

Explaining the Increased Variability in Long-Term Interest Rates

Mark W. Watson

Monetary policy affects the macroeconomy only indirectly. In the standard mechanism, changes in the federal funds rate, the Federal Reserve's main policy instrument, lead to changes in longer-term interest rates, which in turn lead to changes in aggregate demand. But the links between the funds rate, long rates, and demand may be far from tight, and this potential slippage is a fundamental problem for monetary policymakers. In particular, long-term interest rates sometimes move for reasons unrelated to short-term rates, confounding the Federal Reserve's ability to control these long-term rates and effect desired changes in aggregate demand. Has the link between long rates and short rates weakened over time, therefore making it more difficult for the Federal Reserve to achieve its macroeconomic policy objectives through changes in the federal funds rate?

Such questions naturally arise when one observes the behavior of long-term interest rates. For example, Figure 1 plots year-to-year changes in ten-year Treasury bond yields from 1965 through 1998. (The volatile period of the late 1970s and early 1980s has been masked to highlight differences between the early and later periods.) The most striking feature of the plot is the increase in the variability of long-term rates in the recent period relative to the earlier period. Indeed, the standard deviation of long rates essentially doubled across the two time periods. What caused this increase in variability? Did a change in

■ I thank Michael Dotsey, Robert Hetzel, Thomas Humphrey, Yash Mehra, Pierre Sarte, and Ross Valkanov for comments on a previous draft of this article. This research was supported by National Science Foundation grant SBR-9730489. The views herein are the author's and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.


the behavior of short-term interest rates (caused, for example, by a change in Federal Reserve policy) lead to this dramatic increase in long-rate variability? Or, rather, is this change in variability caused by changes in factors unrelated to short-term rates, often described under the rubric of “term” or “risk” premia?

In what follows, we study the behavior of short-term interest rates over the two sample periods, 1965–1978 and 1985–1998, highlighted in Figure 1. It focuses on two key questions. First, has the short-term interest rate process changed? Second, can these changes in the behavior of short-term interest rates explain the increased volatility in long-term interest rates? The answer to both of these questions is yes; our findings suggest no weakening of the link between short rates and long rates and thus no weakening of the link between the Federal Reserve’s policy instrument and its ultimate objectives.

The variability in long-term interest rates is tied to two distinct features of the short-rate process: (1) the variability of “shocks” or “innovations” to short-term interest rates, and (2) the persistence (or half-life) of these shocks. In the standard model of the term structure, changes in the variability of short-rate innovations lead to proportional changes in the variability of the long rate. Thus, holding everything else constant, doubling the standard deviation of the innovation in short-term interest rates would lead to doubling the standard deviation of long rates evident in Figure 1.

The relationship between short-rate persistence and long-rate variability is more complicated. To explain this relationship it is useful to consider an example in which the short-term interest rate process can be described by an autoregressive model with one lag (an AR(1)). Let ρ denote the autoregressive coefficient associated with the process. When $\rho = 0$, short rates are serially uncorrelated, and shocks have only a one-period effect on the short-term interest rate. In contrast, when $\rho = 1$, short rates follow a random walk so that shocks to the current value of short rates lead to a one-for-one change in all future short rates. When long-term interest rates are viewed as discounted sums of expected future short-term rates, these different values of ρ imply very different behavior for long-term rates. For example, when $\rho = 0$, a change in the current short rate has no implications for future values of short rates, so long rates move very little. In contrast, when $\rho = 1$, any change in the current short rate is expected to be permanent and all future short rates are expected to change. This change in expected future short rates leads to a large change in the long-term rate. Values of ρ between 0 and 1 are intermediate between these two extremes, but in a subtle way that will turn out to be important for explaining the increased variability in long-term interest rates evident in Figure 1. In particular, for long-lived bonds, a short-rate process with $\rho = 0.9$ generates long rates that behave much more like those associated with $\rho = 0$ than with $\rho = 1$. Put another way, changes in the autoregressive parameter ρ have large effects on the behavior of long-term rates only when ρ is very close to 1. Such a result is familiar from studies of consumption behavior using

**Figure 1 Ten-Year Treasury Bond Yields
Annual Differences**



the present-value model, where the variability of changes in consumption increase dramatically as income approaches a “unit-root” process (Deaton 1987, Christiano and Eichenbaum 1990, Goodfriend 1992, and Quah 1992).

As a preview of the empirical results in later sections, we find that the variability of short-term interest rate shocks was *smaller* in the later sample period than in the earlier period. If there were no other changes in the short-rate process, this decline in short-rate variability should have led to a *fall* in the standard deviation of long-term interest rates of approximately 50 percent, as opposed to the 100 percent *increase* shown in Figure 1. However, we also find evidence of an increase in persistence: for example, the estimate of the largest autoregressive root in the short-rate process (the analogue of ρ from the AR(1) model) increased from 0.96 in the early period to nearly 1.0 in the later period. By itself, the increase in persistence should have led to a three-fold increase in the standard deviation of long rates. Taken together, the decrease in short-rate variability and increase in persistence explain remarkably well the increase in the variability of long rates evident in the data.

The estimated change in the persistence of the federal funds process has important implications for the Federal Reserve's leverage on long-term rates. For example, the estimated autoregressive process for the early sample period implies that a 25 basis point increase in the federal funds rate will lead to only a 3 basis point increase in ten-year rates. The autoregressive process for the later period implies that the same increase in the federal funds rate will lead to a 15 basis point increase in ten-year rates. Alternatively, the increase in persistence makes it possible to achieve a given change in the long rate with a much smaller change in the federal funds rate. The "cost" of increased leverage is the implicit commitment not to reverse changes in the federal funds rate, that is, to maintain the persistence in the short-rate process. The benefit of increased leverage is the reduced variability in the short-term rate. These costs and benefits are discussed in detail by Woodford (1999), who argues that it may be beneficial for the monetary authority to commit to making only persistent changes in its policy instrument.

The article is organized as follows. Section 1 documents changes in the variability of both long-term and short-term interest rates from the 1960s to the present. Here we document the decrease in variability of short-term interest rates (the federal funds rate and three-month Treasury bill rates) but an increased variability in longer-term rates (one-, five-, and ten-year Treasury bond rates). The relative increase in variability is shown to depend on the horizon of the interest rate—it is much higher for ten-year bonds than for one-year bonds, for example.

Section 2 studies changes in the persistence of short-term interest rates over the two sample periods. It begins by using a hypothetical AR(1) model for short-term interest rates to quantify the potential effects of short-rate persistence on the variability of long-term interest rates. The calculations are carried out using a standard model linking long rates to short rates—the expectations model with a constant term/risk premium. In this model, changes in long-term interest rates reflect changes in current and future values of short-term interest rates. The persistence of short-term interest rates is important because it affects the forecastability of short-term rates and thus the effect of changes in the short rate on long rates. The results indicate that, when ρ is very near 1, a relatively small change in ρ can lead to a large change in the variability of long-term interest rates.

Also in Section 2 we present empirical estimates of the short-term interest rate processes for the early and later sample periods using monthly values of the federal funds rate. These estimated processes show a fall in the variance of the short rate but an increase in persistence. Statistical inference about persistence is complicated by the near unit-root behavior of the short rate. This behavior leads to bias in the ordinary least squares (OLS) estimates and a nonstandard sampling distribution for test statistics for shifts in the process across the two sample periods. The article corrects the OLS estimates for bias using a procedure

developed in Stock (1991) and develops a new statistical test for a change in an autoregression that can be applied when data are highly persistent.

