Pricing Long Bonds: Pitfalls and Opportunities

Philip H. Dybvig
and
William J. Marshall*

First draft: November, 1994.
This draft: November, 1995.

Abstract

Recent issues such as Disney's 100-year bonds have made it especially important to value long-maturity fixed claims accurately,[1] and valuing long claims has always been important for immunizing portfolios used to fund long obligations such as certain defined-benefit pension payments. Due to the subtleties of convexity, parameter uncertainty, and the impact of default risk, it is easy to underestimate the value of long-maturity bonds. This article examines the pitfalls in pricing long bonds and the correct pricing that points to investment opportunities. Formal analysis of these problems reveals large errors that can arise from the usual intuitive arguments using the expectations hypothesis and point estimates of unknown parameters.

*Philip H. Dybvig is Boatmens Bancshares Professor of Banking and Finance at Washington University in Saint Louis. William J. Marshall is the Chief Operating Officer at NISA Investment Advisors, LLC.

[1]See, for example, Fred Vogelstein, "The 100-year Bond is Coming Back, But Is It Good?" *Wall Street Journal*, November 22, 1995.

Introduction

Pricing long-maturity bonds is an important part of managing fixed-income portfolios, and it is also important for the determination of how to price and hedge a fixed liability, for example in defined-benefit pension programs or the actuarially expected payment on a pool of insurance policies. Compared with pricing short bonds, pricing long bonds is a particular challenge. Not only are there fewer comparable bonds at longer maturities, but also the analysis is more complex, given that simple rules of thumb such as the expectations hypothesis may be far off. The flip side of this challenge is that there may be profit opportunities available for investors who can do a better job of pricing these bonds than the rest of the market. The purpose of this paper is to use recent theoretical advances to shed light on the practical problem of pricing long bonds.

Our analysis identifies several possible pitfalls in evaluating long bonds. Interestingly, in each case, the natural error would cause us to undervalue the long bond or equivalently to expect yields or forward rates that are too high. It is our perception that long bonds are probably valued too little by the market, and that this represents a potential trading opportunity for investors, at a scale appropriate for their tolerance of risk taking.

Arbitrage arguments are at the core of modern asset pricing theory, and our analysis of the first pitfall is an application of simple arbitrage arguments. The first pitfall is overestimating the importance of later forward rates on coupon bond prices. We consider the pricing of a 100-year bond by noting it has the same cash flows as holding a 50-year bond and making a forward commitment to buy a 50-year bond 50 years from now. Therefore, absence of arbitrage implies that the yields on the three bonds are linked. We find that if the yield on the 100-year bond is even as little as 25 basis points higher than the yield on the 50-year bond, then the implied yield on the forward commitment is unreasonably high. These arguments suggest that increasing bond yields at long maturities may not be justified, even if the forward rates are somewhat higher than current rates.

A second possible pitfall is overstating the importance of default risk. Intu-

itively, it may seem that a 100-year bond is much more risky than a 50-year bond. However, the extension to 100 years is *conditional* on the bond making through the first 50 years, since the two bonds are identical if the bond defaults before 50 years. Since the expected creditworthiness of the firm after 50 years given no default before then is probably not significantly less than the initial creditworthiness, no special default risk premium is required for the extension. In the extreme case in which the firm is sure to be dissolved before 50 years, there is no difference between a 50-year bond and a 100-year bond, and they should have exactly the same yield. In fact, in an intermediate case in which a firm that survives 50 years will have better prospects than it has initially, the required default premium on a 100-year bond might even be lower than on a 50-year bond.

The analysis of the remaining pitfalls requires more machinery, and is based on valuation using the *Local Expectations Hypothesis* (LEH) championed by Cox, Ingersoll, and Ross.[2] The third pitfall is the bias introduced by using the expectations hypothesis. This would be a good approximation at short maturities, but results in discount and forward rates that are much too high at longer maturities.

The fourth pitfall is a subtle variation of the third pitfall. If we are uncertain about the long-run mean interest rate, using the mean of the distribution instead of the uncertain distribution tends to overstate long forward rates by a large amount. This contrasts with the usual case, common in practice, in which we take some reasonable estimate and substitute it in as if it were the truth: in many contexts this approach is fine.[3] In a specific example we analyze, there is a 90% probability the long-term mean interest rate is around 10%, and only a 10% probability it is around 1%. In that example, the yield on a 100-year zero-coupon bond is around 3.6%, which is much less than the expected long-term mean of around 9%.

[2]See Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross (1981) "A Reexamination of Traditional Hypotheses about the Term Structure of Interest Rates" *Journal of Finance* **36**, 769–799. Further mention of Cox, Ingersoll, and Ross refers to this paper.

[3]However, another interesting case in investments in which this is definitely a bad procedure is in using simple sample statistics in mean-variance analysis, which would direct us to take extreme positions based mostly on the luck of historical draws.

The analysis of the various pitfalls is augmented by some theoretical results, some of which are similar to what is already in the literature.[4] The theoretical analysis in this paper is bolstered by numerical examples based on a class of models with reasonable assumptions about interest rate dynamics.

