Journal of Empirical Finance 8 (2001) 403–426
www.elsevier.com/locate/econbase

Does an intertemporal tradeoff between risk and return explain mean reversion in stock prices?

Chang-Jin Kim [a], James C. Morley [b,*], Charles R. Nelson [c]

[a] Department of Economics, Korea University, South Korea
[b] Department of Economics, Washington University, Campus Box 1208, One Brookings Drive, St. Louis, MO 63130-4899, USA
[c] Department of Economics, University of Washington, 302 Savery Hall, Box 353330, Seattle, WA 98195-3330, USA

Accepted 8 June 2001

Abstract

When volatility feedback is taken into account, there is strong evidence of a positive tradeoff between stock market volatility and expected returns on a market portfolio. In this paper, we ask whether this intertemporal tradeoff between risk and return is responsible for the reported evidence of mean reversion in stock prices. There are two relevant findings. First, price movements not related to the effects of Markov-switching market volatility are largely unpredictable over long horizons. Second, time-varying parameter estimates of the long-horizon predictability of stock returns reject any systematic mean reversion in favour of behaviour implicit in the historical timing of the tradeoff between risk and return. © 2001 Elsevier Science B.V. All rights reserved.

JEL classification: G12; G14
Keywords: Volatility feedback; Mean reversion; Markov switching; Time-varying parameter

1. Introduction

More than a decade has passed since Fama and French (1988) and Poterba and Summers (1988) reported that price movements for market portfolios of common

* Corresponding author. Tel.: +1-314-935-4437; fax: +1-314-935-4156.
E-mail address: morley@wueconc.wustl.edu (J.C. Morley).

PII: S0927-5398(01)00034-2

stocks tend to be at least partially offset over long horizons. This behaviour, labeled “mean reversion,” runs contrary to the random walk hypothesis of stock prices. Subsequent studies by Richardson and Stock (1989), Kim et al. (1991), and Richardson (1993) have challenged the statistical significance of the mean reversion evidence. However, as Summers (1986) points out, statistical tests employed in studies of the random walk hypothesis are ultimately constrained by their low power against mean-reverting alternatives. Thus, point estimates, which according to Fama and French (p. 247) imply “25–45% of the variation of 3–5-year stock returns is predictable from past returns,” may still require some behavioural explanation.

In this paper, we take the economic magnitude of the reported evidence of mean reversion at face value and ask whether it can be explained by an intertemporal tradeoff between risk and return for the stock market as a whole. We add empirical content to our explanation by limiting our definition of risk to the general level of market volatility. Specifically, we consider an empirical model of the tradeoff between market volatility and expected returns on a market portfolio, originally due to Turner et al. (1989) (hereafter, “the TSN model”). The model uses Markov-switching regimes to capture the effects of large changes in market volatility. A Markov-switching specification of market volatility has been used elsewhere, including Schwert (1989a), Schaller and van Norden (1997), and Maheu and McCurdy (2000). However, the formulation used in Turner et al. has the distinctive feature that it implicitly accounts for volatility feedback in measuring the intertemporal tradeoff between risk and return.

Volatility feedback is the idea that an unanticipated change in the level of market volatility will have an immediate impact on stock prices as investors react to new information about future discounted expected returns. In particular, if the level of market volatility is persistent, then the current price index and future discounted expected returns should move in opposite directions. Thus, it is important to account for volatility feedback in order to avoid obscuring any underlying positive tradeoff between market volatility and expected returns. As discussed in Kim et al. (2001), estimates for the TSN model provide strong support for both volatility feedback and a positive tradeoff between market volatility and expected returns. French et al. (1987) and Campbell and Hentschel (1992) also find similar results for alternative specifications of market volatility. The conclusion of these papers is that the predictable behavioural response of risk-averse investors to large changes in market volatility explains a statistically significant portion of stock price movements. The question here, then, is whether these movements are responsible for the reported evidence of mean reversion in stock prices.

