

Random Switching Times Among Randomly Parameterized Regimes of Random Interest Rate Rate Scenarios

James G. Bridgeman, FSA

University of Connecticut 196 Auditorium Rd.U-3009 Storrs CT 06269-3009

bridgeman@math.uconn.edu

October 31, 2006

Abstract

Most of the usual stochastic interest rate models are designed and calibrated to provide plausible behavior in expected value and variance, which are the raw material for first approximation pricing and hedging of financial instruments. Standards of practice and often regulators require actuaries to stress-test financial positions over long projection horizons against extreme interest rate paths. The behavior of extreme paths in the usual stochastic interest rate models is not nearly so plausible as the behavior of their expected values and variances. This paper proposes a new class of models that deliver more plausible extreme paths while preserving the usual expected value and variance behavior. Along the way we derive the closed form solution for the traditional mean-reverting lognormal process, including the drift compensation.

1 Background

The random interest rate model that we will present evolved out of a practitioner's efforts to make sense of the traditional mean-reverting lognormal process model for interest rates in a decision-making stress-test context. The author discussed this historical background informally and with no mathematical detail in [2].

After a thorough formal background narrative in this section (still *sans* the mathematics) to motivate the model this paper will proceed in later sections to fill in the mathematics and to develop more fully the theory that emerges from the pragmatic modeling. The model that results, a generalization of the Black-Karasinski model ([4] page 155), may have interest beyond the decision-making stress-test context it came from. Moreover, the same logic will lead to corresponding generalizations of other models that incorporate mean-reversion.

1.1 Asset adequacy decisions

Throughout the 1990's and into the 2000's the author provided the asset adequacy opinion for a substantial life insurance company that did business exclusively in one Asian country (not Japan) but that also filed the NAIC statement and opinion with a U.S. state of domicile. A simple but very large duration mismatch with very little optionality (in practice) dominated the asset adequacy position.

At first the duration mismatch created little asset adequacy concern (and caused no failures of NAIC required interest rate scenarios.) Interest rates were high relative to pricing assumptions; the predominant products had high margins in elements other than the interest rate spread, including large volumes of highly profitable riders; and the reserves had high margins, including large amounts of catastrophe and contingency reserves required by the host country but not attributable to any specific insured contingency. (Under an election permitted by the state of domicile, the company reported reserves in the NAIC statement according to host country reserve requirements).

Out of prudence we explored what sort of future interest rate scenario would impair current asset adequacy in light of the duration mismatch. It turned out that only something as bad as a future decline in interest rates to 2% or less that persisted at that level for 10 to 15 years or more could ruin the current asset adequacy. With rates at the 8% level at first, surely we could judge the assets adequate? But on what logic or facts could we characterize the ruinous scenario as "more than moderately adverse"? The question took on more urgency as interest rates came down over the years.

Professional guidance at the time declared that "moderate adversity" set the threshold for asset adequacy but did not define "moderate adversity" in terms any more instructive than the phrase itself. For asset adequacy purposes, i.e. for purposes of deciding whether or not to put up additional reserves, should we or should we not deem "more than moderately adverse" a future decline in interest rates from 8% to 2% or less that stayed at that level for 10 to 15 years or more? A few years later, what about a decline from 6% to 2% and staying there? In the end, we had to ask what about staying *here* at 2% or 2.5% for another decade or more?

1.2 Random scenarios

Early in the experience, with interest rates still above 6%, we decided to validate the judgmental conclusion that the mismatch did not impair asset adequacy by running random interest rate scenarios generated as an option in the standard vendor cash flow testing software that we used. At that time the host country had insufficient volume of medium or long-dated instruments relative to potential demand to constitute a meaningful market. Given the company's policy almost entirely to avoid currency mismatch, short and floating rate instruments composed almost the whole portfolio. Our model could treat the host country yield curve as flat, consisting only of short (one or two year) or floating rates.

In effect, we modeled just a single, one to two year risk-free rate at each time in the model projections. This vastly simplified the interest rate modeling process and, fortuitously, allowed our attention to focus on other too-often neglected aspects of the interest rate modeling process.

An asset adequacy opinion requires cash flow stress-testing of the portfolio. Stress-testing entails the use of physical (i.e. real-world) interest rate scenarios rather than risk-neutral ones. This in turn requires the practitioner (at least in the opinion of this practitioner) to form a judgment about the appropriateness of the generated set of future interest rate scenarios as a representation (singly) of possible future worlds and (collectively) of the plausible range of possible future worlds.

The fact that an interest rate generator comes packaged with a respected vendor's software or has a name (say, "mean-reverting lognormal") that often appears in the literature fails by itself to satisfy that requirement. It requires in addition some thoughtful analysis of the set of scenarios that come out of the generator against the criterion "appropriately represents possible future worlds and the plausible range of possible future worlds." Of course such analysis almost surely will include thoughtful reference to, if not actual grounding in, history - the actual past real-world.

Since we didn't have a yield curve to worry about, as long as we compensated for drift the lognormal random process for the change in interest rates provided in the vendor cash flow testing software had impeccable risk-neutral credentials. As a representation of a plausible range of possible future real-worlds, however, it created a series of dilemmata for us.

1.3 Mean-reversion or no mean-reversion?

With no mean-reversion, if we set the volatility parameter value at historical levels it created a set of scenarios that included some with interest rates ultimately reaching 40% or 50% or even higher and some with interest rates ultimately reaching 1% or lower and staying in that neighborhood for many years. We didn't want to reject a parameter value just because of the latter outcome. That would amount to defining in advance as impossible the very scenarios that we needed to decide whether or not to worry about. But we felt quite comfortable about eliminating from consideration as completely implausible any parameter value that could produce even a couple of paths out of 1000 reaching into the 40% to 50% range. (The host country had no history, cultural inclinations, or monetary institutions even remotely suggestive of hyperinflation.)

It seemed inappropriate as a way out of the dilemma just arbitrarily to censor extreme scenarios, either by rejecting them and resampling or by an artifice such as reflection from a censorial barrier (at 20% or 25%, say). If a model produced absurd scenarios at any noticeable frequency with a given set of parameters then we took it as evidence that either the parameters or the model failed to represent appropriately the plausible range of future worlds and needed to be replaced.

Alternatively, we could choose an unhistorically small value for the volatility

parameter and easily assure an historically-plausible range of ultimate interest rates across all scenarios. But, even aside from being indefensible on historical grounds, a small volatility parameter value produced such a tight range of interest rates across all scenarios in near, intermediate, and even long (but short of ultimate) time horizons that just on inspection it failed adequately to represent a reasonable range of possible future worlds.

Clearly we needed to use the mean-reversion option in the vendor interest rate generator in order to eliminate any practical probability of producing scenarios that contained outlandishly high levels of interest rates. Perhaps that would produce a model that looked plausible in all other respects, and from which we could draw a conclusion about the likelihood of a future that contained the ruinous drop-low-and-stay-low-for-a-decade-or-more scenario that worried us.

1.4 What mean-reversion target parameter and mean-reversion speed parameter values?

Risk-neutral logic says that the current value of a market variable provides the best basis to predict its future value. The current interest rate at 8% when we started down this road almost demanded an alternative (lower) real-world assumption for the value of the mean-reversion target parameter. The historical average over the longest time for which we could find records of the rate we were trying to simulate provided the most obvious (and lowest historically-based) alternative value, around 6%, for the mean-reversion target parameter. Compared to higher then-current values of the market interest rate we could congratulate ourselves on our conservatism.

At first, by trial and error we set the volatility parameter and the mean-reversion speed parameter values jointly to produce

- a. an observed model volatility (i.e. standard deviation of the log-change in simulated interest rates *after* mean-reversion) equal to the historical volatility (standard deviation of log-change in the historical interest rate series),

and

- b. observed standard deviations of model interest rates (across scenarios at each point in time) that converged over time to the standard deviation of the set of historical interest rates.

In effect, we treated one long actual historical series of interest rates as an empirical sample in two different ways: once as a sample from the distribution of possible sequential changes in interest rates and once as a sample from the distribution of possible levels of interest rates.