Some analytical results are developed in an Appendix, with more intuitive accounts in the text. The Appendix also contains a detailed description of the examples used in the text. We are grateful to money managers who do not learn about quantitative term structure models like the ones in the Appendix, since those managers are the ones who will be providing our clients with profit opportunities!

Making Long Bonds out of Shorter Bonds and Shorter Bond Commitments

One useful device for thinking about pricing long bonds is to think about an arbitrage of replicating the long bond by a shorter bond and one or more forward commitments to buy bonds in the future. For example, if we are trying to price a new riskless par 100-year bond and there is an existing riskless par 50-year bond, the cash flows for the 100-year bond are identical to the cash flows to a 50-year bond with the same yield and face and a forward commitment to buy an identical 50-year bond, 50 years from now: the coupon payments are identical in the two cases and the repayment of principal on the first 50-year bond is just sufficient to finance the purchase of the new 50-year bond. In this way, fairness of the pricing of the 100-year bond can be examined as fairness of implicit pricing of the forward commitment to buy the bond 50 years from now, given that the first 50-year bond is priced fairly. Obviously, we could also think about replicating the 100-year bond by purchase of four 25-year bonds, one now and three forward purchases in 25, 50, and 75 years.

[4]See especially Cox, Ingersoll, and Ross, and Dybvig, Philip H., Jonathan E. Ingersoll, Jr., and Stephen A. Ross "Long Forward Rates Can Never Fall," *Journal of Business*, forthcoming. Further mention of Dybvig, Ingersoll, and Ross refers to this paper.

If, as is usually the case, the 100-year bond is not priced to have the same yield as the 50-year bond, some adjustment must be made. The simplest adjustment, which is normally adequate, replaces either the par 50-year bond or the par 100-year bond with a premium or discount bond that equates yield on initial purchase price. For example, suppose there exist simultaneously 100-year and 50-year par riskless bonds, with respective annual coupon rates of 5.5% and 5.25%, face values both \$1000, and therefore semi-annual coupons of \$27.5 and \$26.25, respectively. To modify the 50-year bond to be more directly comparable with the 100-year bond, we change its coupon to be the same as for the 100-year bond and we adjust its final face value to keep its initial price and yield unchanged. The formal calculation per \$1000 of original face value is that the new face value F must satisfy

$$(1) \quad \$1000 = \frac{1}{.0525/2}\left(1 - \frac{1}{(1+.0525/2)^{2\times 50}}\right)\$27.50 + \frac{1}{(1+.0525/2)^{2\times 50}}F,$$

which involves only ordinary bond calculations. The new face value is \$412.1451. This might seem surprisingly low, but most of the value of the 50-year bond is in early coupon payments, not in repayment of principal, so a big reduction of the final repayment of face value is required to offset the increase in coupon payments. The only real assumption we have made to do this calculation is that the yield would be the same for the premium bond, which should be a good approximation. Alternatively, if we have some idea about the shape of the yield curve (or better yet the yield on a 50-year discount bond), we could increase the precision of our analysis.

Thus, buying the 100-year bond is the equivalent of buying a premium 50-year bond similar to what is marketed plus a forward commitment to buy a discount 50-year bond 50-years from now. The forward commitment is to buy a bond at a price of \$412.1451, with a semi-annual coupon of \$27.50 and a face value at maturity of \$1000. Such a bond has a yield of around 13.37%, which is probably unreasonably large for reasons we discuss below, and therefore the purchase is a good deal. Although it is not necessarily the case that this long bond will outperform short investments in the long haul, it does seem to be less expensive for its amount of interest rate exposure than other bonds in the market (in the example). Therefore, its purchase would

be recommended in concert with other positions that create the appropriate overall risk exposure in the portfolio.

Default risk

Perhaps surprisingly, this analysis is still reasonable with little modification in the presence of default risk. If the 50-year bond and the 100-year bond have the same face value and yield, it is probably reasonable to assume they will be treated similarly in bankruptcy. In that case, purchasing the 100-year bond is equivalent to purchasing the 50-year bond today plus a forward commitment to purchase the other 50-year bond, 50 years from now, *but only in the contingency that the first bond did not disappear because of a bankruptcy or other event.* If the yield on the 50-year and 100-year bonds differ, then it may makes sense to try to do some adjustment for what happens in bankruptcy, although such an adjustment will not be very large given that bond defaults tend to be partial and do not happen very often empirically.

Our conclusion from this analysis is that since it is the creditworthiness of the bond 50 years hence *conditional on its survival to that point*, it is probably incorrect to assert that a 100 year bond is intrinsically riskier than a 50 year bond. With rare exceptions, we probably do not know much about the prospective health of the issuer even 20 years hence conditional on survival to that point, let alone 50 years hence conditional on survival to that point.