To test this explanation of the mean reversion evidence, we incorporate the TSN model of the intertemporal tradeoff between risk and return directly in a regression test of the long-horizon predictability of stock returns due to Jegadeesh (1991) (hereafter, “the Jegadeesh test”). This allows us to estimate the remaining long-horizon predictability of stock price movements that are not directly related to the effects of Markov-switching market volatility. Using CRSP data on value-weighted and equal-weighted portfolios of all NYSE-listed stocks over the period of 1926–1996, we find that the residual unexplained price movements display much less long-horizon predictability than overall returns. For example, the largest estimate of long-horizon predictability for the value-weighted portfolio is reduced from implying 4-year mean reversion of more than 30% to implying 4-year mean reversion of only 2%. Meanwhile, the estimates of long-horizon predictability are always statistically insignificant at conventional levels.

To verify our explanation of the mean reversion evidence, we develop a time-varying parameter version of the Jegadeesh test. This allows us to estimate changes in long-horizon predictability of stock returns over the 1926–1996 sample. Such changes are of interest because of the historical timing of the tradeoff between risk and return implied by estimates of stock market volatility. In particular, the probability inferences from the TSN model depict seemingly periodic 3–4-year volatility regime shifts during the 1930s and 1940s, followed by much less regular regime shifts during the postwar period. Meanwhile, a broadly similar historical pattern for market volatility is portrayed in classic studies by Officer (1976), French et al. (1987), and Schwert (1989a,b). Thus, given our finding for the previous test that unexplained price movements are largely unpredictable over long horizons, only 1930s and 1940s price movements should be consistent with mean reversion over 3–4-year horizons. By contrast, postwar price movements should be more consistent with the random walk hypothesis than with mean reversion. This is precisely what we find with the time-varying parameter estimates of the long-horizon predictability of stock returns. Specifically, estimates reflect both an apparent tendency for price movements to be offset over 3–4-year horizons during the 1930s and 1940s and the disappearance of any such tendency during the postwar period. This finding provides further support for our explanation of the reported mean reversion evidence and argues against any systematic mean reversion due to fads. It also explains why Fama and French (1988), Poterba and Summers (1988), and Kim et al. (1991) find that the mean reversion evidence is extremely sensitive to the inclusion of 1930s and 1940s data in estimation.

Our regression approach can be contrasted with the simulation approach typically used to test explanations of the mean reversion evidence. For instance, Cecchetti et al. (1990) use the simulation approach to demonstrate that the mean reversion evidence is consistent with a general equilibrium model of stock prices with a Markov-switching endowment process and consumption smoothing. Meanwhile, the statistical literature on mean reversion, including Richardson and Stock (1989), Kim et al. (1991), Richardson (1993) and, more recently, Kim and Nelson (1998) and Kim et al. (1998), can be thought of as using the simulation approach to demonstrate that the reported mean reversion evidence is consistent with the random walk hypothesis. By contrast, we use the regression approach to reach a stronger conclusion: the intertemporal tradeoff between risk and return not only can explain mean reversion in stock prices, but it actually does.

The rest of this paper is organized as follows. Section 2 presents the details behind the regression tests employed in this paper. Section 3 reports empirical results for monthly data from CRSP. Section 4 concludes.

2. Tests of long-horizon predictability

Following Jegadeesh (1991), we employ regression tests of the long-horizon predictability of stock returns in which the dependent variable is a 1 month return, the independent variable is a lagged multiple month return, and the coefficient on the lagged return is negative given mean reversion. Richardson and Smith (1994) demonstrate the similarity of the Jegadeesh test to the overlapping autoregression test used in Fama and French (1988) and the variance ratio test used in Poterba and Summers (1988). They show that each test statistic can be represented as a weighted-average of sample autocorrelations for returns, with the difference between each statistic being the weights. Given the rough equivalence of these three tests, there are two reasons why we choose the Jegadeesh test in particular. First, Jegadeesh shows that, within a class of regression tests that also includes the overlapping autoregression test, his test has the highest asymptotic power against Summers' (1986) fads model of mean reversion.[1] Second, implementation of the Kalman filter for time-varying parameter analysis is most straightforward for the Jegadeesh test since it avoids the imposed MA error structure of an overlapping autoregression.