In Section 3 the variance of long-term interest rates is calculated using the expectations model together with the estimated processes for the short rate. These calculations show that the changes in the estimated short-rate process lead to increases in long-rate variability quite similar to the change found in the long-rate data.

Finally, Section 4 discusses the robustness of the empirical conclusions to specifics of the econometric specification, and Section 5 concludes. Econometric details concerning tests for changes in the persistence of the short-rate process are given in the Appendix.

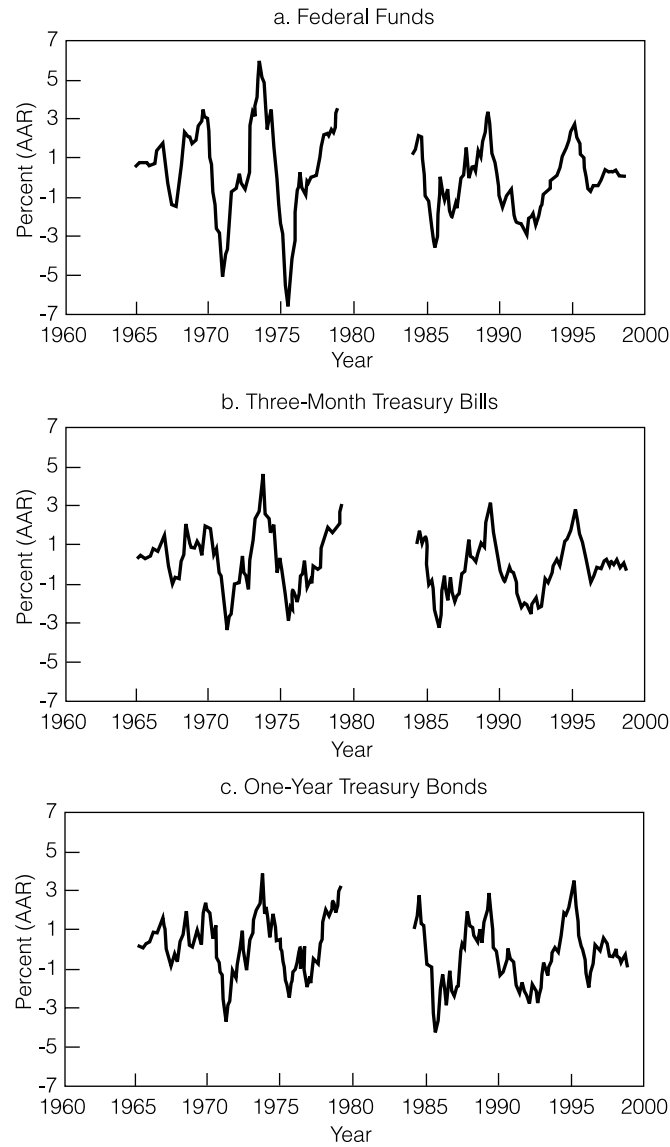
1. CHANGES IN THE VARIABILITY OF U.S. INTEREST RATES

The first task is to examine shifts in the volatility of market interest rates. Figure 2 plots year-over-year changes in six different interest rates over 1965 to 1998. As in Figure 1, the data from 1979 to 1984 are masked to highlight differences between the early sample period (1965:1–1978:9) and the more recent period (1985:1–1998:9). The interest rates range from very short maturity (the federal funds rate) to long maturity (ten-year Treasury bonds and AAA corporate bonds).¹ Each series is a monthly average of daily observations of the interest rates measured in percentage points at annual rates. Table 1 presents standard deviations for changes in interest rates over different sample periods. Panel a reports results for the year-over-year changes plotted in Figure 2, panel b reports results for monthly changes ($R_t - R_{t-1}$), and panel c reports standard deviations of residuals from estimated univariate autoregressions.

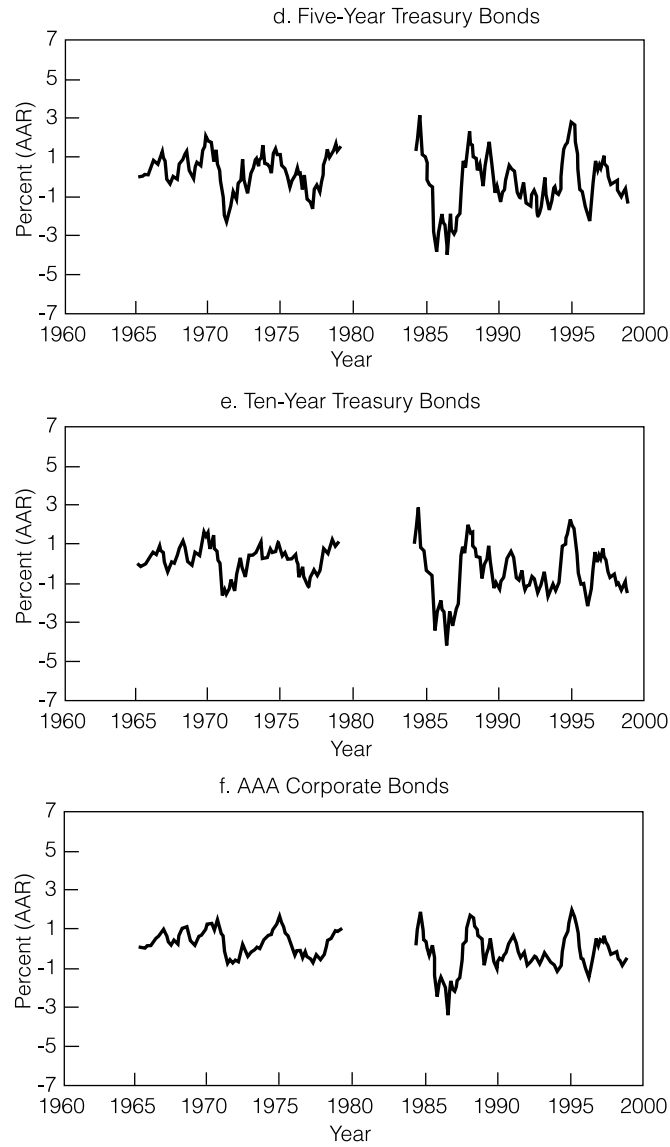
As seen in the figures and table, the volatility of long-term rates is much higher in the recent period than in the 1965–1978 sample period, but that is not the case for short-term rates. For example, from panel a of Table 1, the standard deviation of year-over-year changes in ten-year Treasury bond rates increased from 0.69 (69 basis points) in the 1965–1978 period to 1.29 (129 basis points) in the 1985–1998 period. A similar large increase is evident for AAA Corporate bond rates and for five-year Treasury bonds. At the shorter end of term structure, volatility did not increase. Indeed there is a substantial fall in the variability of the federal funds rate from 2.44 (244 basis points) to 1.50 (150 basis points).

The remainder of the table investigates the robustness of this conclusion about volatility both with respect to sample period and data transformation.

¹ All of the data are from the DRI database. The series are FYFF, FYGM3, FYGT1, FYGT5, FYGT10, and FYAAAC.