Yield versus Average Spot Rate

The *expectations hypothesis* states that the forward rate at which we would agree now to borrow or lend at some future date is equal to the expected future spot rate. Cox, Ingersoll, and Ross have pointed out that the expectations hypothesis cannot hold for all maturities and forward dates. This result is an implication of the differences in convexities and Jensen's inequality, and it is a close relative of all the remaining theoretical results in our paper. They

argue persuasively that the natural canonical assumption is the *local expectations hypothesis*, which asserts that all fixed-income securities have the same short-term expected return. By now, the local expectations hypothesis is the usual assumption made in practice by people who are pricing fixed-income derivatives. The two forms of the expectation hypothesis are very similar at short maturities, but the difference between the two grows more and more important as maturity increases.

Under the expectations hypothesis, the yield on a discount bond is equal to the average (over time and random realizations) of the spot rate from the quotation date until maturity. However, under the more reasonable assumption of the local expectations hypothesis, the price is the average (across random realizations) of the present value of \$1 at the end, discounted in each realization by the rolled-over spot rate. The convexity (nonlinearity) of the present value in in the rolled over spot rate means that the price is greater than the price corresponding to the average of the rolled over spot rates (over realizations and time). Or, a larger price corresponds equivalently to a lower yield. This result is proven in Theorem 1 in the Appendix.

Figure 1 shows this average spot rate versus the yield for a numerical example with reasonable parameter values (described in detail in the Appendix). In the example, the average expected spot rate over 100 years is 8.6% but the yield on the corresponding discount bond is only 6.9%. A similar pattern emerges at less exreme maturities as well: in the same example, the average spot rate over 30 years is 7.7% but the yield on the corresponding discount bond is only 6.7%. These numbers illustrate the general result that people pricing long bonds using the expectations hypothesis are setting the yields too high (and therefore are setting the prices too low). The expectations hypothesis works well (and the yield is nearly the same as the expected average spot rate) at short maturities, but the two are becoming significantly different even at intermediate maturities of 10-20 years.

Long-maturity Forward Rates

It is useful to examine forward rates at long maturities since they are the interest rates that are implicit in the forward purchase of a bond. In parallel with the result about discount bond pricing, the forward rate should be less than the expected spot rate provided interest rates at different points in time are tend to move up and down together. Theorem 2 in the Appendix formalizes this result.

In the same example we explored in Figure 1, the forward rates and expected future spot rates are shown in Figure 2. In the example, the expected spot rate in 50 years is 9.0% but the associated forward rate is only 7.1%. And, the expected spot rate in 20 years is 8.5% while the associated forward rate is only 7.0%.

Looking at forward rates formally as is done in the Appendix offers another insight. Since the forward rates depend on a weighted expectation of a spot rate at that time (see (5) in the Appendix), it seems likely that what we know now should have little effect on the forward rate.[5] For example, a forward rate 25 years in the future is six presidential administrations away (under the current U.S. system): what do we know about the actions of the Fed, the Congress, and the economy as a whole that far in the future? We conclude that forward rates should not move around very much. We are skeptical of the economic content of the amount of movement of rates at the long end, and our skepticism is consistent with what seem to be attractive opportunities for expected profit.

The assumption that interest rates tend to move together may seem to contradict recent experience, in which we have sometimes seen a Fed decision to tighten money followed by an increase in spot rates (the direct effect of tightening) with a decrease in long rates (due to reduced expectations of inflation). However, this is the exception that proves the rule: in general,

[5]This is a little subtle, since the weights depend on current information as well as later information. However, unless we have a lot of changing expectations about what will happen a long time from now, the effect of current information on the relative weights is small.

increases in Fed credibility keep down long rates *and* short rates, since when the Fed is most credible only a very small increase in short rates is needed to reaffirm investors' confidence in the Fed's resolve.

Coupon Bond Rates

So far we have focused on how yields and forward rates are affected by the error in using the traditional expectations hypothesis. Yields on par coupon[6] bonds are also affected, although the size of the effect is smaller because coupon rates are weighted averages that place less weight (and as maturity increases vanishing weight) on longer maturities. Most of a coupon bond's value is contained in early cash flows, even as the maturity of the bond increases without bound. Since what happens at larger maturities has a small impact on value, it also has a small impact on coupon bond yields. In the same example underlying Figures 1 and 2, the par coupon bond yield curve is plotted in Figure 3, along with the coupon bond yield curve corresponding to the traditional expectations hypothesis. The coupon bond yield curve is quite flat beyond 30 years. At 30 years, the coupon yield is 6.5% while the corresponding computed value given the expectations hypothesis is 7.1%. At 100 years, the coupon yield is 6.5% while the corresponding par coupon yield constructed using the expectations hypothesis is 7.2%. While less dramatic than the errors in using the expectations hypothesis to compute forward rates, 70 basis points is still economically significant.

[6] We refer to *par* coupon bonds to identify uniquely the yield, since coupon bonds with the same maturity and different coupon rates will generally have different yields unless the yield curve is flat. Intuitively, discount bonds with lower coupons have higher duration and their yields will be influenced more by longer forward rates than bonds selling at premiums.