We consider three variations on the Jegadeesh test. Each variation is designed to address a different issue and can be thought of in terms of what is assumed about the specification of the mean return and coefficient in the regression equation. The first variation provides a formal benchmark for our extensions and has a constant mean and fixed coefficient specification. The second variation tests for the long-horizon predictability of stock price movements that are not directly related to the effects of Markov-switching market volatility and has a time-varying mean and fixed coefficient specification. The third variation tests for changes in the long-horizon predictability of stock returns and has a constant mean and time-varying coefficient specification.

[1] See Jegadeesh (1991) for details. Briefly, he uses the approximate slope criterion to determine the optimal aggregation intervals for the dependent and independent variables in terms of power against mean reversion. For Summers' (1986) fads model of mean reversion with a variety of parameter values, Jegadeesh finds that the optimal aggregation interval for the dependent variable is always 1 month.

2.1. Constant mean and fixed coefficient specification

The first variation is actually the original specification employed in Jegadeesh (1991). The regression equation is given as follows:

$$r_t - \mu = \beta(k) \sum_{j=1}^{k} (r_{t-j} - \mu) + \varepsilon_t, \tag{1}$$

where r_t is the 1-month continuously compounded return on a market portfolio, μ is the mean of r_t, k is the holding period (in months) for lagged returns, and ε_t is a serially uncorrelated error term. Under the null hypothesis H_0: $\beta(k) = 0$, sometimes referred to as the random walk hypothesis, the market return is serially uncorrelated with constant expected value μ.[2] Under the alternative hypothesis H_A: $\beta(k) \neq 0$, the market return is predictable using past returns, with $\beta(k) < 0$ corresponding to mean reversion. Note that the coefficient is most comparable to the regression coefficient for a $k/2$-month overlapping autoregression since both reflect an almost identical set of sample autocorrelations (Richardson and Smith, 1994). Therefore, given that previous studies find the strongest evidence of mean reversion for 2–5-year overlapping autoregressions, we should expect to find the strongest evidence of mean reversion for holding periods between 4 and 10 years. Since we take the reported evidence of mean reversion at face value, we intentionally focus on holding periods in this range, even though this stacks the evidence in favour of finding long-horizon predictability (Richardson, 1993). For the same reason, we purposely do not adjust reported estimates for a negative small sample bias in $\beta(k)$ under the null hypothesis (Jegadeesh, 1991). For estimation, we use ordinary least squares (OLS).

2.2. Time-varying mean and fixed coefficient specification

The second variation extends the Jegadeesh test by nesting the TSN model under the null hypothesis H_0: $\beta(k) = 0$. The regression equation is given as follows:

$$r_t - \mu_t = \beta(k) \sum_{j=1}^{k} (r_{t-j} - \mu_{t-j}) + \varepsilon_t, \tag{2}$$

where ε_t has a two-state Markov-switching variance:

$$\varepsilon_t \sim \text{i.i.d.}\, N(0, \sigma_{\varepsilon t}^2), \sigma_{\varepsilon t}^2 = \sigma_{\varepsilon 0}^2 (1 - S_t) + \sigma_{\varepsilon 1}^2 S_t, \sigma_{\varepsilon 1}^2 > \sigma_{\varepsilon 0}^2,$$
$$S_t = \{0,1\}, \Pr[S_t = 0 \mid S_{t-1} = 0] = q, \Pr[S_t = 1 \mid S_{t-1} = 1] = p.$$

[2] We refer to the general version of the random walk hypothesis that allows a positive drift and time dependence for higher moments, including the variance.