1.5 The outcome

The resulting parameters produced a set of random interest rate scenarios from the model that confirmed our asset-adequacy conclusions. It contained a handful of ruinous (barely) scenarios but not nearly enough of them to require reserve strengthening, or even to raise risk-based capital concerns. The whole set of scenarios looked reasonable on first inspection. For example, no scenarios had interest rates reaching 40% to 50% or more.

Apparently we had reached our goal: a model that did not define away our ruinous scenarios by its very structure or parameters and that produced a plausible looking total set of scenarios, so that we could take some confidence in the low probabilities indicated for our ruinous scenarios.

1.6 A dilemma

On closer review, however, these parameters proved to have narrowed the range of interest rates too much, despite matching the historical mean and standard deviation in interest rates. At any selected time, the probability across scenarios to have interest rates in the low to mid teens fell short of the historical frequency of interest rates in the low to mid teens by just enough to cause concern.

In other words our model produced the right mean and standard deviation of interest rates and appeared to produce a thin enough upper tail to make outlandishly high interest rates vanishingly unlikely. But the mean reversion parameter ruled out the creation of broad enough shoulders in the interest rate distribution to reproduce accurately the historical probability of low to mid-teen interest rates. Perhaps an alternative model that did produce broad enough shoulders also would produce higher probabilities for the ruinous down-and-stay-down scenarios that worried us?

Alternatively, we could take the risk-neutral view and just set the value for the mean-reversion target equal to the current interest rate at the time (something higher than the historical average interest rate). Then using exactly the same logic as in section 1.4 to determine jointly the volatility and mean-reversion speed parameter values we got exactly the right probability of interest rates in the low to mid teens to match the historical value.

But with these parameter values, no ruinous scenarios occurred at all. With the higher value for the mean-reversion target parameter, the mean-reversion speed parameter pulled all scenarios up out of the ruinous range.

1.7 A worse dilemma

The mean-reversion mechanism caused the preceding dilemma. More trial and error to investigate that mechanism revealed a troubling phenomenon apparently associated with any mean-reversion model no matter what we did.

Essentially any combination of parameter values that reproduced model values consistent with historical values for the observed volatility of interest rates and for the average and standard deviation of achieved interest rate levels could

not possibly produce anything beyond low probabilities for the down-and-stay-down-for-a-decade-or-more scenarios that we worried about. While a lot of scenarios might touch the low interest rate levels that worried us, only a highly unlikely scenario could repeatedly counter the upward pull of the mean-reversion enough times to stay at low levels long enough to create ruinous results in the model.

Even visually, too many scenarios just oscillated in large swings around the target value. Where were the long runs apparent in any historical interest rate series?

We were in the position of trying to take comfort from the results of a model whose very structure made it impossible to produce more than an occasional scenario that could trouble us. It seemed inappropriate to base our judgment on a set of model outcomes guaranteed not to produce results leading to asset inadequacy conclusions. But models without mean-reversion had produced only absurd-looking sets of future worlds. What to do?

1.8 A way out

If we needed a mean-reversion target (to make extremely high levels of interest rates almost impossible) but the rigidity of the mean-reversion target value created other problems, why not add some motility to the mean-reversion target value? Why not make the mean-reversion target value something that changes from time to time?

Come to think of it, that might make economic sense. Sometimes we live in a world that favors higher interest rates, other times in a world that favors lower interest rates. That doesn't mean we live in a world that always pulls random interest rates toward an average level somewhere between high and low. Maybe we live in world that sometimes pulls random interest rates toward a level higher than average and other times toward a level lower than average.

That describes a mean reversion target whose value sometimes changes. In other words some core economic structure beneath the fluctuating economic environment might itself sometimes change.

1.9 When should the value of the mean-reversion target parameter change?

To specify exactly how long "sometimes" lasts would seem to introduce another kind of rigidity. So we made the length of time that the mean reversion target stays at any given value a random variable to be simulated along with the interest rates as each random scenario in the model unfolds over time.

An exponential random variable (the interarrival time for events in a Poisson process) seemed like a logical first choice to model the time between changes to the value of the mean-reversion target. We decided to generalize slightly and use an Erlang distribution (gamma distribution with integer α parameter; the interarrival time for α events to occur in a Poisson process) so as not to prejudge completely the shape of the distribution.

1.10 What should the value of the mean-reversion target parameter change to?

To specify the new value for the mean-reversion target whenever a time to change arises in a random scenario would seem to introduce yet another source of rigidity. What would be a good "high" value? A good "low" value? For that matter, why should the core structure underlying the economic environment just keep changing back and forth between just one "high" interest rate tendency and just one "low" interest rate tendency? Or even among a "high", "low, and "normal" value? Or any other number of values?

To choose arbitrarily some specific possible values for the mean reversion target might prejudice inadvertently the very question at issue: whether our ruinous scenarios can arise with enough probability in the model to warrant a decision to strengthen reserves.

So we chose not to choose. A new random variable established the new value for the mean-reversion target whenever a time to change arose in a random scenario. Nothing other than another lognormal distribution suggested itself for this new random variable representing the possible values of the mean-reversion target.

Note that this new lognormal variable represented the actual value of the mean-reversion target after it changes, not the process of change from old value to new. We reasoned that a major shift in the long-term central tendency of the economy caused by (what? war, peace, inflation, deflation, elections won or lost, appointments, retirements, globalization, protectionism, etc.) might have more to do with what happened to the world than with what the world looked like before it happened.

1.11 The resulting model

That's how we wound up with a mean-reverting lognormal random model of interest rates in which the value of the mean-reversion target was itself a random variable. We determined the value of the mean-reversion target (independently for each scenario) by independent (within each scenario) Erlang random variables determining the times when it would change values in that scenario and independent (for each scenario and time of change) lognormal random variables determining the value it would take each time that it changed.

Only when preparing for the 2001 Valuation Actuary Symposium talk [2] on these matters did Mary Hardy's paper [3] come to hand and we learned that this sort of thing has a name: a regime-switching model. Each time the value of our mean-reversion target changed to a new value the model was switching to a new regime.

The value for its mean-reversion target parameter characterized each of the possible regimes. Call it the *regime random variable*. Our model has a continuum of possible regimes defined by a lognormal distribution for the possible values of the regime random variable.

Each step on the way was dictated purely by a practitioner's desire

- a. not to rely on any model that for structural reasons might rule out inadvertently any real likelihood of the very scenarios that concerned him, but
- b. still to use a model that created an overall set of scenarios that in aggregate reproduced key statistics of past interest rates and looked plausible overall as a set of possible futures.

1.12 Calibrating the model - expected values

By this point we had more parameters to choose than obviously salient historical statistics to reproduce.

At first we continued to use the long-term historical average value of the interest rate for both the initial value of the mean-reversion target parameter and for the mean (*not* the μ parameter) of the lognormal regime random variable that determined the new value after a regime change. As long as the then-current beginning interest rate exceeded the historical average interest rate this provided conservatism. Section 1.14 discusses what we did when the starting interest rate became lower than the historical average.

Next we settled arbitrarily on the Erlang that has mean 20 years and mode 10 years (i.e $\alpha = 2, \beta = 10$) for the regime-switching interval. First, the mean assured that when the regime changed to an unfavorable mean-reversion target parameter value it had a fair chance of staying there long enough to cause damage. Second, the mode assured that a respectable number of scenarios actually would change regimes soon enough to worry about the new mean-reversion target parameter value.

To the latter point we also reasoned that at the current starting time we have no way to know how far we stand along the interval since the last historical regime change. The regime and its changes are unobservable variables. So we modeled the first regime change in each scenario to occur at a randomly selected time (uniform distribution) within the first Erlang random period generated in that scenario.

(When illustrating all this for U.S. interest rates in section 3 we will explore external rationale for selecting the Erlang parameters, such as Becker's suggestion [1], pages 7-12, that U.S. interest rate evolution couples (weakly) to the presidential election cycle.)

1.13 Calibrating the model - variances

We still had to select values for the mean reversion speed parameter and the volatility parameter for the mean-reverting lognormal process, and for the volatility parameter σ of the lognormal regime random variable that determined the mean-reversion target value for each new regime. The trick was to forget what we knew about what mean-reversion parameters are supposed to look like.