Parameter Uncertainty

When we have uncertainty about process parameters, our usual response is to use a reasonable value or some expected value. Due to the convexity of long bond prices in the interest rate, this may not give a reasonable answer for pricing long bonds. In particular, if there is uncertainty about what interest rates will tend to in the long run, then the yield on increasingly long discount bonds will be no larger than the smallest value in the distribution of possible values for the long-term average rate, as is shown in Theorem 3 in the Appendix. This result combines the spirit of the convexity argument of Cox, Ingersoll, and Ross and the long-term interest rate results of Dybvig, Ingersoll, and Ross.

This result seems very provocative. So long as there is any probability that the government will eventually learn how to control inflation to the point that nominal rates will stay as low as 1%, then the yield on sufficiently long discount bonds cannot be larger than about 1%. Of course, the practical question is one of how long "sufficiently long" must be. Figures 4, 5, and 6 are analogous to Fgures 1, 2, and 3, but for an example with parameter uncertainty. In this new example, there is a 10% probability that the long-term mean interest rate is around 1% and a 90% probability that the long-term mean is around 10%. In this example, the yield on a 100-year discount bond is 3.6%, although the corresponding average spot rate is 8.7%. Furthermore, the average spot rate 50 years from now is 9.1% but the corresponding forward rate is 3.2%. Also, the 100-year par coupon rate computed from the traditional expectations hypothesis is 7.3% but the model value is only 5.5%.

A similar but less extreme pattern arises at less extreme maturities. Over 30 years, the average expected spot rate is 7.8% while the corresponding discount bond yield is 6.2%. The expected future spot rate 20 years out is 8.6%, while the corresponding forward rate is only 6.4%. The traditional analysis would give a 30-year par coupon yield of 7.1%, while the better model gives 6.2%. Perhaps more troubling is that the traditional analysis predicts a gently upwards-sloping yield curve (since the expected future spot rate is increasing), while the yield curve for coupon bonds declines significantly.

Thus, if there is even a relatively small probability that interest rates will be brought into control over the long run, this can have a significant effect on the pricing of long bonds. It is a big mistake simply to plug in an average or best estimate for the parameter!

Conclusion

We have explored some of the tools that economic theory offers us when we are faced with pricing long-maturity bonds. This analysis has pointed to several potential pitfalls in valuing long bonds, each of which might mislead us to undervalue the bond. The remedy is a return to first principles when modeling bond prices.

Appendix: Formal Term Structure Theory and Examples

Discount rates

To focus on pure questions of valuation, we will ignore taxes and transaction costs.[7] Modern asset pricing theory tells us that bonds and other assets can

[7]Transaction costs probably have an immaterial effect on bond valuation, especially for longer bonds. Taxes may tend to increase the value of longer bonds, due to greater opportunity for tax timing, but this seems like a trivial consideration for most investors. Taxes may also lead to the creation of clienteles, due to differential tax treatment of short and long positions. While there is empirical evidence of tax clienteles in historical data on Treasury Bills, Bonds, and Notes, casual evidence suggests that the development of the market for stripped coupons and principal has minimized the clientele effects. Of course, there are important tax effects when pricing across asset categories, for example in comparing municipal bonds with Treasury issues, so our analysis applies only to one category at a time. For more on tax clienteles and bond pricing, see Schaefer, Stephen M. (1982), "Taxes and Security Market Equilibrium," in William F. Sharpe and Cathryn M. Cootner, eds., *Financial Economics: Essays in honor of Paul Cootner*, Englewood

be priced as conditional expectations. A claim paying only P_T at date T has a value P_t at an earlier time $t < T$ if

$$(2) \quad P_t = E_t[P_T \exp(-\int_{s=t}^{T} r_s ds)],$$

where E_t indicates expectation conditional on information at time t, r_s is the short interest rate at time s, and $\exp(-\int_{s=t}^{T} r_s ds)$ is the discount factor from time t to T corresponding to rolling over investments at the short interest rate. The expectation E_t is risk-adjusted as necessary. If the local expectations hypothesis (*LEH*) holds (all assets have the same short-term expected rate of return) no risk adjustment is needed. This seems to be a good approximation, and we will assume the LEH throughout. If the LEH did not hold, each of the results would have a natural analog in the risk-neutral probabilities.[8]

For riskless bonds, the valuation equation (2) has a particularly simple form, since for each coupon or repayment of face value P_T is constant, and we can obtain the total value of a bond by adding together the values of the individual cash flows. In particular, the price at time t of 1 at a later time T is given by

$$(3) \quad D_{t,T} = E_t[\exp(-\int_{s=t}^{T} r_s ds)],$$

with corresponding yield $y^D_{t,T} = -\log(D_{t,T})/(T-t)$. Important properties of bond prices follow from the convexity of bond prices in interest rates, which was emphasized by Cox, Ingersoll, and Ross and is now familiar to many practitioners. One implication is that the yield now on a discount bond is

Cliffs, NJ, Prentice Hall, 159–178, Dybvig, Philip H. and Stephen A. Ross (1986), "Tax Clienteles and Asset Pricing," *Journal of Finance* 41, 751–762, and Brown, Stephen A., and Philip H. Dybvig (1986) "The Empirical Implications of the Cox, Ingersoll, Ross Theory of the Term Structure of Interest Rates," *Journal of Finance* **41**, 617–630.