Figure 2 Annual Differences of Interest Rates

As shown on the table, this conclusion does not depend on the precise dates used to define the “early” and “recent” periods. These 1965–1978 and 1985–1998 dates were chosen somewhat arbitrarily, and the same volatility results hold for a wide range of cutoff dates used to define the sample periods. Consequently,

Figure 2 Annual Differences of Interest Rates

defining the early period as 1955–1978 and the recent period as 1992–1998 leads to the same conclusions. However, results do change if the volatile period of the late 1970s and early 1980s is included: from Table 1 interest rates were much more volatile in this period than they were either before 1979 or after

Table 1 Standard Deviations of Interest Rate Changes
(Percent at Annual Rates)

a. Year-over-Year Differences						
Sample Period	Interest Rate					
	FedFunds	3M-TB	1Y-TB	5Y-TB	10Y-TB	Corp
1965:1–1978:9	2.44	1.50	1.40	0.89	0.69	0.60
1985:1–1998:9	1.50	1.37	1.54	1.40	1.29	1.00
1978:10–1984:12	4.12	3.29	3.12	2.35	2.06	1.81
1955:1–1978:9	2.02	1.33	1.30	0.83	0.62	0.51
1992:1–1998:9	1.28	1.17	1.38	1.24	1.07	0.81

b. First Differences						
Sample Period	Interest Rate					
	FedFunds	3M-TB	1Y-TB	5Y-TB	10Y-TB	Corp
1965:1–1978:9	0.44	0.37	0.37	0.26	0.20	0.13
1985:1–1998:9	0.23	0.21	0.28	0.30	0.27	0.21
1978:10–1984:12	1.36	1.10	1.03	0.68	0.57	0.47
1955:1–1978:9	0.38	0.32	0.32	0.22	0.17	0.11
1992:1–1998:9	0.15	0.16	0.23	0.26	0.23	0.17

c. AR Innovations						
Sample Period	Interest Rate					
	FedFunds	3M-TB	1Y-TB	5Y-TB	10Y-TB	Corp
1965:1–1978:9	0.38	0.35	0.34	0.25	0.19	0.12
1985:1–1998:9	0.21	0.18	0.25	0.27	0.25	0.19
1978:10–1984:12	1.23	0.94	0.87	0.56	0.50	0.40
1955:1–1978:9	0.35	0.30	0.29	0.21	0.16	0.10
1992:1–1998:9	0.14	0.15	0.21	0.23	0.20	0.16

Notes: Entries are the sample standard deviations of the series over the sample period given in the table's first row. Year-over-year differences are $R_t - R_{t-12}$, first differences are $R_t - R_{t-1}$, and AR innovations are residuals from AR(6) models that incorporate a constant term.

1984. Finally, the results from different panels show that the same qualitative conclusion follows when year-over-year differences are replaced with monthly differences or with residuals from univariate autoregressions. For example, the standard deviation of the residuals in a univariate autoregression for ten-year Treasury bond rates increased from 19 basis points in 1965–1978 to 25 basis points during 1985–1998 (see panel c). The corresponding standard deviation for the federal funds rate fell from 38 to 21 basis points.

Since the variability of short-term rates was smaller in the later sample period than in the early period, it is clear that changes in the variability of short rates cannot explain the increased variability of long rates. We will have to look elsewhere, and with that in mind, the next section investigates changes in persistence in the short-rate process.

2. CHANGES IN THE PERSISTENCE OF U.S. INTEREST RATES

Before examining the empirical results on the persistence of short-term interest rates, it is useful to review the mechanism that links changes in short-rate persistence with changes in long-rate variability. This mechanism can be described using a simple expectations model of the term structure. Thus, let R_t^h denote the yield to maturity on an h -period pure discount bond, and assume that these yields are related to short-term rates by

$$R_t^h = \frac{1}{h} \sum_{i=0}^{h-1} E_t R_{t+i}^1,$$

where R_t^1 is the corresponding rate on a one-period bond. This relation can be interpreted as a risk-neutral arbitrage relation. Now, suppose that short-term rates follow an AR(1) process

$$R_t^1 = \rho R_{t-1}^1 + \varepsilon_t$$


so that $E_t R_{t+i}^1 = \rho^i R_{t+i}^1$ for $i \geq 1$. Then

$$R_t^h = R_t^1 \left[\frac{1}{h} \sum_{i=0}^{h-1} \rho^i \right]$$

so that long rates are proportional to short rates, with a factor of proportionality that is an increasing function of the persistence parameter ρ . Complications to the model (incorporation of term/risk premia, allowance for coupon payments, etc.) change details of the link between long rates and short rates. They do not, however, change the key feature of the model—namely, that long rates depend on a sequence of expected future short rates and that the variance of this sequence depends critically on the persistence of shock to short-term rates.

Of crucial importance is the quantitative impact of short-rate persistence on long-rate variability. Figure 3 gives a sense of this impact. Using the expectations relation given above, it plots the standard deviation of year-over-year changes in R_t^h (that is, $R_t^h - R_{t-12}^h$) as a function of ρ . Results are shown for different maturities h , and the scale of the plot is fixed by setting the innovation

**Figure 3 Annual Differences in Interest Rates
Implied Standard Deviation from AR(1) Model**



variance of short rates (σ_ε^2) equal to 1. The plot shows the functions for values of ρ between 0.95 and 1.00, which is the relevant range for the monthly data studied in this article. For short maturities (small values of h) ρ does not have much of an effect on the standard deviation interest rate. For example, as ρ increases from 0.95 to 1.00, the standard deviation of one-period rates increases by a factor of 1.1 (from 3.1 to 3.5). However, ρ has a large effect on the variability of long-term interest rates. When $h = 120$ (a ten-year bond when the period is a month), then as ρ increases from 0.95 to 1.00, the resulting standard deviation of long rates increases by a factor of 7 (from 0.5 to 3.5). Moreover, the rate of increase in the standard deviation increases with the value of ρ . Thus, the implied changes in the volatility of long rates across sample periods will depend both on the level of ρ and on its change.

Having considered the analytical importance of persistence, we now examine the empirical evidence on it. Table 2 contains estimates of the persistence in short-term rates for the two sample periods. Results are presented for both the federal funds and the three-month Treasury bill rate. Univariate autoregressions are fit to the series, and persistence is measured by the largest root of the implied autoregressive process. This largest autoregressive root determines the effect of

Table 2 Largest Autoregressive Roots for Short-Term Interest Rates

Interest Rate	Sample Period						Chow Test	
	1965:1–1987:9			1985:1–1998:9			F_ρ	P-Value
	ρ_{ols}	ρ_{mub}	90% CI	ρ_{ols}	ρ_{mub}	90% CI		
Federal Funds	0.97	0.96	0.91-1.01	0.98	1.00	0.94-1.02	1.19	0.30-0.64
3-Month TBill	0.96	0.98	0.93-1.02	0.98	0.99	0.94-1.02	1.17	0.31-0.64

Notes: ρ_{ols} is the OLS estimate of ρ constructed from an AR(6) model that included a constant term. ρ_{mub} is the median-unbiased estimator of ρ constructed from the Dickey-Fuller τ^μ statistic as described in Stock (1991). The 90 percent confidence interval is also computed from τ^μ using Stock’s procedure. F_ρ is the Chow F-statistic for testing for change in ρ across the two sample periods. The column labeled P-value shows the upper and lower bound for the F-statistic P-value using the procedure described in the Appendix.

shocks on long horizon forecasts of short rates and therefore summarizes most of the information about the link between short rates and long-term interest rate variability. We denote the parameter by ρ as in the AR(1) model discussed above.

