\frac{d\bar{\omega}}{d\bar{\omega}} - \infty f(\infty)\frac{d0}{d\bar{\omega}} + \int_{\bar{\omega}}^{\infty} \big(\omega f'(\omega) + f(\omega)\big)\frac{d\omega}{d\bar{\omega}}d\omega \tag{31}$$

But via similar logic, this is just:

$$h'(\bar{\omega}) = \bar{\omega} f(\bar{\omega}) = \bar{\omega} F'(\bar{\omega}) \tag{32}$$

To get some intuition for this, return to the discrete uniform distribution, and remember that a derivative is basically just the change. This won't be exact given the discrete nature and the fact that derivatives are relevant for small changes and continuous variables, but it'll give us an idea. Suppose we have a uniform distribution over 1 to 5. Suppose we are interested in the partial expectation from 2 to 5, i.e.:

$$g(2) = \sum x = 2^5 \frac{x}{5} = \frac{1}{5}(2 + 3 + 4 + 5) = \frac{14}{5} \tag{33}$$

Now calculate the partial expectation from 3 to 5:

$$g(3) = \sum x = 3^5 \frac{x}{5} = \frac{1}{5}(3 + 4 + 5) = \frac{12}{5} \tag{34}$$

The difference is:

$$g(3) - g(2) = -\frac{2}{5} \tag{35}$$

But this is of course just the negative of the density, $\frac{1}{5}$, times the starting point, $\bar{x} = 2$, which is what we have for (30).

We could also do this in reverse. Suppose we are interested in:

$$h(3) = \sum_{x=1}^{3} \frac{x}{5} = \frac{1}{5}(1 + 2 + 3) = \frac{6}{5} \tag{36}$$

Now calculate the partial expectation where the upper bound is 4:

$$h(4) = \sum_{x=1}^{4} \frac{x}{5} = \frac{1}{5}(1 + 2 + 3 + 4) = \frac{10}{5} \tag{37}$$

The difference is:

$$h(4) - h(3) = \frac{4}{5} \tag{38}$$

This is just the new point of evaluation, 4, times the density. The starting point would be 3, however, which would give us $\frac{3}{5}$, not $\frac{4}{5}$. That's an issue with the discrete nature and not mapping perfectly into the derivative of a continuous random variable. But you can see that the formula Leibniz rule gives us in this case where we are differentiating with respect to an integral bound actually makes sense – we are calculating the change in the partial expectation when we change the bound.

4 Notational Issues

The papers in the CSV literature use different and often bad notation. For example, they will often assume that ω is has distribution $\Phi(\cdot)$ and density $\phi(\cdot)$. This is fine, but if you are thinking of Φ and ϕ as being normal distribution and density functions, respectively, you have to be a bit careful because the log-normal is a transformation of the normal. I'll try to be consistent and will sometimes use different notation than the relevant papers where I find them confusing, but I will probably fail.