For any given value of σ , by trial and error we set the mean reversion speed parameter and the volatility parameter for the mean-reverting lognormal process jointly to produce

- a. an observed model volatility (i.e. standard deviation of the log-change in simulated interest rates *after* mean-reversion) equal to the historical volatility (standard deviation of log-change in the historical interest rate series),

and

- b. observed standard deviations of model interest rates (across scenarios at each point in time) that converged to some stable value within a reasonable period of time, five to ten years depending on a visual judgment of the reasonability of the resulting set of scenarios.

The random spread of values for the mean-reversion target parameter combined with requirement b. to generate a high value for the mean-reversion speed parameter. In turn this recombined with the random spread of values for the mean-reversion target parameter to create the broad shoulders that we needed to see in the observed distribution of model interest rates (across scenarios at each point in time) in order to reproduce the historical frequency of low to mid-teen interest rates.

Finally, we iterated this process through different values of σ until the stable value in b. above for the observed standard deviations of model interest rates matched the standard deviation of the set of historical interest rates.

The usual mean-reversion model has only the volatility parameter in the mean-reversion process to achieve a large enough standard deviation of observed interest rates to match the historical value. This new model achieved it primarily through the volatility parameter σ of the lognormal variable that determined each new regime.

1.14 Surprising payoffs

Naturally, because we built it that way, the model now reproduced more of the key aspects of the historical interest rate distribution and time series. But more benefits accrued.

Unplanned, but welcome, some heteroskedasticity turned up in the observed model volatility (i.e. standard deviation of the log-change in simulated interest rates *after* mean-reversion) for each scenario. Why? Observed model volatility has some correlation with time elapsed since the last regime-switch.

Immediately post-switch the likelihood for the scenario interest rate to sit far from the new mean-reversion target parameter value goes up. With it the likelihood for the mean-reversion speed parameter to induce several large changes in the scenario interest rate goes up. This effect did not become big enough to reproduce historical heteroskedasticity, but finding any heteroskedasticity at all in a process driven by a constant volatility parameter opens the mind.

As hoped, the model turned out to have real potential to produce ruinous scenarios, even with starting interest rates higher than the historical average. As expected, the probability of ruinous scenarios remained low enough to warrant a judgment of asset adequacy for high starting interest rates.

The real payoff came as the actual interest rates declined over the years. When the starting interest rate in the model lies below the historical average interest rate, what to do about the assumed values for the initial mean-reversion target parameter and for the mean of the lognormal regime random variable that determines the new target parameter value after a switch?

First, we could find no reason set the mean for the lognormal regime random variable any lower just because of a lower current starting interest rate (except to the extent that the current value slightly reduced the long term historical average interest rate). To do that would implicitly introduce regimes of regimes. The lognormal regime random variable as originally calibrated already anticipates the possibility of regimes of low interest rates. It does not itself need to assume to a lower range of possibilities (i.e. a lower expected value for a randomly selected mean-reversion target parameter value) just because the starting interest rate falls from one year to the next.

Second, we faced a more interesting question what to do about the initial value for the mean-reversion target parameter. When the starting interest rate exceeded the historical average it seemed easy and conservative to set the initial value for the mean-reversion target parameter down at the historical average level. In the reverse situation it came harder and seemed anti-conservative to leave it up at the historical average level, in this case higher than the current starting interest rate. But drop it all the way down to the current starting interest rate? A bad model in times of high interest rates can't become completely good just because interest rates fall.

The current value of the mean-reversion target parameter, the regime random variable, is never observable. Since it represents only a central tendency, to conclude that its value has changed in a year by an amount equal to the change in the current interest rate over that same year clearly overrates the evidential value contained in just one year's change in the interest rate. The question calls for credibility logic.

We experimented variously with credibility-weighted averages of the current interest rate and the long term historical average and with shorter term averages of most-recent history, in effect a moving-average approach from year to year. More important than the specific choices, the model structure itself gave us a very specific logical framework to explore in a very disciplined way the issues that actuaries around the world faced as interest rates declined. What counts as "moderately adverse" going forward when the current situation itself is adverse by historical reckoning? How reasonable is a conclusion that at certain times a scenario with interest rates frozen at their current levels forever in fact represents "more than moderate adversity"?

There may be no universal answer to such questions. It seems noteworthy, however, that the same model logic that we had created originally in order to minimize any inadvertent anti-conservatism in a high interest rate environment

provided a very natural way to moderate any inadvertent tendency toward excessively punishing conservatism in a low interest rate environment, and to do so without the need to introduce extraneous or *ad hoc* considerations.

As interest rates fell year by year the combination of lower starting interest rates and gradually lowering values (dampened by credibility) for the initial mean-reversion target, gradually increased the frequency of the ruinous scenarios in the model. Management started early to introduce corrective changes, where possible, in part because the model started showing the increased frequency of ruinous scenarios earlier than most. When those frequencies began to require reserve strengthening, it came in measured doses as credibility of a lower initial mean-reversion target grew from year to year.

If lower interest rates last forever the company will have pre-funded the emerging shortfall in a measured way over time. On the other hand, if as almost any common-sense forecast would indicate the host country interest rates eventually climb back out the pit then the magnitude of the reserve-strengthening automatically will have been limited, without the need for *ad hoc* expedients, to levels less punishing to policyholder prices and to shareholder capital than immediately assuming that a low current level of interest rates must last forever. And all will have been accomplished within a logical framework essentially consistent across a wide range of interest rate environments, both high and low.

1.15 One last perplexity

We encountered a small but vexing technical problem the very first time we tried the mean-reversion models and we never satisfactorily resolved it. It was a small effect that did not threaten our basic approach or conclusions but it did nag at our confidence that we completely understood the model and that the software did not have a bug.

Remember (section 1.2) that we needed real-world scenarios, not risk-neutral ones. So we had to compensate for the lognormal drift. As soon as we introduced mean-reversion the compensation for drift got screwy.

The usual "minus one-half the square of the volatility" now overcompensated and drove the average for the model interest rates (across scenarios at each point in time) below what you would expect from the starting value and the values of the mean-reversion target and speed parameters. In particular, the average model interest rates dropped over time to levels below the value of the mean-reversion target.

We dealt with this by trial and error to come up with a good drift compensation term by visual inspection of model output graphics. But any change to the values of the volatility parameter and the mean-reversion speed parameter disrupted the drift compensation all over again and required more trial and error to correct it. This went on for years.

Even the best trial and error drift compensation values we could come up with still annoyed us. If they resulted in the average model interest rates ultimately converging over a long time to the value of the mean-reversion target parameter (which they should for a real-world model) then they caused the

average model interest rates to hover just a little too far, for far too long a time, above the trajectory you would expect, running from starting value to mean-reversion target parameter value, based on the mean-reversion speed parameter.

The most frustrating situations occurred when the starting value equalled the value of the mean-reversion target parameter. Then the average interest rates across scenarios drifted up a little and then back down ultimately to the mean-reversion target parameter value. It was as if the drift compensation did not fully kick in until many years into the scenarios.

Since concern about low interest scenarios led us to these models in the first place it annoyed us to have even this small upward bias to the interest rates, in early years especially, still embedded in the model. Often, we chose a trial and error value for the drift compensation high enough to eliminate even this small early upward bias, and just accepted an average interest rate (across scenarios) that bewilderingly drifted a little too low, creeping over time to ultimate levels below the mean-reversion target parameter value.

Section 2.2 below came as sunshine to scatter the miasma.

2 Mathematics of the model

Remember (section 1.2) that we built a real-world interest rate model for stress-testing purposes, not a risk-neutral one, and that term structure essentially did not exist for us because of certain unique facts in the particular stress-testing situation. It may be that this one-factor model with further development can anchor two- or three- factor models, real-world or risk-neutral. This paper sticks to the one-factor single rate model.

In sections 2.1 through 2.3 we derive the closed form solution for the traditional mean-reverting lognormal process. Then in subsequent sections we make the first steps to apply a similar analysis to the regime-switching model with randomized regimes.

The model must produce scenarios (or paths) of a single interest rate r_t across times t starting with a given value r_0 at time $t = 0$ and proceeding in equal finite increments of time that we will denote by dt in order facilitate passage to a continuous model when desired.