The first entry for each sample period is the OLS estimate of ρ (denoted ρ_{ols}) computed from an AR(6) model. (The next section will discuss the robustness of results to the lag length in the autoregression.) The values of ρ_{ols} are very large both for the two interest rates and the two sample periods. The implication is that short rates were apparently highly persistent in both sample periods. There is some evidence of a small increase in ρ in the latter sample period: the value of ρ_{ols} increases from 0.97 to 0.98 for federal funds and from 0.96 to 0.98 for three-month Treasury bills. However, interpreting these changes is difficult because of statistical sampling problems associated with highly persistent autoregressions. These problems are well known in autoregressions with unit roots, but similar problems also arise when roots are close to unity. To aid the reader, we digress with a short statistical primer before discussing the other entries in Table 2.

When values of ρ are close to 1 and the sample size is moderate (as it is here), then the sampling distributions of OLS estimators and test statistics differ markedly from the distributions that arise in the classical linear regression model. In particular, ρ_{ols} is biased, and the usual t-statistics have non-normal distributions. One cannot construct confidence intervals for ρ in the usual way. Of course, as long as ρ is strictly less than 1, the usual asymptotic statistical arguments imply that these difficulties disappear for a “suitably” large sample size. Unfortunately, the sample size used in this article (like that commonly used in empirical macroeconomic research) is not large enough for the conventional asymptotic normal distributions (based on stationarity assumptions)

to provide an accurate approximation to the sampling distribution of the usual OLS statistics. We must use alternative and more accurate approximations.

In empirical problems when ρ is close to 1 (say in the range 0.90–1.01) and the sample size is moderate (say less than 200 observations), econometricians have found that “local-to-unity” approximations provide close approximations to the sampling distribution of OLS statistics.² In the present context, these approximations will be used to construct unbiased estimators of ρ , confidence intervals for ρ , and Prob-values in tests for changes in ρ over the two sample periods. Specifically, “median-unbiased” estimates and confidence intervals for ρ are constructed from the Dickey-Fuller τ^μ statistic using the procedures developed in Stock (1991).³ Tests for changes in ρ across the two sample periods are carried out using the usual Chow-F statistic. This statistic is computed as the Wald statistic from changes in the values of ρ_{ols} over the two sample periods. The regressions are estimated separately in each sample period, so that all of the coefficients are allowed to change, but the Wald statistic tests for a change in the largest root only. (Changes in the other autoregressive parameters will have little effect on the variance of long rates, so we focus the test on the largest root.) The statistical significance of the Chow statistic can be determined using Prob-values computed from the local-to-unity probability distributions. These alternative Prob-values are described in detail in the Appendix. As the Appendix shows, the Prob-value depends on the true, and unknown, value of ρ . Thus, rather than reporting a single Prob-value, we report an upper and lower bound.

With this background, the reader can now understand other entries in Table 2. The unbiased estimates are reported in the column labeled ρ_{mub} (the *mub* subscript stands for “Median UnBiased”), and these are followed by the 90 percent confidence interval for ρ . The point estimates suggest that persistence was higher in the second period; for example, using the federal funds rate, the value of ρ_{mub} increased from 0.96 to 1.00. However, the confidence intervals show that there is a rather wide range of values of ρ that are consistent with the data—the confidence interval, which for federal funds in the first period is 0.91–1.01, shifts up to 0.94–1.02 in the second period. The overlap in these confidence intervals suggests that the apparent shift in ρ is not highly statistically significant, and this conjecture is verified by the Chow-statistic, which has a Prob-value that falls between 0.30 and 0.64. Thus, there is some evidence of a shift in the largest root, in a direction consistent with the behavior of long-term rates, but the shift is small and the exact magnitude is difficult to determine because of sampling error. However, when ρ is near 1, small changes in its value can cause large changes in the variability of long-term interest rates.

² Important early references in econometrics include Cavanagh (1985), Phillips (1987), and Stock (1991).

³ The median-unbiased estimator, which will be denoted ρ_{mub} , has the property that $\text{Prob}(\rho_{mub} \leq \rho) = \text{Prob}(\rho_{mub} \geq \rho) = 0.5$.

3. IMPLICATIONS OF THE CHANGES IN THE SHORT-RATE PROCESS ON LONG-RATE VARIABILITY

The changes in long-rate volatility associated with the changes in the short-rate process depend on the specifics of the model linking short rates to long rates. Before we compute the variability of long rates associated with the estimated short-rate processes from the last section, three issues need to be addressed in the present context.

First, the data used here, while standard, are not ideal. The data are not point sampled but rather are monthly observations of daily averages. The bonds contain coupon payments, which were missing in the simple theory presented above. The calculations presented below are based on two approximations. First, the process for one-month rates is estimated using the federal funds data. This is a rough approximation that uses a monthly average of daily rates as a monthly rate. As it turns out, similar results obtain if the federal funds process was replaced with the estimated process for the three-month Treasury bills, so the precise choice of short rate does not seem to matter much. The second approximation adjusts the present-value expectations model for coupon payments using the approximation in Shiller, Campbell, and Schoenholz (1983). Specifically, the expectational equation for long rates becomes

$$R_t^h = \frac{1 - \beta}{1 - \beta^h} \sum_{i=0}^{h-1} \beta^i E_t R_{t+i}^1,$$

where $\beta = 0.997$.

The second issue involves the expectations theory described above. That model used an AR(1) driving process for short rates, and constructed expectations using this process. The univariate process for short rates is more complicated than an AR(1) process; moreover, one can form short-rate expectations using a richer information set than one containing only lags of short rates. Extending the calculations to account for a higher-order univariate AR process is straightforward, as the exercise merely involves computing the terms $E_t R_{t+i}^1$ from a higher-order AR model. However, to account for a wider information set is more problematic. A standard and powerful approach to this problem is to construct bounds on the implied variance of long rates from the short process, using, for example, the approach in Shiller (1981). Unfortunately, this approach requires stationarity of the underlying data, so the bounds are likely to be inaccurate for the highly persistent data studied here. West (1988) proposes bounds for the expectational present-value model based on the innovations in the univariate processes and shows that these bounds hold for integrated as well as stationary processes. But as it turns out, West's results hold only for the

infinite horizon model, and the model here is finite horizon.⁴ Another approach is simply to specify a more general information set and carry out the analysis using, say, a vector autoregression (VAR) instead of a univariate autoregression. However, the statistical analysis becomes increasingly complicated in a VAR with highly persistent variables. For all of these reasons, the analysis here will be carried out using a univariate AR.

Finally, the calculations reported here ignore all term/risk premia and other deviations from the simple expectations theory. As mentioned above, even in more complicated versions of the models, the first-order impact of short-rate persistence on long-rate variability occurs through the expected present-value expression from the version of the model used here.