[8]One way of defining risk-neutral probabilities (aka *martingale probabilities* or *artificial probabilities*) is as an equivalent probability assignment that makes the LEH true.

less than the interest rate we expect on average between the quote date and maturity. Formally,

Theorem 1 *Given the LEH, the quoted discount bond yield $y^D_{t,T}$ at t on a bond maturing at T is always less than or equal to expected average interest rate $E_t[\int_{s=t}^T r_s]/(T-t)$ between the quote and maturity, with equality only if there is no uncertainty in the sense that $\int_{s=t}^T r_s$ is known at t.*

PROOF By the definitions, (3), and Jensen's inequality,

$$
\begin{aligned}
(4) \quad y^D_{t,T} &= -\log(D_{t,T})/(T-t) \\
&= -\log(E_t[\exp(-\int_{s=t}^T r_s ds)])/(T-t) \\
&\leq E_t[-\log(\exp(-\int_{s=t}^T r_s ds))]/(T-t) \\
&= E_t[\int_{s=t}^T r_s ds]/(T-t),
\end{aligned}
$$

where the inequality (from Jensen's inequality) is strict if $\int_{s=t}^T r_s ds$ is not constant conditional on information at t. ∎

Forward rates

The forward rate quoted at time t for a later time T is given by

$$
\begin{aligned}
(5) \quad f_{t,T} &= -\frac{\partial \log(D_{t,T})}{\partial T} \\
&= -\frac{1}{D_{t,T}} \frac{\partial D_{t,T}}{\partial T} \\
&= \frac{E_t[r_t \exp(-\int_{s=t}^T r_s ds)]}{E_t[\exp(-\int_{s=t}^T r_s ds)]},
\end{aligned}
$$

where the last expression follows from the discount bond valuation formula (3). The forward rate is a weighted average of future spot rates where the

weights are the discount factors. When interest rates are high between t and T, then the forward rate at T doesn't carry much weight since in present value terms the forward commitment isn't worth much. Since interest rates tend to move up and down with inflation (even if the expected real return to short investments is not constant), this can also be interpreted as saying it is not very relevant what will happen when the forward borrowing or lending is in heavily deflated currency.

Here is the formal result that says that the forward rate is less than the expected spot rate provided interest rates tend to move together.

Theorem 2 *Given the LEH and the assumption that* $E[r_T | \int_{s=t}^{T} r_s]$ *is increasing in* $\int_{s=t}^{T} r_s$ *(i.e. rates tend to move together), then the forward rate* $f_{t,T}$ *is less than the expected spot rate* $E_t[r_T]$.

PROOF The result follows immediately from (5) since the weighted average puts relatively more weight than the simple expectation on states in which $\int_{s=t}^{T} r_s$ is small and therefore on average r_T is small. ∎

Parameter uncertainty

Theorem 3 *Given the LEH, suppose the long-term average interest rate*

$$(6) \quad \overline{r} \equiv \lim_{T \uparrow \infty} \frac{1}{T-t} \int_{s=t}^{T} r_s ds$$

exists almost surely as a random variable whose value may not be known at t. *(Note that the value of* $\overline{r}$ *is independent of the starting time* t *used in the definition.) Then if the limiting long discount yield* $\lim_{T \uparrow \infty} y_{t,T}^{D}$ *exists at time* t, *it is no larger than the smallest value in the support of* $\overline{r}$ *(i.e.* $\operatorname{essinf}(\overline{r})$*).*

PROOF By definition of $\overline{r}$ and its essential infimum, for each $\varepsilon > 0$, there exists $\hat{T} > 0$ and a set θ of states of nature with probability $p > 0$, such

that for all times $T > \hat{T}$ and all states $\omega \in \theta$, $(\int_{s=0}^{T} r_s ds)/T > \text{essinf}(\overline{r}) + \varepsilon$. Consequently, for all $T > \hat{T}$,

$$
\begin{aligned}
(7) \quad y_{t,T}^{D} &= -\log(E_0[\exp(-\int_{s=t}^{T} r_s ds)])/(T-t) \\
&\leq -\frac{\log(p\exp(-\int_{s=t}^{T}(\text{essinf}(\overline{r})+\varepsilon)ds))}{T-t} \\
&= -\frac{\log(p)}{T-t} + \text{essinf}(\overline{r}) + \varepsilon
\end{aligned}
$$

which tends to $\text{essinf}(\overline{r}) + \varepsilon$ as $T \uparrow \infty$, which implies that $\lim_{T\uparrow\infty} y_{t,T}^{D} \leq \text{essinf}(\overline{r}) + \varepsilon$. Since this is true for all epsilon, $\lim_{T\uparrow\infty} y_{t,T}^{D} \leq \text{essinf}(\overline{r})$, as we are to show. ∎