That is, the conditional variance $\sigma_{\varepsilon t}^2$ switches between "high" and "low" volatility regimes according to an unobserved Markov-switching state variable S_t with transition probabilities q and p.[3] The time-varying mean μ_t has the following three terms:

$$\mu_t = \mu_0 + \mu_1 \Pr[S_t = 1 \mid r_{t-1}, r_{t-2}, \dots] \\ + \delta(S_t - \Pr[S_t = 1 \mid r_{t-1}, r_{t-2}, \dots]), \tag{3}$$

where the parameters μ_0 and μ_1 and the conditional probability $\Pr[S_t = 1 \mid r_{t-1}, r_{t-2}, \dots]$ determine the expected return under the null hypothesis H_0: $\beta(k) = 0$, while the parameter δ and the revision between the true state and its conditional probability $(S_t - \Pr[S_t = 1 \mid r_{t-1}, r_{t-2}, \dots])$ determine a volatility feedback effect, discussed below. Note that the specification of Eq. (3) represents a simple linear transformation of the "learning" model developed in Turner et al. (1989). For estimation, we use maximum likelihood and an extended version of the filter for Markov switching presented in Hamilton (1989). Details can be found in Appendix A.1.

The components of the time-varying mean μ_t warrant further discussion. First, consider the expected return component: $\mu_0 + \mu_1 \Pr[S_t = 1 \mid r_{t-1}, r_{t-2}, \dots]$. We assume the expected return is a simple linear function of expectations about level of market volatility. Thus, given a positive tradeoff between volatility and expected return, both μ_0 and μ_1 should be positive. That is, a positive and increasing conditional expectation of the level of volatility should correspond to a positive and increasing return, ceteris paribus. Note that we avoid imposing a strict proportionality on the relationship between expected return and expected volatility. In particular, the marginal impact of an increase in the expectation of market volatility can be different from the overall impact of having a positive level of volatility. Second, consider the volatility feedback component: $\delta(S_t - \Pr[S_t = 1 \mid r_{t-1}, r_{t-2}, \dots])$. Volatility feedback can arise whenever investors acquire new information about volatility. In this paper, we follow Turner et al. (1989) and proxy this new information by the difference between the true unobserved volatility regime S_t and its conditional probability $\Pr[S_t = 1 \mid r_{t-1}, r_{t-2}, \dots]$. Then, if volatility regimes are persistent (i.e., the sum of the transition probabilities is greater than one: $q + p > 1$), the new information embodied in the revision term $(S_t - \Pr[S_t = 1 \mid r_{t-1}, r_{t-2}, \dots]) \neq 0$ produces a corresponding change in the dis-

[3] Numerous other studies, including Schwert (1989a), Schaller and van Norden (1997), Mayfield (1999), and Maheu and McCurdy (2000) have used Markov switching to capture large changes in market volatility. The best justification for Markov-switching volatility, however, comes from a paper by Hamilton and Susmel (1994). They develop a Markov-switching autoregressive conditional heteroskedasticity (SWARCH) model of weekly stock returns. From their results, it appears that, once Markov-switching regime changes are accounted for, most ARCH effects die out at the monthly return horizon considered in this paper.

counted sum of future expected returns on the market portfolio. From Campbell and Shiller's (1988) log-linear approximate present-value identity, this change in the discounted sum of future expected returns is equivalent to an opposite movement in the market price index. Thus, given a positive tradeoff between volatility and expected returns, the volatility feedback coefficient δ should be negative. That is, news about higher future volatility should correspond to an immediate decline in the price index, producing a lower return, ceteris paribus. Note that the volatility feedback effect δ should be easier to detect than the partial effect μ_1 since volatility feedback embodies a change in the discounted sum of all future expected returns. Other studies that consider volatility feedback include French et al. (1987) and Campbell and Hentschel (1992). The algebraic derivation of a formal model of volatility feedback under the assumption of Markov-switching volatility is provided in Kim et al. (2001).