With these limitations in mind, consider now the implied variability in long-term rates. The results are summarized in Figure 4 and in Table 3, which shows the implied variability of interest rates computed from the expectations model, using the estimated short-rate process over the different sample periods and for different values of ρ . Results are shown for four maturities. Each panel of Figure 4 shows the variability of year-over-year changes in the interest rate implied by the estimated AR(6) model for the federal funds rate, where the estimates are derived by imposing the value of ρ shown on the x axis. Results are shown for both sample periods. Highlighted on the graphs are the results that impose the OLS and the median-unbiased estimates of ρ from Table 2. (A circle denotes the value of ρ_{ols} ; a square denotes ρ_{mub} .) In each panel, the variance function for the second period lies below the function for the first period. This shift is caused by the decrease in variance of the AR errors estimated for the second period. The vertical distance between the curves shows the change in variance for a given value of ρ . To compute the variance across periods, the value of ρ in each sample period must be specified. In terms of the figures, the vertical displacement of the plotted circles gives the change in variability across the two periods using the OLS estimates of ρ (ρ_{ols}). The displacement of the squares gives the change using the median-unbiased estimator (ρ_{mub}). The implied standard deviation for the four maturities in both sample periods and for ρ_{ols} and ρ_{mub} are given in Table 3. For comparison, the table also gives the period-specific sample standard deviations for the federal funds rate and the rates on one-, five-, and ten-year Treasury bonds.

There are substantial differences in the results across the four panels in Figure 4. For one-month rates (panel a), variability is essentially independent

⁴ In West's present-value model $y_t = E_t \sum_{i=0}^h \beta^i x_{t+i}$, the key restriction is that $E_t \beta^h x_{t+h}$ converges to zero in mean square as $h \rightarrow \infty$. This suggests that West's bounds will provide a good approximation in the finite horizon model so long as $E_t \beta^h x_{t+h}$ is small. Thus, the quality of the approximation will depend on the size of $(\beta\rho)^h$, where ρ is the largest autoregressive root. In the term structure model $\beta = 0.997$, and the x_t process is highly persistent, with a largest autoregressive root of, say, $\rho = 0.99$. Thus, for $h = 120$, $(\beta\rho)^h = 0.16$, which implies that $E_t \beta^h x_{t+h}$ will often be substantially different from 0.

Table 3 Standard Deviation of Annual Changes in Interest Rates


Maturity	Sample Period					
	1965:1–1978:9			1985:1–1998:9		
	Actual	Implied by		Actual	Implied by	
		ρ_{ols}	ρ_{mub}		ρ_{ols}	ρ_{mub}
1 Month	2.44	2.39	2.39	1.50	1.36	1.41
12 Month	1.40	1.80	1.76	1.54	1.31	1.46
60 Month	0.89	0.57	0.52	1.40	0.76	1.22
120 Month	0.69	0.31	0.29	1.29	0.46	0.98

Notes: Entries show the actual (sample value) and implied standard deviations of year-to-year changes in interest rates ($R_t^h - R_{t-12}^h$) for different horizons. Entries labeled Actual are taken from Table 1 and are the sample values for the federal funds rate and the rates for one-, five-, and ten-year Treasury bonds. The columns labeled ρ_{ols} and ρ_{mub} were computed using the expectations model and the estimated AR(6) processes using the federal funds rate over the sample periods shown and imposing the values of ρ_{ols} and ρ_{mub} listed in Table 2. These values correspond to the circles and squares shown in Figure 4.

of ρ and thus the model predicts a substantial decrease in variability during the second period. Since the federal funds rate data were used to estimate the short-rate process, this decrease in variability is essentially equal to the sample values—see the first row of Table 3. For one-year rates (panel b of Figure 4 and the second row of Table 3), variability is also predicted to decrease in the second period, but the decrease is far less than for one-month rates and depends on which estimator is used for ρ . (The implied decrease in the standard deviation is 49 basis points using ρ_{ols} and 30 basis points using ρ_{mub} .) In the sample, there was a small increase (14 basis points) in the standard deviation of one-year interest rates. At longer maturities (panels c and d of Figure 4 and the last two rows of Table 3), variability is predicted to increase in the second period, and again, the amount of the increase depends on the estimator of ρ that is used. The increase is not particularly large using ρ_{ols} (less than 20 basis points); however, it is much larger using ρ_{mub} (70 basis points). The small bias correction incorporated in ρ_{mub} results in this large difference because it pushes the second-period estimate of ρ very close to 1 and because the variance function is rapidly increasing in this region.

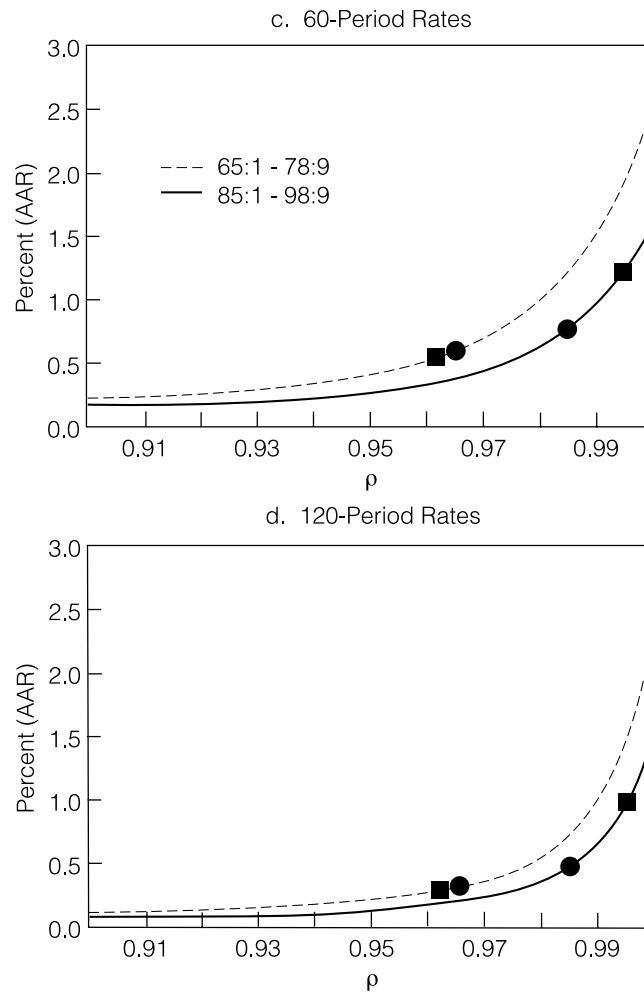
While the estimated difference in persistence, as measured by ρ_{mub} , explains much of the increase in variability in long-term interest rates, much of that variability is still unexplained. For example, in the first sample period the model's implied standard deviation for five-year rates is 52 basis points, while the sample standard deviation of actual five-year rates is 89 basis points. This leaves a “residual” component, orthogonal to short-term rates, with a standard

**Figure 4 Annual Differences in Interest Rates
Implied Standard Deviation from Fitted AR Models**
(Circle = ρ_{ols} , Square = ρ_{mub})



deviation of 72 basis points ($72 = \sqrt{89^2 - 52^2}$) representing the difference between the actual five-year rates and the value implied by the expectations model. Interestingly, a residual component of similar size (69 basis points) is necessary in the second sample period. (A somewhat larger residual is required for ten-year rates.) Thus, although the simple expectations model with constant term/risk premia and simple information structure leaves much of the variability

**Figure 4 Annual Differences in Interest Rates
Implied Standard Deviation from Fitted AR Models**
(Circle = ρ_{ols} , Square = ρ_{mub})



in long rate unexplained in both sample periods, it does explain the lion's share of the increase in variability across the two sample periods.