Numerical examples mentioned in the text

The numerical examples alluded to earlier are based on assumptions that are broadly consistent with what we know about interest rates and the term structure. First, the variance of the short rate is larger when the short rate is larger. Second, the variance of the short rate is not just a function of the short rate but has an independent component. Third, the particular dynamics of the short rate process are motivated by exploratory data analysis by the authors and by the empirical analysis of Kearns.[9] Specifically, the spot interest rate process for quarterly interest rates is given by

$$(8) \quad r_t = \max(0, kr^* + (1-k)r_{t-1} + \sigma_t\sqrt{r_{t-1}}z_t),$$

where

$$(9) \quad \sigma_t = (k_\sigma \sigma^* + \sigma_{t-1}(1-k_\sigma))\exp(\sigma_\sigma z_t^\sigma - {\sigma_\sigma}^2/2)$$

[9]See Kearns, Philip, 1993, *Volatility and the Pricing of Interest Rate Derivative Claims (Treasury Bills)*, Ph.D. Thesis, University of Rochester.

and the z_t's and z_t^σ's are jointly i.i.d. unit normal. Except for the simulation with uncertain r^*, the parameter choices (for the quarterly interest rate) are as follows: $r_0 = 5\%/4$, $r^* = 9\%/4$, $k = 0.1/4$, $\sigma_0 = 0.01/(\sqrt{0.06}\sqrt{4})$, $\sigma^* = 0.01/(\sqrt{0.06}\sqrt{4})$, and $\sigma_\sigma = 0.1/\sqrt{4}$, where these expressions include any normalization for moving from annual to quarterly and any normalization for setting the standard deviation at the "typical" value. In the simulation with uncertain r^*, all parameters are the same except for r^* which is a random variable equal to $1\%/4$ with probability 0.1 and $10\%/4$ with probability 0.9.

In terms of fitting what we know about interest rate data, the main weakness of the model appears at the short end, since we do not have enough factors or a short enough time interval to model the entire shape of the yield curve.[10] However, the single factor model should be fine for describing the long-term behavior we are focusing on in this paper.

Bond prices and average interest rates in each of these models are computed using Monte Carlo simulation with $1,000,000$ draws. Each simulation uses the same seed for the random number generator.

[10]For our purposes, the models that start with fitting the initial yield curve would be inappropriate, since they start with the yield curve (or forward rates) themselves and we are interested in the question of how those forward rates should be determined. Furthermore, those models typically do not have any way of ruling out or diagnosing implicit extreme assumptions about interest rate movements long in the futures (see Dybvig, Philip H. (1988) "Bond and Bond Option Pricing Based on the Current Term Structure," mimeo, Washington University in Saint Louis.

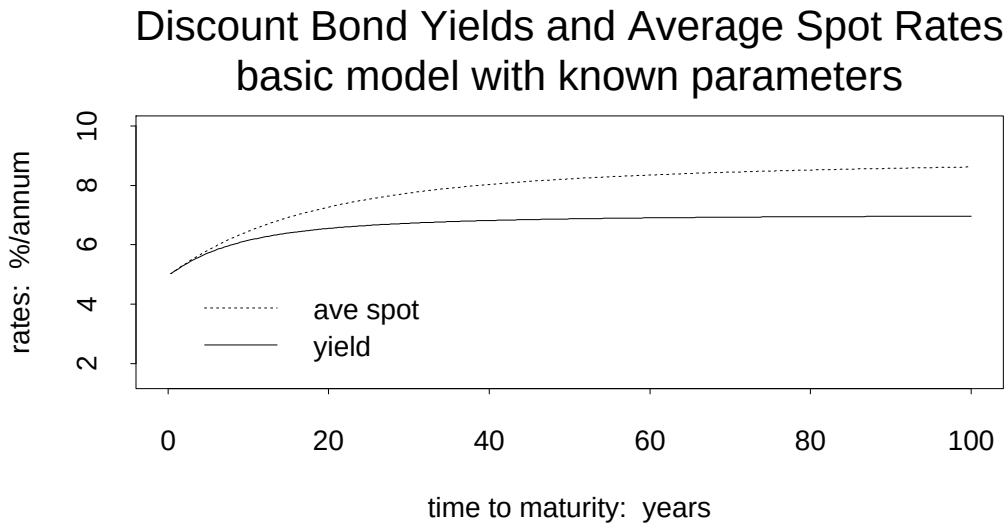

Figure 1: The expectations hypothesis asserts that the yield on a zero-coupon bond is the average, over time and random realizations, of the spot rate between the quotation date and maturity. Under the more reasonable local expectations hypothesis, this is a good approximation for short maturities but not for longer maturities, as illustrated in an example. The example, detailed in the Appendix, has a reasonable interest rate process with volatility dependent on the interest rate level and on a separate process.