2.1 The traditional practitioner's mean-reverting lognormal model

In this simplest model r_t in each scenario is generated iteratively by the process

$$\begin{aligned} \ln(r_t) &= \left[\ln(r_{t-dt}) + \sigma\sqrt{dt}\mathbf{N}_t + D_t dt \right] (1-F)^{dt} \\ &\quad + \ln(T) \left[1 - (1-F)^{dt} \right] \end{aligned} \quad (2.1.1)$$

$$\begin{aligned} &= \ln(r_{t-dt}) + \left[1 - (1-F)^{dt} \right] [\ln(T) - \ln(r_{t-dt})] \\ &\quad + (1-F)^{dt} D_t dt + (1-F)^{dt} \sigma\sqrt{dt}\mathbf{N}_t \end{aligned} \quad (2.1.2)$$

where \mathbf{N}_t are independent standard normal variables for each t in each scenario and the following are parameters common across all scenarios:

- $F < 1$ is the annualized value of the mean-reversion speed parameter,
- T is the mean-reversion target parameter value for the interest rate,
- D_t is a drift compensation term at each t (specified in advance and common across all scenarios, but allowed to vary with t),
- σ is the annualized volatility parameter for the lognormal process.

For Monte Carlo simulations, the software implements (2.1.1) or (2.1.2) and the journey described in section 1 begins.

For analytic purposes, (2.1.2) allows the backward-difference expression

$$d\ln(r_t) = \left[1 - (1-F)^{dt} \right] [\ln(T) - \ln(r_{t-dt})] + (1-F)^{dt} D_t dt + (1-F)^{dt} \sigma\sqrt{dt}\mathbf{N}_t \quad (2.1.3)$$

If one wants to reproduce some specific observed annualized volatility value

$$\sigma_{obs} = \frac{std.dev. [d\ln(r_t)]}{\sqrt{dt}}$$

in the interest rate series, then it is necessary and sufficient to set the annualized volatility parameter

$$\sigma = \frac{\sigma_{obs}}{(1-F)^{dt}} \quad (2.1.4)$$

This direct calculation eliminates the need to incorporate the desired historical σ_{obs} into a trial and error routine such as described in sections (1.4) and (1.13) (item a. in each case).

In the continuous case, letting $dt \rightarrow 0$ in (2.1.3) and observing that by l'Hôpital's rule

$$\lim_{dt \rightarrow 0} \frac{\left[1 - (1-F)^{dt} \right]}{dt} = -\ln(1-F)$$

we arrive at

$$d\ln(r_t) = \{-\ln(1-F) [\ln(T) - \ln(r_t)] + D_t\} dt + \sigma d\mathbf{W}_t \quad (2.1.5)$$

where \mathbf{dW}_t is the standard Brownian motion. Equation (2.1.5) will be recognized as the Black-Karasinski stochastic interest rate model, a variant of the Black-Derman-Toy ([4], page 155).

2.2 The surprise in the drift compensation

As discussed in section 1.15 a practitioner certainly wants a real-world interest rate model to have

$$\lim_{t \rightarrow \infty} \mathbb{E}[r_t] = T \quad (2.2.1)$$

because if not then she can't understand or interpret the mean-reversion target parameter value T in real-world terms. In fact, to be sure that she understands what the mean-reversion target parameter and speed parameter values mean in the real-world, it will make the the practitioner happiest if the model admits for all t the crystal clear expression

$$\mathbb{E}[r_t] = r_0^{(1-F)^t} T^{1-(1-F)^t} \quad (2.2.2)$$

for the expected value across all scenarios for the interest rate at time t . Equation (2.2.2) allows no doubt about the real-world interpretation of the mean-reversion target and speed parameters. What does it take to guarantee that (2.2.2) will hold?

Use (2.1.1) recursively, substituting into the right hand side the expression for $\ln(r_{t-dt})$ that would result if you used $t - dt$ instead of t on the left hand side. Repeat all the way back until the only value for r that appears on the right hand side is r_0 . You will arrive at the expression (there's nothing for it but to walk through a step or two until you see the pattern)

$$\begin{aligned} \ln(r_t) &= \ln(r_0) (1-F)^{\frac{t}{dt} dt} + \sigma \sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1-F)^{sdt} \\ &\quad + \ln(T) \left[1 - (1-F)^{dt} \right] \sum_{s=1}^{\frac{t}{dt}} (1-F)^{(s-1)dt} \\ &\quad + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt} \end{aligned} \quad (2.2.3)$$

which simplifies, noting that the second \sum is a geometric series, to

$$\begin{aligned} \ln(r_t) &= \ln(r_0) (1-F)^t + \sigma \sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1-F)^{sdt} \\ &\quad + \ln(T) \left[1 - (1-F)^t \right] + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt} \end{aligned} \quad (2.2.4)$$

Exponentiating both sides

$$r_t = r_0^{(1-F)^t} T[1-(1-F)^t] e^{\sigma\sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1-F)^{sdt} + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt}} \quad (2.2.5)$$

The $\mathbf{N}_{t-(s-1)dt}$ terms in the exponential represent independent standard normal random variables and everything else in (2.2.5) is constant. The indicated linear combination of the independent standard normal random variables must be normal itself so (2.2.5) provides the closed form solution to the mean-reverting lognormal, just another more complicated lognormal whose parameters appear in (2.2.5). Since the expected value of a lognormal with $\mu = 0$ is $e^{\frac{1}{2}\sigma^2}$ (2.2.5) implies that

$$\begin{aligned} \mathbb{E}[r_t] &= r_0^{(1-F)^t} T[1-(1-F)^t] e^{\frac{1}{2}\sigma^2 dt \sum_{s=1}^{\frac{t}{dt}} (1-F)^{2sdt} + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt}} \\ &= r_0^{(1-F)^t} T[1-(1-F)^t] e^{\frac{1}{2}\sigma^2 dt (1-F)^{2dt} \frac{1-(1-F)^{2t}}{1-(1-F)^{2dt}} + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt}} \end{aligned} \quad (2.2.6)$$

This means that (2.2.2) holds if and only if

$$\sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt} = -\frac{1}{2}\sigma^2 (1-F)^{2dt} \frac{1-(1-F)^{2t}}{1-(1-F)^{2dt}} \quad (2.2.7)$$

Now we want (2.2.2) to hold for all t so (2.2.7) must hold for all t . In particular, (2.2.7) must hold when t gets replaced by $t-dt$, so rearrange (2.2.7) as follows:

$$\begin{aligned} D_t (1-F)^{dt} &= -\frac{1}{2}\sigma^2 (1-F)^{2dt} \frac{1-(1-F)^{2t}}{1-(1-F)^{2dt}} - \sum_{s=2}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt} \\ &= -\frac{1}{2}\sigma^2 (1-F)^{2dt} \frac{1-(1-F)^{2t}}{1-(1-F)^{2dt}} \\ &\quad - (1-F)^{dt} \sum_{s=2}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{(s-1)dt} \\ &= -\frac{1}{2}\sigma^2 (1-F)^{2dt} \frac{1-(1-F)^{2t}}{1-(1-F)^{2dt}} \\ &\quad - (1-F)^{dt} \sum_{s=1}^{\frac{t-dt}{dt}} D_{(t-dt)-(s-1)dt} (1-F)^{sdt} \end{aligned}$$

Now use (2.2.7) at $t - dt$ to evaluate the sum, giving

$$\begin{aligned}
D_t (1 - F)^{dt} &= -\frac{1}{2}\sigma^2 (1 - F)^{2dt} \frac{1 - (1 - F)^{2t}}{1 - (1 - F)^{2dt}} \\
&\quad + (1 - F)^{dt} \frac{1}{2}\sigma^2 (1 - F)^{2dt} \frac{1 - (1 - F)^{2(t-dt)}}{1 - (1 - F)^{2dt}} \\
&= -\frac{1}{2}\sigma^2 \frac{(1 - F)^{2dt}}{1 + (1 - F)^{dt}} \\
&\quad \times \left[\frac{1 - (1 - F)^{2t}}{1 - (1 - F)^{dt}} - (1 - F)^{dt} \frac{1 - (1 - F)^{2(t-dt)}}{1 - (1 - F)^{dt}} \right]
\end{aligned}$$

Notice that both terms represent geometric series, so

$$\begin{aligned}
D_t (1 - F)^{dt} &= -\frac{1}{2}\sigma^2 \frac{(1 - F)^{2dt}}{1 + (1 - F)^{dt}} \\
&\quad \times \left[\sum_{s=0}^{2t-dt} (1 - F)^s - (1 - F)^{dt} \sum_{s=0}^{2(t-dt)-dt} (1 - F)^s \right] \\
&= -\frac{1}{2}\sigma^2 \frac{(1 - F)^{2dt}}{1 + (1 - F)^{dt}} \left[1 + (1 - F)^{2t-dt} \right], \text{ and finally} \\
D_t &= -\frac{1}{2}\sigma^2 \frac{(1 - F)^{dt}}{1 + (1 - F)^{dt}} \left[1 + (1 - F)^{2t-dt} \right] \tag{2.2.8}
\end{aligned}$$

Equation (2.2.8) explains our bewildering problems (section 1.15) trying to come up with a constant drift compensation term by trial and error. Who knew? No constant can compensate for drift in a mean-reverting model. Equation (2.2.8) gives the correct drift compensation, varying with time, for (2.2.2) to hold in a mean-reverting model.