The results derived here, based on a simple version of the expectations theory of the term structure, are consistent with results derived by other researchers using reduced-form time-series methods. For example, the expectations theory, together with a process for the short-term interest, can be used to calculate

the change in the long rate associated with a given change in the short rate. Using the first-period estimates (and the values of ρ_{mub} shown in Table 2) the model predicts that a 25 basis point change in the federal funds rate would lead to a 3 basis point change in the ten-year bond rate. The second-period estimates imply that the same 25 basis point change in the federal funds rate would lead to a 15 basis point change in the long rate. Mehra (1996) estimates a reduced-form time-series model (a vector error correction model) of long rates and inflation over the 1957–1978 and 1979–1995 sample periods. His estimated models predict that a 25 basis point change in the federal funds rate led to a 3 to 7 basis point change in long rates in the early period and a 7 to 12 basis point change in the later period.

4. ROBUSTNESS OF RESULTS

This section discusses the robustness of the article's main findings to specification of lag length in the autoregression and to choice of sample period. The empirical results are summarized in Table 3. The first row in each panel of the table shows the results from the specification used in the last section, so these results are the same as reported in Table 2. Each of the following rows summarizes results from a different specification of either lag-length or sample period. Panel a of Table 4 shows results for the federal funds rate and panel b shows results for the three-month Treasury bill rate.

The AR lag length of 6 used in the baseline specification was suggested by the Akaike Information Criteria (AIC) and by t-tests on the autoregressive coefficients. Much shorter lag lengths were suggested by the Schwartz criteria (BIC). Table 2 shows results from specifications using 2, 4, and 8 lags. Each of these alternative specifications yield first-period estimates of ρ that are lower than the estimates from the AR(6) model; second-period estimates are essentially unchanged. The first-period differences in ρ_{ols} are small, but the differences are more substantial for the ρ_{mub} . Ignore for the moment the large amount of sampling error associated with these estimates. Even so, the new point estimates have little effect on the variance of long-term rates. From Figure 3, the long-rate variance function is relatively flat over the range of first-period ρ estimates given in Table 3. Thus, from Figure 3, the implied first-period standard deviation of long-term interest rate changes is 0.11 when $\rho = 0.93$ and increases to only 0.18 as ρ increases to 0.96. (The first ρ figure is the value of ρ_{mub} from the AR(4) first-period model; the second figure is the corresponding value of ρ_{mub} in the AR(6) model.) Both of these specifications imply a much larger second-period standard deviation (1.48 and 0.975 for the AR(4) and AR(6) models, respectively) since the second-period values of ρ_{mub} are very close to 1.0 in both specifications. Thus lag-length choice appears to have little effect on the qualitative conclusions.

Table 4 Largest Autoregressive Roots for Different Specifications

a. Federal Funds Rate								
Specification Change from Baseline	First Sample Period			Second Sample Period			Chow Test	
	ρ_{ols}	ρ_{mub}	90% CI	ρ_{ols}	ρ_{mub}	90% CI	F_ρ	P-Value
None	0.97	0.96	0.91-1.01	0.98	1.00	0.94-1.02	1.19	0.30-0.64
AR(2)	0.96	0.92	0.86-1.00	0.99	1.00	0.95-1.02	2.46	0.13-0.43
AR(4)	0.96	0.93	0.87-1.01	0.99	1.00	0.95-1.02	2.34	0.14-0.45
AR(8)	0.96	0.95	0.90-1.01	0.99	1.00	0.95-1.02	1.65	0.22-0.55
SD 1955	0.98	0.98	0.96-1.01	0.98	1.00	0.94-1.02	0.06	0.81-0.93
SD 1992, AR(2)	0.96	0.92	0.86-1.00	0.98	1.01	0.95-1.04	1.55	0.24-0.57

b. Three-Month Treasury Bill Rate								
Specification Change from Baseline	First Sample Period			Second Sample Period			Chow Test	
	ρ_{ols}	ρ_{mub}	90% CI	ρ_{ols}	ρ_{mub}	90% CI	F_ρ	P-Value
None	0.96	0.98	0.93-1.02	0.98	0.99	0.94-1.02	1.17	0.31-0.64
AR(2)	0.95	0.97	0.92-1.01	0.99	1.00	0.95-1.02	1.73	0.21-0.54
AR(4)	0.95	0.95	0.90-1.01	0.99	1.00	0.95-1.02	2.66	0.12-0.40
AR(8)	0.95	0.97	0.91-1.01	0.98	0.99	0.94-1.02	1.85	0.20-0.52
SD 1955	0.98	1.00	0.97-1.01	0.98	0.99	0.94-1.02	0.01	0.93-0.98
SD 1992, AR(2)	0.98	1.00	0.97-1.01	0.98	0.99	0.94-1.02	0.00	0.96-0.99

Notes: The first column shows the change in the specification from the baseline AR(6) model incorporating a constant (from Table 2). The baseline specification was estimated over the sample periods 1965:1–1978:9 and 1985:1–1998:9. AR(p) denotes an AR(p) model when a constant was used. “SD 1955” denotes a specification with the first sample period from 1955:1–1978:9. “SD 1992, AR(2)” denotes an AR(2) with second sample period from 1992:1–1998:9.

The choice of sample period has a more important effect. The baseline sample periods 1965:1–1978:9 and 1985:1–1998:9 were chosen to eliminate the large variability in interest rates during the late 1970s and early 1980s. With this volatile period eliminated, two samples of equal size were chosen (with 1998:9 being the last sample period available when this research was started). There is no compelling reason, other than equating statistical power in each sample, why the early and recent samples should be of equal size. The last two rows of the table show results from increasing the early sample period (by changing the beginning date to 1955:1) and decreasing the recent sample period (by changing the beginning date to 1992:1). Since the 1992–1998 sample period is very short, an AR(2) model was used for this specification. Evidently, the choice of the second period has little effect on the estimates of ρ , but the choice of first sample period does. Estimates of ρ are larger for both interest rates in the extended sample period 1955–1978 than in the 1965–1978 period. This increase should not be surprising given the behavior of interest rates over

the 1955–1978 period, where the dominant feature of the data is an increase in the “trend” level of interest rates. However, since this article’s analysis focuses on the behavior of long rates as they are affected by expected future short rates, the question is whether investors in the late 1950s anticipated this trend rise in interest rates, as would be suggested by ex post fitted values from the univariate autoregression. Such prescience seems unlikely.

5. SUMMARY AND DISCUSSION

We have documented the increase in the variability of long-term interest rate changes during the 1985–1998 period relative to the 1965–1978 period. In contrast, the variability of short-term interest rates decreased in the later period. A possible explanation for this differential behavior is a change in the persistence of changes in short-term rates: expectations theories of the term structure imply that such shifts in persistence will have a large effect on the variability of changes in long-term rates but have little effect on the variability of changes in short rates. Point estimates of the largest autoregressive root for short rates show an increase in persistence that is large enough to explain the increased variability in long rates. However, the short-rate persistence parameter is imprecisely estimated, so that it is impossible to reach definitive conclusions based on this analysis. The lack of precision raises two issues: one related to statistical technique and one related to learning about changes in central bank policy.