2.3. Constant mean and time-varying coefficient specification

The third variation extends the Jegadeesh test by allowing the long-horizon predictability of stock returns to change over time. The regression equation is given as follows:

$$r_t - \mu = \beta_t(k) \sum_{j=1}^{k} (r_{t-j} - \mu) + \varepsilon_t, \tag{4}$$

where ε_t is a serially uncorrelated error term. To identify the time-varying coefficient $\beta_t(k)$, we must impose structure on its evolution. In this paper, we choose a random walk process:

$$\beta_t(k) = \beta_{t-1}(k) + v_t, \tag{5}$$

where v_t is a serially uncorrelated error term, independent of ε_t. Garbade (1977) and Engle and Watson (1987) argue that a random walk provides a good empirical model of the univariate behaviour of regression coefficients in many situations by allowing for permanent changes in regression coefficients. At the same time, it is fairly robust to misspecification.[4] The random walk process also allows a constant coefficient as a special case when the variance of v_t collapses to zero. For estimation, we use maximum likelihood and the Kalman filter. Details can be found in Appendix A.2.

[4] See Garbade (1977) for a Monte Carlo investigation of the consequences of misspecification. Briefly, he shows that a random walk process detects parameter instability even when the truth is either a one-time discrete jump in the parameter or a persistent, but stationary, first-order autoregressive process. In addition, he points out that the graphical representations of parameter estimates tend to reflect the true nature of parameter instability, not just its presence.

As with the constant coefficient cases, $\beta_t(k) \neq 0$ could reflect either fads or predictable changes in equilibrium expected returns. However, allowing $\beta_t(k)$ to change over time is important because it potentially allows us to discriminate between the two cases. In particular, if equilibrium expected returns are related to the level of market volatility, then the historical pattern for market volatility portrayed in Officer (1976), French et al. (1987), and Schwert (1989a,b) implies that $\beta_t(k)$ will change during the postwar period to reflect a disappearance in the apparent long-horizon predictability of stock returns evident in the 1930s and 1940s. By contrast, if fads are responsible for any systematic mean reversion in stock prices, then there is no implication that $\beta_t(k)$ will change over time and $\beta_t(k) \neq 0$ should strictly hold throughout the postwar period.

3. Empirical results

3.1. Data

To test for long-horizon predictability, we use stock return data from the CRSP file. The data, available for the sample period of January 1926 to December 1996, are the total monthly returns on the value-weighted portfolio and on the equal-weighted portfolio of all NYSE-listed stocks, where "total" denotes capital gain plus dividend yield as calculated by CRSP. Following Fama and French (1988), we deflate nominal returns by the monthly CPI (not seasonally adjusted) for all urban consumers from Ibbotson Associates to get a measure of real returns. We convert to continuously compounded returns by taking natural logarithms of simple gross returns.

3.2. Constant mean and fixed coefficient results

Table 1 reports OLS estimates for the constant mean and fixed coefficient specification and holding periods of 48, 72, 96, and 120 months.[5] The results confirm what has been previously reported in the literature and are reported here to provide a benchmark for our extensions. First, for the full 1926–1996 sample, the reported economic magnitude of long-horizon stock return predictability is large in most cases, although the estimates are only statistically significant at conventional levels in a few cases. To help think about the "economic magnitude" of the reported estimates, consider the implied mean reversion of a price shock

[5] All OLS estimates were calculated in EViews. Following Jegadeesh (1991), we use White's (1980) heteroskedasticity-consistent standard errors.

Table 1
Constant mean and fixed coefficient specification: OLS estimates and Wald breakpoint test, 1926–1996

k	$\beta(k)$ 1926–1996	$\beta(k)$ 1926–1946	$\beta(k)$ 1947–1996	Wald breakpoint statistic	$\chi^2(1)$ p-value
Value-weighted portfolio					
48	−0.0076 (0.0094)	−0.0150 (0.0152)	−0.0015 (0.0074)	0.6375	0.4246
72	−0.0064 (0.0064)	−0.0338 (0.0175)**	0.0030 (0.0060)	3.9506	0.0469
96	0.0000 (0.0046)	−0.0269 (0.0221)	0.0031 (0.0044)	1.7719	0.1831
120	0.0025 (0.0038)	−0.0013 (0.0199)	0.0023 (0.0037)	0.0312	0.8598
Equal-weighted portfolio					
48	−0.0096 (0.0098)	−0.0102 (0.0142)	−0.0086 (0.0093)	0.0085	0.9265
72	−0.0172** (0.0089)	−0.0324 (0.0168)**	−0.0043 (0.0079)	2.2920	0.1300
96	−0.0102* (0.0062)	−0.0373 (0.0181)**	−0.0034 (0.0060)	3.1672	0.0751
120	−0.0023 (0.0057)	−0.0021 (0.0250)	−0.0026 (0.0056)	0.0004	0.9840