D_t clearly decreases in absolute value with increasing t and

$$\lim_{t \rightarrow \infty} D_t = -\frac{1}{2}\sigma^2 \frac{(1 - F)^{dt}}{1 + (1 - F)^{dt}}, \text{ provided that } F > 0, \tag{2.2.9}$$

which is the stable link between (2.2.8) and the usual lognormal value $-\frac{1}{2}\sigma^2$. In the continuous case with $dt \rightarrow 0$ in (2.2.8) we get

$$\begin{aligned}
\lim_{dt \rightarrow 0} D_t &= -\frac{1}{2}\sigma^2 \frac{1}{2} \left[1 + (1 - F)^{2t} \right] \\
&= -\frac{1}{4}\sigma^2 \left[1 + (1 - F)^{2t} \right] \text{ for each } t \\
\text{and } \lim_{t \rightarrow \infty} \lim_{dt \rightarrow 0} D_t &= -\frac{1}{4}\sigma^2, \text{ provided that } F > 0. \tag{2.2.10}
\end{aligned}$$

The limiting value $-\frac{1}{4}\sigma^2$ of the drift compensation in the continuous case is independent of the annualized mean-reversion speed $F > 0$.

When the practitioner's ideal drift compensation values (2.2.8) have been used then we know that (2.2.7) holds and its value can be substituted for the drift compensation sum in formula (2.2.5) giving a slightly simpler practitioner's form for the scenario interest rate at each time:

$$r_t = r_0^{(1-F)^t} T[1-(1-F)^t] e^{\sigma\sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1-F)^{sdt} - \frac{1}{2}\sigma^2 dt (1-F)^{2dt} \frac{1-(1-F)^{2t}}{1-(1-F)^{2dt}}} \quad (2.2.11)$$

We could have entitled this whole section "Expected value of the interest rate distribution produced by a mean-reverting lognormal process." Equation (2.2.6) gave $\mathbb{E}[r_t]$ for each t for any arbitrary specification of the sequence $\{D_s : s \leq t\}$ of drift terms. Our next section can take up the variance calculation.

Before that, to complete the expected value presentation, in the continuous case with $dt \rightarrow 0$ in (2.2.6) we get

$$\begin{aligned} \mathbb{E}[r_t] &= r_0^{(1-F)^t} T[1-(1-F)^t] e^{\frac{1}{2}\sigma^2 \int_0^t (1-F)^{2s} ds + \int_0^t D_{t-s} (1-F)^s ds} \\ &= r_0^{(1-F)^t} T[1-(1-F)^t] e^{-\frac{1}{4} \frac{\sigma^2}{\ln(1-F)} [1-(1-F)^{2t}] + \int_0^t D_{t-s} (1-F)^s ds} \end{aligned} \quad (2.2.12)$$

The value for $\mathbb{E}[r_t]$ in (2.2.6) - for the discrete process - and (2.2.12) - for the continuous process - converges respectively to

$$\lim_{t \rightarrow \infty} \mathbb{E}[r_t] = T e^{\left(\frac{1}{2}\sigma^2 dt \frac{(1-F)^{2dt}}{1-(1-F)^{2dt}} + dt \frac{(1-F)^{dt}}{1-(1-F)^{dt}} \lim_{t \rightarrow \infty} D_t \right)} \quad (2.2.13)$$

$$\lim_{t \rightarrow \infty} \mathbb{E}[r_t] = T e^{-\left(\frac{1}{4}\sigma^2 + \lim_{t \rightarrow \infty} D_t \right) \frac{1}{\ln(1-F)}} \quad (2.2.14)$$

provided that $F > 0$ and that D_t has a limit as $t \rightarrow \infty$. Actually a slightly weaker condition on D_t would work, namely that as t increases D_t should increasingly sum or integrate as if it had a limit. At any rate, the condition for the weak *desideratum* (2.2.1) to hold is that (2.2.9), respectively (2.2.10), hold or at least increasingly behave in sums or integrals as if they hold.

2.3 Variance of the interest rate distribution produced by a mean-reverting lognormal process

Since they contain just constants and independent standard normal random variables, (2.2.5) or (2.2.11) allow us to calculate at each t the variance $\mathbb{V}[r_t]$ across scenarios of the random interest rates r_t the same way it allowed us to calculate expected values. Recall that a lognormal random variable with $\mu = 0$

has variance equal to $e^{\sigma^2} (e^{\sigma^2} - 1)$. Then (2.2.5) gives

$$\begin{aligned}
\mathbb{V}[r_t] &= \left(r_0^{2(1-F)^t} T^2 [1-(1-F)^t] e^{\sigma^2 dt \sum_{s=1}^{\frac{t}{dt}} (1-F)^{2sdt} + 2dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt}} \right) \\
&\quad \times \left(e^{\sigma^2 dt \sum_{s=1}^{\frac{t}{dt}} (1-F)^{2sdt}} - 1 \right) \\
&= \left(r_0^{2(1-F)^t} T^2 [1-(1-F)^t] e^{\sigma^2 dt (1-F)^{2dt} \frac{1-(1-F)^{2t}}{1-(1-F)^{2dt}} + 2dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt}} \right) \\
&\quad \times \left(e^{\sigma^2 dt (1-F)^{2dt} \frac{1-(1-F)^{2t}}{1-(1-F)^{2dt}}} - 1 \right) \tag{2.3.1}
\end{aligned}$$

In cases when the practitioner's ideal drift compensation values (2.2.8) have been used then we can use (2.2.11) for the variance calculation, or just refer to (2.2.7), which is equivalent to (2.2.8), to simplify (2.3.1). Either way the result is the practitioner's formula for the variance across scenarios for the interest rate at time t .

$$\mathbb{V}[r_t] = r_0^{2(1-F)^t} T^2 [1-(1-F)^t] \left(e^{\sigma^2 dt (1-F)^{2dt} \frac{1-(1-F)^{2t}}{1-(1-F)^{2dt}}} - 1 \right) \tag{2.3.2}$$

The limit of this as $t \rightarrow \infty$ is

$$\lim_{t \rightarrow \infty} \mathbb{V}[r_t] = T^2 \left(e^{\sigma^2 dt \frac{(1-F)^{2dt}}{1-(1-F)^{2dt}}} - 1 \right), \text{ provided that } F > 0, \tag{2.3.3}$$

knowledge of which would have saved endless trial and error model runs, as described in section 1.4, looking for joint values of F and σ to reproduce in the limit the historical variance of the interest rate. In fact, if the historically observed variance of the interest rate is σ_{obs} , then the value $\sigma = \frac{\sigma_{obs}}{(1-F)^{dt}}$ from (2.1.4) can be substituted into (2.3.3), along with the requirement that $\lim_{t \rightarrow \infty} \mathbb{V}[r_t] = \sigma_{obs}^2$, to give

$$\sigma_{obs}^2 = T^2 \left(e^{\sigma_{obs}^2 \frac{dt}{1-(1-F)^{2dt}}} - 1 \right), \text{ provided that } F > 0. \tag{2.3.4}$$

This can be solved for the value of the annualized mean-reversion parameter value F that will reproduce σ_{obs}^2 ,

$$F = 1 - \left\{ 1 - \frac{\sigma_{obs}^2 dt}{\ln \left[\left(\frac{\sigma_{obs}}{T} \right)^2 + 1 \right]} \right\}^{\frac{1}{2dt}} \tag{2.3.5}$$

with no trial and error required. For the continuous model $dt \rightarrow 0$, inaccessible to trial and error anyway, the identity $\lim_{\epsilon \rightarrow 0} (1 + a\epsilon)^{1/\epsilon} = e^a$ applied to (2.3.5) produces

$$F = 1 - e^{-\frac{1}{2}\sigma_{obs}^2 \frac{1}{\ln\left[\left(\frac{\sigma_{obs}}{T}\right)^2 + 1\right]}} \quad (2.3.6)$$

for the annualized mean-reversion parameter value F to reproduce σ_{obs}^2 .