The first issue concerns using the behavior of long rates to infer the persistence of the short-rate process. This is appropriate if long rates and short rates are connected by the present-value model. This procedure is used in Valkanov (1998), where the model’s implied cointegration between long and short rates yields improved estimators for ρ . Valkanov then uses the improved estimator to overcome inference problems identified by Elliott (1998) in his critique of cointegration methods. Indeed, in a comment on a preliminary draft of this article, Valkanov (1999) uses his method to construct estimates of ρ together with 90 percent confidence intervals for the time periods 1962:1–1978:8 and 1983:1–1991:2 using data on the federal funds rate and ten-year Treasury bonds. He finds an estimate of ρ of 0.96 (with a 90 percent confidence interval of 0.93–0.98) in the early period and an estimate of 0.99 (with a 90 percent confidence interval of 0.99–1.00) in the later period (Valkanov 1999, Table 2c). His point estimates are essentially identical to the values of ρ_{mub} reported in our Table 2, but as expected from the use of a more efficient procedure, his confidence intervals are considerably narrower than the results presented in Table 2.

The large sampling uncertainty associated with estimates of the short-rate persistence suggests that the market will learn about changes in persistence very slowly from observing short-term interest rates. A central bank interested

in increasing the persistence of short-term interest rates (for the reason suggested in Woodford [1999], for example) would have to follow this policy for a considerable time to convince a market participant who relied only on econometric evidence that such a change had indeed taken place. For example, suppose that the federal funds process changed from one with a largest root of 0.96 to one with a largest root of 0.99, and after ten years in the new regime an econometrician tested the null hypothesis that $\rho = 0.96$ versus the alternative that $\rho > 0.96$ using a standard t-test at the 5 percent significance level. The econometrician would (correctly) reject this null only about 50 percent of the time. (That is, the power of the test using ten years of data is roughly 0.50.) Thus, it is likely the econometrician would have to observe the new federal funds process for quite some time before he concluded that the process had changed. This failure immediately to recognize policy shifts highlights the importance of other devices (institutional constraints, public statements, etc.) to more quickly convince a wary public that such shifts have occurred.

This article has presented econometric evidence suggesting that changes in the federal funds rate are more persistent now than they were in the 1960s and 1970s. Why did this change occur? We can offer but a few remarks on this important question. Here is one possible explanation. Suppose we decompose the funds rate into a real rate and an inflation component. If movements in the real rate are transitory, then the persistence in the funds rate will be driven by the inflation component. Therefore, an increase in the persistence of inflation possibly explains the increased persistence in the funds rate. This explanation, however, does not seem promising. For example, the values of ρ_{mub} computed using CPI inflation *fell* from 0.98 in the earlier sample period to 0.92 in the later period. As a result, inflation seems to have become *less* persistent, and this implies that some of the explanation must lie in the persistence of the real component of the funds rate. There is growing econometric evidence that the Federal Reserve's "reaction function" linking the federal funds rate to expected future inflation and real activity has been quite different under Chairmen Volker and Greenspan than under the previous three chairmen. For example, Clarida, Galí, and Gertler (1999) present evidence suggesting that the Federal Reserve responded more aggressively to expected future inflation after 1979 than in the previous two decades. Their evidence also suggests that the Federal Reserve more aggressively smoothed the funds rate in this latter period, consistent with the increased persistence found here. Changes in this reaction function undoubtedly contain the key to explaining the increased persistence in the federal funds rate process.

APPENDIX
A. Computing Prob-Values for the Chow Test Statistic for the Largest Autoregressive Root

This Appendix describes the method used to compute the Prob-values for tests of changes in the largest autoregressive root of a univariate autoregression. The specification is the AR(p) autoregression

$$x_t = \mu + u_t$$

with

$$u_t = \sum_{i=1}^p \phi_i u_{t-i} + \varepsilon_t,$$

where x_t denotes the level of the interest rate, μ is a constant denoting the average level of the process in the stationary model, and u_t is a stochastic term. The u_t process can be rewritten as

$$u_t = \rho u_{t-1} + \sum_{i=1}^{p-1} \pi_i (u_{t-i} - u_{t-i-1}) + \varepsilon_t,$$

where $\rho = \sum_{i=1}^p \phi_i$ and $\pi_i = -\sum_{j=i+1}^p \phi_j$. The parameter ρ is thus the sum of the AR coefficients. When one root of the AR polynomial $1 - \sum_{i=1}^p \phi_i z^i$ is close to 1 and all of the other roots are larger than 1, then ρ is also approximately equal to the inverse of the root closest to unity. In this case ρ is usually called the “largest” root because its inverse is the largest eigenvalue of the companion matrix of the model VAR(1) representation.

We study the behavior of statistics in a setting where ρ is modeled as close to 1.0, written as

$$\rho_T = 1 + \frac{c}{T}.$$

The artificial dependence of ρ on the sample size T facilitates the analysis of continuous asymptotic limits as $T \rightarrow \infty$.⁵ To simplify notation, we will present the AR(1) model, so that $\pi_i = 0$, for $i = 1, \dots, p - 1$. For the test statistics

⁵ To see this, contrast the discontinuous results

$$\lim_{T \rightarrow \infty} \rho^T = \begin{cases} 0 & \text{when } |\rho| < 1 \\ 1 & \text{when } \rho = 1 \\ \infty & \text{when } \rho > 1 \end{cases}$$

with the continuous result

$$\lim_{T \rightarrow \infty} (\rho_T)^T = e^c \text{ when } \rho_T = 1 + c/T.$$

used in this article, the inclusion of extra lags has no effect on the limiting distribution, and in this sense the presentation here is without loss of generality. Following the discussion of the limiting distribution of the Chow test statistic, Appendix A2 discusses the numerical procedure used to compute the Prob-values shown in Tables 2 and 3.

A1 Asymptotic Distribution in the AR(1) Model

Assume

$$u_t = \rho_T u_{t-1} + \varepsilon_t,$$

where u_0 is a finite fixed constant, $t = 1, \dots, T$, and ε_t is a martingale difference sequence with $E(\varepsilon_t^2 | \varepsilon_{t-1}, \varepsilon_{t-2}, \dots) = 1$, and with $\sup_t E\varepsilon_t^4 < \infty$, where $\rho_T = 1 + \frac{c}{T}$.

Let $\hat{\rho}_1$ denote the OLS estimator of ρ constructed from the regression of x_t onto $(1, x_{t-1})$ using the early sample period $t = 1, \dots, T_1$, and let $\hat{\rho}_2$ denote the corresponding estimator constructed using the later sample period $t = T_2, \dots, T$. Assume

$$\lim_{T \rightarrow \infty} \frac{T_1}{T} = \tau_1$$

and

$$\lim_{T \rightarrow \infty} \frac{T_2}{T} = \tau_2$$

with $0 < \tau_1 < \tau_2 < 1$. Denote the sample means by

$$\bar{x}_{1,T} = \frac{1}{T_1} \sum_{t=1}^{T_1} x_t$$

$$\bar{x}_{2,T} = \frac{1}{T - T_2 + 1} \sum_{t=T_2}^T x_t$$

and the demeaned series by

$$x_{1,t}^\mu = x_t - \bar{x}_{1,T}$$

$$x_{2,t}^\mu = x_t - \bar{x}_{2,T}.$$

The limiting behavior of these series is related to the behavior of the diffusion process $J_c(s)$, generated by

$$dJ_c(s) = cJ_c(s)ds + dW(s)$$

for $0 \leq s \leq 1$, where $W(s)$ is a standard Wiener process. In particular,

$$\frac{1}{\sqrt{T}} x_{1,[sT]}^\mu \Rightarrow J_c(s) - \tau_1^{-1} \int_0^{\tau_1} J_c(r)dr \equiv J_{1,c}^\mu(s) \text{ for } 0 < s \leq \tau_1$$

$$\frac{1}{\sqrt{T}}x_{2,[sT]}^\mu \Rightarrow J_c(s) - (1 - \tau_2)^{-1} \int_{\tau_2}^1 J_c(r)dr \equiv J_{2,c}^\mu(s) \text{ for } \tau_2 \leq s < 1.$$