Estimates are calculated using the continuously compounded total monthly real return on the value-weighted portfolio and the equal-weighted portfolio of all NYSE-listed stocks. Data are available for the period of January 1926 to December 1996, with sample periods adjusted to account for lagged variables. White's (1980) heteroskedasticity-consistent standard errors are reported in parentheses and are used to calculate the Wald statistics for a breakpoint in the coefficient in 1947.

* t-statistic for H_0: $\beta(k)=0$ is significant at 10% level.

** t-statistic for H_0: $\beta(k)=0$ is significant at 5% level.

over a 4-year horizon (see Appendix A.3 for calculation details). For example, the statistically insignificant point estimate of -0.0076 for the value-weighted portfolio ($k=48$) implies 4-year mean reversion of as much as 30%. Meanwhile, the statistically significant point estimate of -0.0172 for the equal-weighted portfolio ($k=72$) implies 4-year mean reversion of over 55%.[6] Second, the reported economic magnitude is extremely sensitive to the sample period. For example, the point estimate of -0.0373 for the equal-weighted portfolio ($k=96$) implies 4-year mean reversion of as much as 85% for the 1926–1946 sample period, while the corresponding point estimate of -0.0034 implies 4-year mean reversion of only 15% for the 1947–1996 sample period. The Wald statistics for a breakpoint in January 1947 suggest that a postwar reduction in the economic magnitude of long-horizon predictability is statistically significant for both the value-weighted portfolio ($k=72$) and equal-weighted portfolio ($k=96$).

[6] While we purposely shy away from explicitly defining a particular threshold level of mean reversion that should be considered economically "large," we believe that 4-year mean reversion between 30% and 55% warrants the description.

3.3. Time-varying mean and fixed coefficient results

Table 2 reports the maximum likelihood estimates for the time-varying mean and fixed coefficient specification and holding periods of 48, 72, 96, and 120 months.[7] The first thing to notice is that, when we account for the intertemporal tradeoff between risk and return, the reported economic magnitude of the residual unexplained long-horizon predictability is much smaller than before. The point estimate of -0.0004 for the value-weighted portfolio $(k=48)$ implies 4-year mean reversion of only 2%, compared to 30% before. Meanwhile, the point estimate of -0.0049 for the equal-weighted portfolio $(k=72)$ implies 4-year mean reversion of 20%, compared to over 55% before. Furthermore, the estimates of unexplained long-horizon predictability are all statistically insignificant at conventional levels.

Testing the Markovian specification of regime switching is hindered by the failure of several assumptions of standard asymptotic distribution theory. Notably, as discussed in Hansen (1992) and Garcia (1998), the transition probabilities q and p are not identified under a null hypothesis of a constant mean and variance. Since the distribution of test statistics are model and data dependent, Hansen argues for the use of computationally intensive simulations to determine the small sample distributions. Garcia, however, derives asymptotic distributions of a likelihood ratio test for different two-state Markov-switching models. The largest asymptotic critical value he reports is 17.52, corresponding to a 1% significance level for a test of a two-state Markov-switching mean and variance model with an uncorrelated and heteroskedastic noise function. If we use this critical value as a rough guide, regime switching appears to be quite significant for stock returns. Likelihood ratio statistics for the null hypothesis H_0: $\sigma_{\varepsilon 0}^2 = \sigma_{\varepsilon 1}^2$, $\mu_1 = \gamma = 0$ of no regime switching range between 123.06 and 288.94 for the value-weighted portfolio $(k=120$ and $k=48$, respectively) and between 184.58 and 381.24 for the equal-weighted portfolio $(k=120$ and $k=48$, respectively). Furthermore, the strong persistence of the regimes (i.e., $q+p>1$) provides support for the Markovian specification of regime switching (Engel and Hamilton, 1990).