To complete the general story, for the continuous model $dt \rightarrow 0$ with arbitrary values for the drift function D_t , (2.3.1) becomes

$$\begin{aligned} \mathbb{V}[r_t] &= \left(r_0^{2(1-F)^t} T^2 [1 - (1-F)^t] e^{-\frac{1}{2} \frac{\sigma^2}{\ln(1-F)} [1 - (1-F)^{2t}] + 2 \int_0^t D_{t-s} (1-F)^s ds} \right) \\ &\quad \times \left(e^{-\frac{1}{2} \frac{\sigma^2}{\ln(1-F)} [1 - (1-F)^{2t}] - 1} \right) \end{aligned} \quad (2.3.7)$$

in a calculation exactly parallel to (2.2.12). When $t \rightarrow \infty$ in this continuous case, with arbitrary D_t ,

$$\lim_{t \rightarrow \infty} \mathbb{V}[r_t] = T^2 e^{-\left(\frac{1}{2}\sigma^2 + 2 \lim_{t \rightarrow \infty} D_t\right) \frac{1}{\ln(1-F)}} \left(e^{-\frac{1}{2}\sigma^2 \frac{1}{\ln(1-F)}} - 1 \right) \quad (2.3.8)$$

provided that $F > 0$ and similar conditions on D_t hold as discussed in connection with (2.2.14). When $t \rightarrow \infty$ the general discrete case (2.3.1) with arbitrary drift values D_t becomes

$$\lim_{t \rightarrow \infty} \mathbb{V}[r_t] = T^2 e^{\left(\frac{1}{2}\sigma^2 dt \frac{(1-F)^{2dt}}{1-(1-F)^{2dt}} + 2dt \frac{(1-F)^{dt}}{1-(1-F)^{dt}} \lim_{t \rightarrow \infty} D_t\right)} \left(e^{\frac{1}{2}\sigma^2 dt \frac{(1-F)^{2dt}}{1-(1-F)^{2dt}}} - 1 \right) \quad (2.3.9)$$

provided that $F > 0$ and similar conditions on D_t hold as discussed in connection with (2.2.13).

This completes the task of making the traditional mean-reverting lognormal a completely open book. Formulas (2.2.5) or (2.2.11), whichever applies, can give any moment you like for the real-world interest rate r_t across scenarios at any time t . They will simplify (if that's the word) by geometric series just as did the formulas for mean and variance.

Sections 2.1, 2.2 and 2.3 now serve both as a basis and as an ideal for developing the regime-switching version of the mean-reverting lognormal.

2.4 Regime-Switching

As described in section 1.8 and 1.11 we introduced multiple possible values for the mean-reversion target parameter T . Let T_t indicate the value of the mean-reversion target parameter at time t . Then the backward difference expression (2.1.3) for $\ln(r_t)$ becomes

$$\begin{aligned} d \ln(r_t) &= \left[1 - (1-F)^{dt} \right] [\ln(T_t) - \ln(r_{t-dt})] \\ &\quad + (1-F)^{dt} D_t dt + (1-F)^{dt} \sigma \sqrt{dt} \mathbf{N}_t \end{aligned} \quad (2.4.1)$$

and the continuous version (2.1.5), now a generalized Black-Karasinski, becomes

$$d \ln(r_t) = \{-\ln(1-F)[\ln(T_t) - \ln(r_t)] + D_t\} dt + \sigma d\mathbf{W}_t \quad (2.4.2)$$

Applying (2.4.1) recursively all the way back to $\ln(r_0)$ gives the regime-switching expression corresponding to (2.2.3):

$$\begin{aligned} \ln(r_t) &= \ln(r_0)(1-F)^t + \sigma \sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1-F)^{sdt} \\ &\quad + \left[1 - (1-F)^{dt}\right] \sum_{s=1}^{\frac{t}{dt}} \ln(T_{t-(s-1)dt}) (1-F)^{(s-1)dt} \\ &\quad + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt} \end{aligned} \quad (2.4.3)$$

Before randomizing the mean-reversion target parameter values T_t it will prove convenient to change indexes in the second sum in (2.4.3) so that the target parameter values appear in forward time sequence:

$$\begin{aligned} \ln(r_t) &= \ln(r_0)(1-F)^t + \sigma \sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1-F)^{sdt} \\ &\quad + \left[1 - (1-F)^{dt}\right] \sum_{s=1}^{\frac{t}{dt}} \ln(T_{sdt}) (1-F)^{t-sdt} \\ &\quad + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt} \end{aligned} \quad (2.4.4)$$

It also will be convenient to let $\{t_j\}$ be the (time-ordered) countable sequence of times when the value of T_t changes and $\{T_j\}$ be the corresponding sequence of values for T_t subsequent to each change. Then T_t is a function of t defined by

$$T_t = T_j \text{ for all } t_j \leq t < t_{j+1} \quad (2.4.5)$$

where we take $t_0 = 0$ so that T_0 unambiguously represents the initial value for the mean-reversion target parameter. Then we can write

$$T_t = \sum_{j=0}^{\infty} \mathbf{1}_{[j, j+1)}(t) T_j \quad (2.4.6)$$

where $\mathbf{1}_{[j, j+1)}(t)$ is the *cádlág* indicator function for the half-open interval $t_j \leq t < t_{j+1}$, equal to 1 when $t_j \leq t < t_{j+1}$ and equal to 0 otherwise. The functional dependence of T_t on t is expressed through these indicator functions. Definition (2.4.5) also leads to the unusual-appearing (relative to 2.4.6) expression

$$\ln(T_t) = \sum_{j=0}^{\infty} \mathbf{1}_{[j,j+1)}(t) \ln(T_j) \quad (2.4.7)$$

We can substitute the value for $\ln(T_t)$ given by (2.4.7) into equations (2.4.1), (2.4.2), and (2.4.4), respectively, to get for the discrete case

$$\begin{aligned} d \ln(r_t) = & \left[1 - (1 - F)^{dt} \right] \left[\sum_{j=0}^{\infty} \mathbf{1}_{[j,j+1)}(t) \ln(T_j) - \ln(r_{t-dt}) \right] \\ & + (1 - F)^{dt} D_t dt + (1 - F)^{dt} \sigma \sqrt{dt} \mathbf{N}_t, \end{aligned} \quad (2.4.8)$$

for the continuous case

$$d \ln(r_t) = \left\{ -\ln(1 - F) \left[\sum_{j=0}^{\infty} \mathbf{1}_{[j,j+1)}(t) \ln(T_j) - \ln(r_t) \right] + D_t \right\} dt + \sigma d\mathbf{W}_t, \quad (2.4.9)$$

and for the recursive resolution

$$\begin{aligned} \ln(r_t) = & \ln(r_0) (1 - F)^t + \sigma \sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1 - F)^{sdt} \\ & + \left[1 - (1 - F)^{dt} \right] \sum_{s=1}^{\frac{t}{dt}} \sum_{j=0}^{\infty} \mathbf{1}_{[j,j+1)}(sdt) \ln(T_j) (1 - F)^{t-sdt} \\ & + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1 - F)^{sdt}. \end{aligned} \quad (2.4.10)$$

Changing the order of summation in the double sum in (2.4.10) gives

$$\begin{aligned} \ln(r_t) = & \ln(r_0) (1 - F)^t + \sigma \sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1 - F)^{sdt} \\ & + \left[1 - (1 - F)^{dt} \right] \sum_{j=0}^{\infty} \ln(T_j) \sum_{s=1}^{\frac{t}{dt}} \mathbf{1}_{[j,j+1)}(sdt) (1 - F)^{t-sdt} \\ & + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1 - F)^{sdt} \end{aligned} \quad (2.4.11)$$