The Chow F-statistic for testing $H_0 : \rho_1 = \rho_2$ is

$$F = \frac{(\hat{\rho}_1 - \hat{\rho}_2)^2}{[\sum_{t=1}^{T_1} (x_{1,t-1}^\mu)^2]^{-1} + [\sum_{t=T_2}^T (x_{2,t-1}^\mu)^2]^{-1}}.$$

The limiting behavior follows from considering the terms

$$U_{1,T} \equiv \frac{1}{T} \sum_{t=1}^{T_1} \varepsilon_t x_{1,t-1}^\mu \Rightarrow \int_0^{\tau_1} J_{1,c}^\mu(s) dW(s) \equiv U_1$$

$$U_{2,T} \equiv \frac{1}{T} \sum_{t=T_2}^T \varepsilon_t x_{2,t-1}^\mu \Rightarrow \int_{\tau_2}^1 J_{2,c}^\mu(s) dW(s) \equiv U_2$$

$$V_{1,T} \equiv \frac{1}{T^2} \sum_{t=1}^{T_1} (x_{1,t-1}^\mu)^2 \Rightarrow \int_0^{\tau_1} (J_{1,c}^\mu(s))^2 ds \equiv V_1$$

$$V_{2,T} \equiv \frac{1}{T^2} \sum_{t=T_2}^T (x_{2,t-1}^\mu)^2 \Rightarrow \int_{\tau_2}^1 (J_{2,c}^\mu(s))^2 ds \equiv V_2.$$

Defining

$$\gamma_{1,T} = T(\hat{\rho}_1 - \rho) \text{ and } \gamma_{2,T} = T(\hat{\rho}_2 - \rho),$$

the F can be written as

$$F = \frac{(\gamma_{1,T} - \gamma_{2,T})^2}{V_{1,T}^{-1} + V_{2,T}^{-1}}.$$

Since

$$\gamma_{1,T} = \frac{\frac{1}{T} \sum_{t=1}^{T_1} \varepsilon_t x_{1,t-1}^\mu}{V_{1,T}} = \frac{U_{1,T}}{V_{1,T}} \Rightarrow \frac{U_1}{V_1} = \gamma_1$$

and

$$\gamma_{2,T} = \frac{\frac{1}{T} \sum_{t=T_2}^T \varepsilon_t x_{2,t-1}^\mu}{V_{2,T}} = \frac{U_{2,T}}{V_{2,T}} \Rightarrow \frac{U_2}{V_2} = \gamma_2$$

by the continuous mapping theorem, then

$$F \Rightarrow \frac{(\gamma_1 - \gamma_2)^2}{V_1^{-1} + V_2^{-1}},$$

which provides a representation for the limiting distribution of F in terms of functionals of the diffusions $J_c(s)$.

A2 Approximating Prob-values

The limiting distribution of F is seen to depend on three parameters τ_1 , τ_2 (through the limits in the integrals), and the value of c (through the mean reversion in the diffusion process J_c). Quantiles of the limiting distribution (and hence Prob-values for the test statistic) can be approximated by repeated simulations of F using a large sample size and for fixed values of τ_1 , τ_2 , and c , and ε_t chosen as $Niid(0, 1)$ random variables. The Prob-values reported in the article resulted from 10,000 replications from a sample size of 500. The parameters τ_1 and τ_2 were chosen as T_1/T and T_2/T , where T_1 denotes the first break point and T_2 denotes the second break point. The distribution also depends on c , which governs how close ρ is to unity. Unfortunately, this parameter cannot be consistently estimated. (Equivalently, in finite samples the distribution of F depends on ρ , and small changes in ρ —like those associated with sampling error—lead to large changes in the quantiles of this distribution.) Thus, selecting the correct distribution of F requires knowledge of c (equivalently, ρ). Since c is unknown, the distribution is computed for a range of values in $-25 \leq c \leq 10$ and the resulting minimum and maximum Prob-value over all of the values of c is reported in the table. Viewing c as unknown, classical approaches (which must hold for all values of the “nuisance parameter” c) would use the upper Prob-value. The lower bound gives the smallest Prob-value that would be obtained if c were known.

REFERENCES

- Cavanagh, C. L. “Roots Local to Unity.” Manuscript. Department of Economics, Harvard University, 1985.
- Christiano, Lawrence J., and Martin Eichenbaum. “Unit Roots in Real GNP—Do We Know, and Do We Care?” *Carnegie-Rochester Conference Series on Public Policy*, vol. 32 (1990), pp. 7–32.
- Clarida, Richard, Jordi Galí, and Mark Gertler. “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*, vol. 115 (February 2000), pp. 147–80.
- Deaton, Angus S. “Life-Cycle Models of Consumption: Is the Evidence Consistent with the Theory?” in Truman F. Bewley, ed., *Advances in Econometrics*, Vol II. Amsterdam: North Holland, 1987, pp. 121–48.
- Elliott, Graham. “On the Robustness of Cointegration Methods when Regressors Almost Have Unit Roots,” *Econometrica*, vol. 66 (January 1998), pp. 149–58.

- Goodfriend, Marvin. "Information-Aggregation Bias," *American Economic Review*, vol. 82 (June 1992), pp. 508–19.
- Mehra, Yash P. "Monetary Policy and Long-Term Interest Rates," Federal Reserve Bank of Richmond *Economic Quarterly*, vol. 82 (Summer 1996), pp. 27–29.
- Phillips, P. C. B. "Toward a Unified Asymptotic Theory for Autoregression," *Biometrika*, vol. 74 (1987), pp. 535–47.
- Quah, Danny. "The Relative Importance of Permanent and Transitory Components: Identification and Some Theoretical Bounds," *Econometrica*, vol. 60 (January 1992), pp. 107–18.
- Shiller, Robert J. "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review*, vol. 71 (June 1981), pp. 421–36.
- _____, John Y. Campbell, and Kermit L. Schoenholitz. "Forward Rates and Future Policy: Interpreting the Term Structure of Interest Rates," *Brookings Papers on Economic Activity*, 1:1983, pp. 173–223.
- Stock, James H. "Confidence Intervals for the Largest Autoregressive Root in U.S. Macroeconomic Time Series," *Journal of Monetary Economics*, vol. 28 (December 1991), pp. 435–59.
- Valkanov, Rossen. "Notes on Watson's 'Explaining the Increased Variability in Long-Term Interest Rates.'" Research Memorandum. Princeton University, 1999.
- _____. "The Term Structure with Highly Persistent Interest Rates." Manuscript. Princeton University, 1998.
- West, Kenneth D. "Dividend Innovations and Stock Price Variability," *Econometrica*, vol. 56 (January 1988), pp. 37–61.
- Woodford, Michael. "Policy Inertia." Manuscript. Princeton University, 1999.