A related issue is whether the two-state Markov-switching variance is sufficient to capture the heteroskedasticity in monthly stock returns. Kim and Nelson (1998) and Kim et al. (1998) suggest that a third volatility regime is necessary, especially when data from the 1930s to 1940s are included in estimation. However, they do not allow for volatility feedback or other changes in the mean. We find that the model given in Eqs. (2) and (3), which does allow for volatility feedback, captures

[7] All maximum likelihood estimation was conducted using the OPTMUM procedure for the GAUSS programming language. Numerical derivatives were used in estimation, as well as for calculation of asymptotic standard errors. Parameters were appropriately constrained (e.g., variances were constrained to be non-negative). Inferences appear robust to a variety of starting values.

Table 2
Time-varying mean and fixed coefficient specification: maximum likelihood estimates, 1926–1996

k	$\beta(k)$ 1926–1996	$\sigma_{\varepsilon 0}$	$\sigma_{\varepsilon 1}$	q	p	μ_0	μ_1	δ	log-likelihood value
Value-weighted portfolio									
48	−0.0004	0.0372	0.1110	0.9931	0.9487	0.0067	−0.0096	−0.0366	1340.69
	(0.0024)	(0.0013)	(0.0086)	(0.0034)	(0.0185)	(0.0023)	(0.0129)	(0.0135)***	
72	−0.0008	0.0369	0.1106	0.9918	0.9282	0.0069	−0.0008	−0.0299	1324.17
	(0.0019)	(0.0014)	(0.0104)	(0.0040)	(0.0252)	(0.0021)	(0.0025)	(0.0117)**	
96	0.0001	0.0340	0.0752	0.9794	0.8941	0.0057	0.0005	−0.0348	1324.39
	(0.0013)	(0.0016)	(0.0071)	(0.0082)	(0.0299)	(0.0026)	(0.0015)	(0.0090)***	
120	0.0007	0.0335	0.0752	0.9783	0.8878	0.0038	0.0052	−0.0338	1289.43
	(0.0014)	(0.0015)	(0.0071)	(0.0085)	(0.0327)	(0.0038)	(0.0107)	(0.0097)***	
Equal-weighted portfolio									
48	−0.0009	0.0441	0.1478	0.9926	0.9536	0.0088	−0.0044	−0.0297	1167.82
	(0.0030)	(0.0016)	(0.0111)	(0.0033)	(0.0166)	(0.0024)	(0.0176)	(0.0163)*	
72	−0.0049	0.0441	0.1511	0.9922	0.9386	0.0091	0.0105	−0.0146	1157.86
	(0.0040)	(0.0017)	(0.0136)	(0.0036)	(0.0219)	(0.0024)	(0.0153)	(0.0171)	
96	−0.0026	0.0378	0.0950	0.9771	0.9296	0.0083	0.0026	−0.0296	1164.86
	(0.0036)	(0.0029)	(0.0090)	(0.0095)	(0.0233)	(0.0032)	(0.0092)	(0.0112)***	
120	−0.0005	0.0370	0.0921	0.9756	0.9234	0.0075	0.0022	−0.0347	1143.97
	(0.0012)	(0.0021)	(0.0071)	(0.0089)	(0.0240)	(0.0028)	(0.0161)	(0.0114)***	

Estimates are calculated using the continuously compounded total monthly real return on the value-weighted portfolio and the equal-weighted portfolio of all NYSE-listed stocks. Data are available for the period of January 1926 to December 1996, with sample periods adjusted to account for lagged variables. Asymptotic standard errors based upon second derivatives are reported in parentheses.

* t-statistics for H_0: $\delta = 0$ is significant at 10% level.