The indicator functions (equal to just 0 or 1) allow a segmentation of the inside sum:

$$\begin{aligned}
\ln(r_t) &= \ln(r_0) (1 - F)^t + \sigma \sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1 - F)^{sdt} \\
&+ \left[1 - (1 - F)^{dt} \right] \sum_{j=0}^{\infty} \ln(T_j) \sum_{s=\lfloor (t \wedge t_j) \frac{1}{dt} \rfloor + 1}^{\lfloor (t \wedge t_{j+1}) \frac{1}{dt} \rfloor} (1 - F)^{t-sdt} \\
&+ dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1 - F)^{sdt}
\end{aligned} \tag{2.4.12}$$

where $\lfloor u \rfloor$ is defined to be the largest integer with $\lfloor u \rfloor < u$ (note the strict inequality, so for an integer $\lfloor n \rfloor = n - 1$; *cádlág* in *cádlág* out). Now bring the factor $\left[1 - (1 - F)^{dt} \right]$ inside the double sum. That makes the inner sums telescope down to just differences of the top and bottom terms

$$\begin{aligned}
\ln(r_t) &= \ln(r_0) (1 - F)^t + \sigma \sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1 - F)^{sdt} \\
&+ \sum_{j=0}^{\infty} \ln(T_j) \left[(1 - F)^{t - \lfloor (t \wedge t_{j+1}) \frac{1}{dt} \rfloor dt} - (1 - F)^{t - \lfloor (t \wedge t_j) \frac{1}{dt} \rfloor dt} \right] \\
&+ dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1 - F)^{sdt}
\end{aligned} \tag{2.4.13}$$

Equation (2.4.13) is the analogue of (2.2.4) in the simple mean-reverting lognormal. The telescoping sums are the closest we can get here to the geometric series that simplified (2.2.4). Exponentiating it, which isn't worth the ink, would give the analogue of (2.2.5), namely a closed form expression for a deterministic regime-switching mean-reverting lognormal. The stage is set to randomize the regimes.

2.5 Randomizing the regimes

In sections 1.08, 1.10 and 1.11 we discussed the practitioner's modeling decision to randomize the mean-reversion target parameters as independent and identically distributed lognormals, also independent of the independent standard normals $\mathbf{N}_{t-(s-1)dt}$. Model this with independent normal random variables $\ln(\mathbf{T}_j)$ with common mean μ_T and variance σ_T^2 .

Still letting the regime-switching points in time $\{t_j\}$ be deterministic (in-

cluding when we are conditioning on random $\{t_j\}$ equation (2.4.13) becomes

$$\begin{aligned}
\ln(r_t) &= \ln(r_0)(1-F)^t + \sigma\sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1-F)^{sdt} \\
&+ \ln(T_0) \left[(1-F)^{t-\lfloor (t \wedge t_1) \frac{1}{dt} \rfloor dt} - (1-F)^t \right] \\
&+ \sum_{j=1}^{\infty} \ln(\mathbf{T}_j) \left[(1-F)^{t-\lfloor (t \wedge t_{j+1}) \frac{1}{dt} \rfloor dt} - (1-F)^{t-\lfloor (t \wedge t_j) \frac{1}{dt} \rfloor dt} \right] \\
&+ dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt} \tag{2.5.1}
\end{aligned}$$

(remember that T_0 is a fixed initial value for the mean-reversion target parameter, not a random variable.). Then $\ln(r_t)$ is a normal random variable at each point in time t with expected value

$$\begin{aligned}
\mathbb{E}[\ln(r_t)] &= \ln(r_0)(1-F)^t + \ln(T_0) \left[(1-F)^{t-\lfloor (t \wedge t_1) \frac{1}{dt} \rfloor dt} - (1-F)^t \right] \\
&+ \mu_T \sum_{j=1}^{\infty} \left[(1-F)^{t-\lfloor (t \wedge t_{j+1}) \frac{1}{dt} \rfloor dt} - (1-F)^{t-\lfloor (t \wedge t_j) \frac{1}{dt} \rfloor dt} \right] \\
&+ dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt} \\
&= \ln(r_0)(1-F)^t + \ln(T_0) \left[(1-F)^{t-\lfloor (t \wedge t_1) \frac{1}{dt} \rfloor dt} - (1-F)^t \right] \\
&\mu_T \left[1 - (1-F)^{t-\lfloor (t \wedge t_1) \frac{1}{dt} \rfloor dt} \right] \\
&+ dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1-F)^{sdt} \tag{2.5.2}
\end{aligned}$$

where the simplification of the μ_T term comes from the facts that the sum telescopes and that for some j large enough $t_j \geq t$ and all the terms after that are just $1 - 1 = 0$.

For the variance

$$\begin{aligned}
\mathbb{V}[\ln(r_t)] &= \sigma^2 dt \sum_{s=1}^{\frac{t}{dt}} (1-F)^{2sdt} \\
&\quad + \sigma_T^2 \sum_{j=1}^{\infty} \left[(1-F)^{t - \lfloor (t \wedge t_{j+1}) \frac{1}{dt} \rfloor dt} - (1-F)^{t - \lfloor (t \wedge t_j) \frac{1}{dt} \rfloor dt} \right]^2 \\
&= \sigma^2 dt (1-F)^{2dt} \frac{1 - (1-F)^{2t}}{1 - (1-F)^{2dt}} \\
&\quad + \sigma_T^2 \sum_{j=1}^{\infty} \left[(1-F)^{t - \lfloor (t \wedge t_{j+1}) \frac{1}{dt} \rfloor dt} - (1-F)^{t - \lfloor (t \wedge t_j) \frac{1}{dt} \rfloor dt} \right]^2
\end{aligned} \tag{2.5.3}$$

That geometric series marks the last exact simplification we'll see.

Of course, with (2.5.3) the mean and variance of r_t at each point in time t now can be expressed as

$$\mathbb{E}[r_t] = e^{\mathbb{E}[\ln(r_t)] + \frac{1}{2}\mathbb{V}[\ln(r_t)]} \tag{2.5.4}$$

$$\mathbb{V}[r_t] = e^{2\mathbb{E}[\ln(r_t)] + \mathbb{V}[\ln(r_t)]} \left(e^{\mathbb{V}[\ln(r_t)]} - 1 \right) \tag{2.5.5}$$

using the expressions from (2.5.2) and (2.5.3) to complete the formulae.

So far so good. Now what about the switching times?

2.6 Randomizing the switching times

In sections 1.08, 1.09, 1.11, and 1.12 we discussed the practitioner's modeling decision to randomize the mean-reversion target parameter switching times \mathbf{t}_j as sums of independent and identically distributed Erlang interarrival times (i.e. the $\mathbf{t}_{j+1} - \mathbf{t}_j$ are the Erlangs) also independent of the independent standard normals $\mathbf{N}_{t-(s-1)dt}$ and the independent lognormal mean-reversion target parameter values \mathbf{T}_j . (Except that the very first switching time $\mathbf{t}_1 = \mathbf{t}_1 - 0$ was taken as an independently random point chosen within the first random Erlang span, see section 1.12).

Substituted into (2.4.8) this gives the backward difference expression for the discrete case

$$\begin{aligned}
d \ln(r_t) &= \left[1 - (1-F)^{dt} \right] \left[\sum_{j=0}^{\infty} \mathbf{1}_{[\mathbf{j}, \mathbf{j}+1)}(t) \ln(\mathbf{T}_j) - \ln(r_{t-dt}) \right] \\
&\quad + (1-F)^{dt} D_t dt + (1-F)^{dt} \sigma \sqrt{dt} \mathbf{N}_t,
\end{aligned} \tag{2.6.1}$$

where $\ln(\mathbf{T}_j)$ reflect the random mean-reversion target parameter values defined in section 2.5 and the indicator functions $\mathbf{1}_{[\mathbf{j}, \mathbf{j}+1)}(t)$ defined deterministically for (2.4.6) now are random variables at each t defined by

$$\mathbf{1}_{[\mathbf{j}, \mathbf{j}+1)}(t) = 1 \text{ if } \mathbf{t}_j \leq t < \mathbf{t}_{j+1} \text{ and } 0 \text{ otherwise.}$$

where \mathbf{t}_j are the random switching times defined in the preceding paragraph.