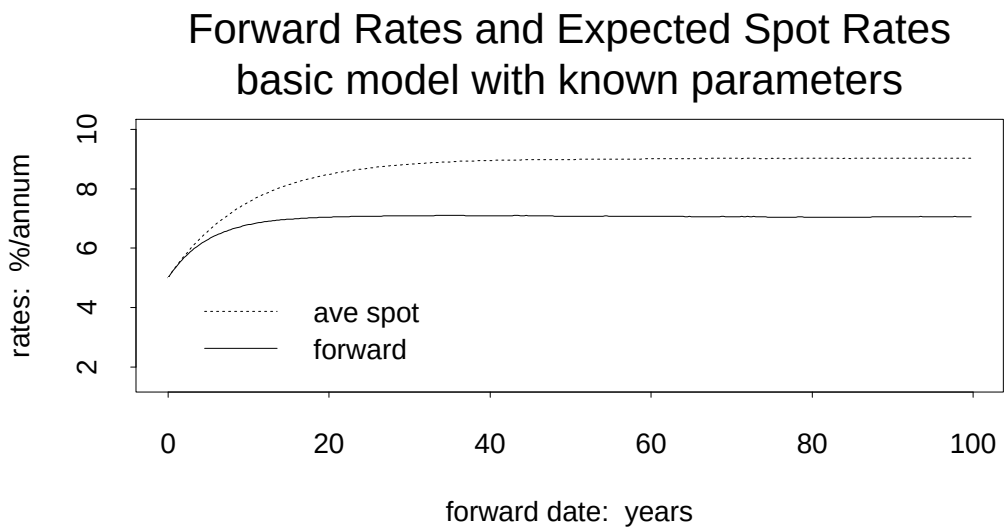

Figure 2: The expectations hypothesis asserts that the forward rate is the expectation of the corresponding future spot rate. Under the more reasonable local expectations hypothesis, this is a good approximation for short maturities but not for longer maturities, as illustrated in the same example used in Figure 1.

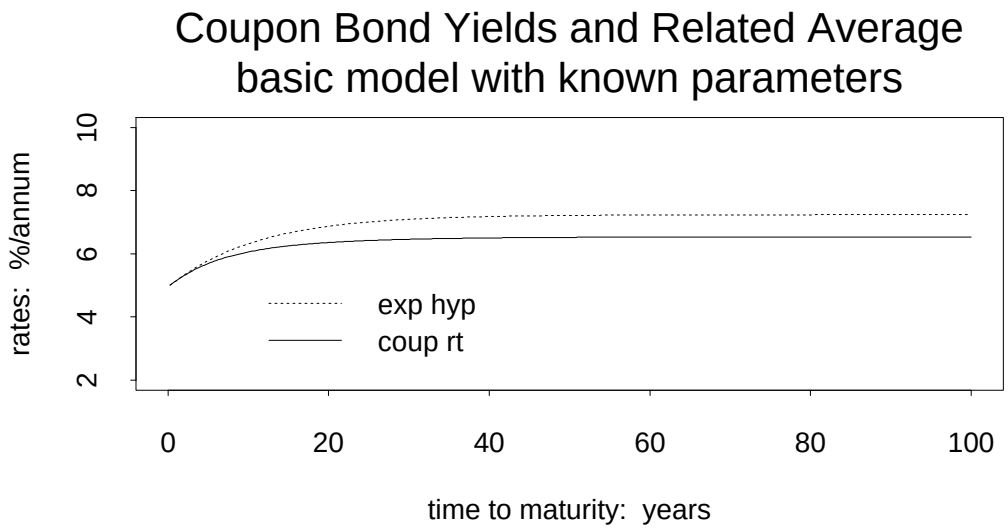

Figure 3: Par coupon bond yields are smaller than what is predicted by the traditional expectations hypothesis. In the same example as the previous two figures, the par coupon yield curve is very flat at the long end and 70 basis points below what would be predicted by the traditional model.