For Monte Carlo simulations, the software implements (2.6.1) and the randomized regime switching journey takes off, as described starting at section 1.11.

For analytic purposes, the recursive resolution developed in (2.4.13) or (2.5.1), now with random switching times inserted, gives

$$\begin{aligned}
\ln(r_t) &= \ln(r_0) (1 - F)^t + \sigma \sqrt{dt} \sum_{s=1}^{\frac{t}{dt}} \mathbf{N}_{t-(s-1)dt} (1 - F)^{sdt} \\
&+ \ln(T_0) \left[(1 - F)^{t - \lfloor (t \wedge \mathbf{t}_1) \frac{1}{dt} \rfloor dt} - (1 - F)^t \right] \\
&+ \sum_{j=1}^{\infty} \ln(\mathbf{T}_j) \left[(1 - F)^{t - \lfloor (t \wedge \mathbf{t}_{j+1}) \frac{1}{dt} \rfloor dt} - (1 - F)^{t - \lfloor (t \wedge \mathbf{t}_j) \frac{1}{dt} \rfloor dt} \right] \\
&+ dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1 - F)^{sdt} \tag{2.6.2}
\end{aligned}$$

where now the \mathbf{t}_j for $j > 0$ are the random variables just described. Forbidding.

But we can get a somewhat simplified expression for the expected value $\mathbb{E}[\ln(r_t)]$ at each t by conditioning on $\{\mathbf{t}_j\}$ and using (2.5.2)

$$\begin{aligned}
\mathbb{E}[\ln(r_t)] &= \ln(r_0) (1 - F)^t + \ln(T_0) \left\{ \left[\mathbb{E} \left[(1 - F)^{t - \lfloor (t \wedge \mathbf{t}_1) \frac{1}{dt} \rfloor dt} \right] - (1 - F)^t \right] \right\} \\
&\quad \mu_T \left\{ \left[1 - \mathbb{E} \left[(1 - F)^{t - \lfloor (t \wedge \mathbf{t}_1) \frac{1}{dt} \rfloor dt} \right] \right] \right\} \\
&\quad + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1 - F)^{sdt} \\
&= \ln(r_0) (1 - F)^t + \mu_T - \ln(T_0) (1 - F)^t \\
&\quad + \mathbb{E} \left[(1 - F)^{t - \lfloor (t \wedge \mathbf{t}_1) \frac{1}{dt} \rfloor dt} \right] [\ln(T_0) - \mu_T] \\
&\quad + dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1 - F)^{sdt} \tag{2.6.3}
\end{aligned}$$

which thankfully involves only \mathbf{t}_1 , although even that in a complicated way. The expectation $\mathbb{E} \left[(1 - F)^{t - \lfloor (t \wedge \mathbf{t}_1) \frac{1}{dt} \rfloor dt} \right]$ in fact represents at each t just a constant $(1 - F)^t$ times the Laplace transform of the random variable $\lfloor (t \wedge \mathbf{t}_1) \frac{1}{dt} \rfloor$ evaluated at $dt \ln(1 - F)$. For $t \rightarrow \infty$ and $F > 0$ it works out well so long as D_t has a limiting value and the Laplace transform of the random variable $\lfloor \frac{\mathbf{t}_1}{dt} \rfloor$ is

finite at $dt \ln(1 - F)$.

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}[\ln(r_t)] &= \mu_T + \lim_{t \rightarrow \infty} dt \sum_{s=1}^{\frac{t}{dt}} D_{t-(s-1)dt} (1 - F)^{sdt} \\ &= \mu_T + dt \frac{(1 - F)^{dt}}{1 - (1 - F)^{dt}} \lim_{t \rightarrow \infty} D_t \end{aligned} \quad (2.6.4)$$

where the limit of the sum is the same thing we saw at (2.2.13) versus (2.2.6) and where $\lim_{t \rightarrow \infty} D_t$ may well be known, for example from (2.2.9) or (2.2.10) if the "practitioner's" drift compensation (2.2.8) has been chosen, or from whatever other choice may have been made. The same discussion about D_t applies as followed (2.2.13) and (2.2.14). The continuous version, comparable to (2.2.14), is available, too. As $dt \rightarrow 0$ (2.6.4) turns into

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}[\ln(r_t)] &= \mu_T + \lim_{t \rightarrow \infty} \int_0^t D_{t-s} (1 - F)^s ds \\ &= \mu_T - \lim_{t \rightarrow \infty} D_t \frac{1}{\ln(1 - F)} \end{aligned} \quad (2.6.5)$$

if D_t has a limiting value.

So, at least the limiting values of $\mathbb{E}[\ln(r_t)]$ can be calculated and they are what you would hope (after going through the simple mean-reverting lognormal material, anyway). Unfortunately, to get $\mathbb{E}[r_t]$, which the practitioner can understand a lot better than $\mathbb{E}[\ln(r_t)]$, requires (even for its limiting value) an expression for the variance $\mathbb{V}[\ln(r_t)]$ of its logarithm, see (2.5.5). And conditioning on $\{\mathbf{t}_j\}$ is not nearly so helpful for the variance.

Conditioning on $\{\mathbf{t}_j\}$ and using (2.5.3)

$$\begin{aligned} \mathbb{V}[\ln(r_t)] &= \sigma^2 dt (1 - F)^{2dt} \frac{1 - (1 - F)^{2t}}{1 - (1 - F)^{2dt}} \\ &\quad + \sigma_T^2 \sum_{j=1}^{\infty} \mathbb{E} \left[\left\{ (1 - F)^{t - \lfloor (t \wedge \mathbf{t}_{j+1}) \frac{1}{dt} \rfloor} - (1 - F)^{t - \lfloor (t \wedge \mathbf{t}_j) \frac{1}{dt} \rfloor} \right\}^2 \right], \end{aligned} \quad (2.6.6)$$

which so far defies closed form analysis, even in the limit at $t \rightarrow \infty$. The problem is that the random variables

$$\left\{ (1 - F)^{t - \lfloor (t \wedge \mathbf{t}_{j+1}) \frac{1}{dt} \rfloor} - (1 - F)^{t - \lfloor (t \wedge \mathbf{t}_j) \frac{1}{dt} \rfloor} \right\} \quad (2.6.7)$$

are not independent. Only the $\mathbf{t}_{j+1} - \mathbf{t}_j$ are independent. Even for the calculation of limiting values the random variables

$$\left\{ (1 - F)^{-\lfloor \mathbf{t}_{j+1} \frac{1}{dt} \rfloor} - (1 - F)^{-\lfloor \mathbf{t}_j \frac{1}{dt} \rfloor} \right\} \quad (2.6.8)$$

are not independent. So nothing telescopes, no geometric series.

If we could cope with $\mathbb{V}[\ln(r_t)]$ then (2.5.4), (2.5.5), and their extensions to higher moments would give us all the moments of r_t for each t just as we had for the simple mean-reverting lognormal. Unfortunately, progress on exact analysis so far dies at (2.6.6).

Many years worth of simulation trials using (2.6.1) give confidence that the perturbations coming from the failure of the random variables (2.6.7) and (2.6.8) to be independent are small and controllable in practice. Having come so far, it is a shame not to be able to conclude with a victory lap of closed form solutions. And even practical "control" of the perturbations leaves a second order (much smaller) version of the vexation described in section 1.15.

3 Illustrations

Organization of some of the output from those many years of simulation trial results so as illustrate both the narrative of section 1 and the mathematics of section 2 will come in the next version of this paper.

4 References

1. Becker, D. N. "Some Observations on U. S. Treasury Interest Rates: 1953-1988" Risks and Rewards Newsletter of the Investment Section of the Society of Actuaries, No. 24 (1995)
2. Bridgeman, J. G. "Session 29IF, Practitioner's Forum", 2001 Valuation Actuary Symposium Proceedings (November 29-30, 2001), Society of Actuaries (2002)
3. Hardy, M. R., "A Regime-Switching Model of Long-Term Stock Returns", North American Actuarial Journal (NAAJ), Vol. 5, No. 2 (April 2001)
4. Ho, T. S. and Lee, S. B., The Oxford Guide To Financial Modeling, Oxford University Press (2004)