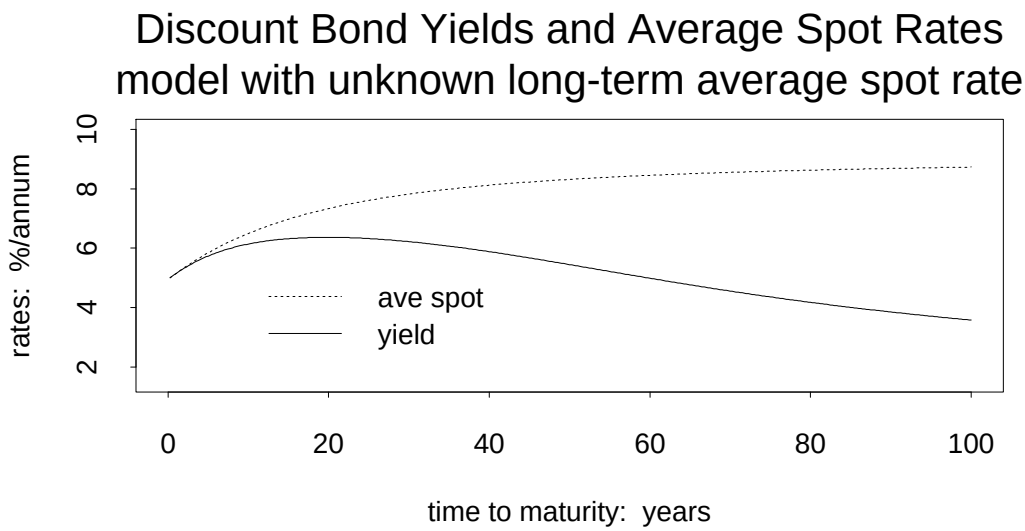

Figure 4: The violation of the expectations hypothesis illustrated in Figure 1 becomes particularly dramatic when interest rate process parameters are not known with certainty. This shows the deviation from the expectations hypothesis in an example with a 10% chance that rates will average around 1% and a 90% chance that rates will average around 10%.

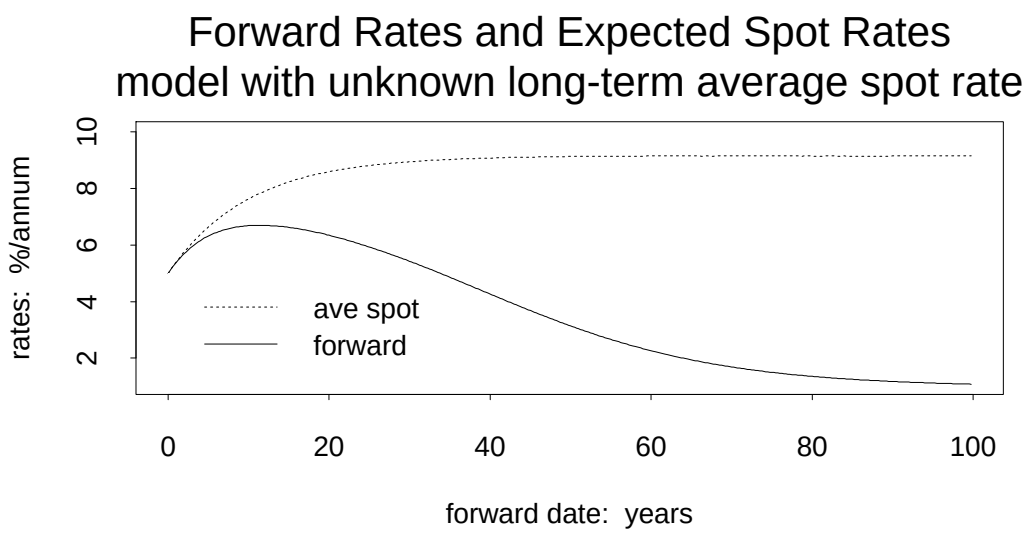

Figure 5: In the same example as in Figure 4, the forward rate drops quickly below the expected spot rate and moves nearer to the lowest possible long-term average spot rate.

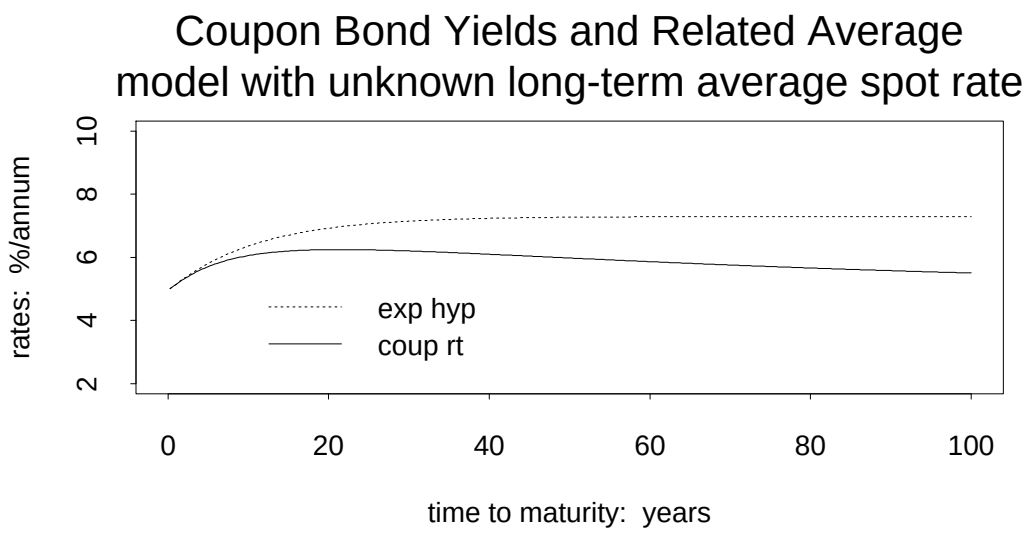

Figure 6: In the same example as in Figure 4 and 5, the coupon bond rate falls at large maturity while the yield curve implied by the expected future spot rate (expectations hypothesis) rises gently.