** t-statistics for H_0: $\delta = 0$ is significant at 5% level.

*** t-statistics for H_0: $\delta = 0$ is significant at 1% level.

most of the negative skewness and leptokurtosis in stock returns. In particular, the standardized residuals for Eq. (2) display much less heteroskedasticity than the residuals for Eq. (1).[8] The Jarque and Bera (1980) $\chi^2(2)$ test statistics based on the third and fourth sample moments fall from 734.57, 853.65, 215.26, and 214.58 to only 16.05, 14.84, 2.28, and 2.61 for the value-weighted portfolio ($k=48$, $k=72$, $k=96$, and $k=120$, respectively) and from 1048.39, 1183.70, 361.89, and 455.28 to only 10.96, 12.67, 19.16, and 28.27 for the equal-weighted portfolio ($k=48$, $k=72$, $k=96$, and $k=120$, respectively). Meanwhile, further addressing heteroskedasticity by adding a third volatility regime should only weaken the evidence of mean reversion for the unexplained price movements. Kim and Nelson (1998) and Kim et al. (1998) find that the evidence of mean reversion for overall stock returns is much weaker when a three-state Markov-switching variance specification is considered. Also, the more flexible the time-varying mean, the more likely it can spuriously explain predictable price movements. Our two-state specification provides the minimal possible flexibility, while still allowing us to test our explanation of the reported mean reversion evidence.

Beyond the unexplained long-horizon predictability and the variance process, the estimates of μ_0, μ_1, and δ, which correspond to the expected return and the volatility feedback effect, are also of considerable interest. Contrary to the findings in Turner et al. (1989), we find that the estimated partial effect μ_1 of an increase in expected volatility is actually positive in most cases, although it is never statistically significant.[9] For the case that corresponds to strongest evidence of mean reversion in Table 1 (i.e., the equal-weighted portfolio with $k=72$), the mean return more than doubles from 0.91% in a perfectly anticipated low volatility regime to 1.96% in a perfectly anticipated high volatility regime. Meanwhile, the estimated volatility feedback effect δ of an unanticipated transition into a high volatility regime on the mean return is always negative, corresponding to a positive tradeoff between volatility and expected returns. The t-statistics for the null hypothesis H_0: $\delta=0$ of no volatility feedback are between -2.56 and -3.87 for the value-weighted portfolio ($k=72$ and $k=96$, respectively) and between -0.85 and -3.04 for the equal-weighted portfolio ($k=72$ and $k=120$, respectively). In terms of economic magnitude, even the smallest point estimates suggest

[8] Since we do not directly observe the true state, we use smoothed, or two-sided, probabilities as the best available estimate of the true state in order to calculate the standardized residuals for Eq. (2). The smoothed probabilities are conditional on all available returns and the maximum likelihood estimates of the hyper-parameters presented in Table 2. As a result of this substitution, our standardized residuals may be less Normal than the true residuals. Also, it should be noted that, while the use of the smoothed probabilities reduces the sample of standardized residuals for Eq. (2), we compare residuals for Eqs. (1) and (2) using the same adjusted sample periods.

[9] Turner et al. (1989) find a negative, but insignificant, partial effect using excess returns from Standard and Poor's composite index for the sample period of January 1946 to December 1987.

that a completely unanticipated transition into a high volatility regime produces an immediate 2.99% decline in the value-weighted portfolio $(k = 72)$ and an immediate 1.46% decline in the equal-weighted portfolio $(k = 72)$.

Fig. 1 displays the filtered and smoothed probabilities of a high volatility regime for holding periods of 48, 72, 96, and 120 months. The filtered, or one-sided, probabilities are conditional on returns up to time t and maximum likelihood estimates of the hyper-parameters presented in Table 2. The smoothed, or two-sided, probabilities are conditional on all available returns and the same maximum likelihood estimates. The main thing to notice about the probabilities is that, for both portfolios $(k = 48$ and $k = 72)$, there are seemingly periodic 3–4-year regime shifts during the 1930s and 1940s. While there are also regime

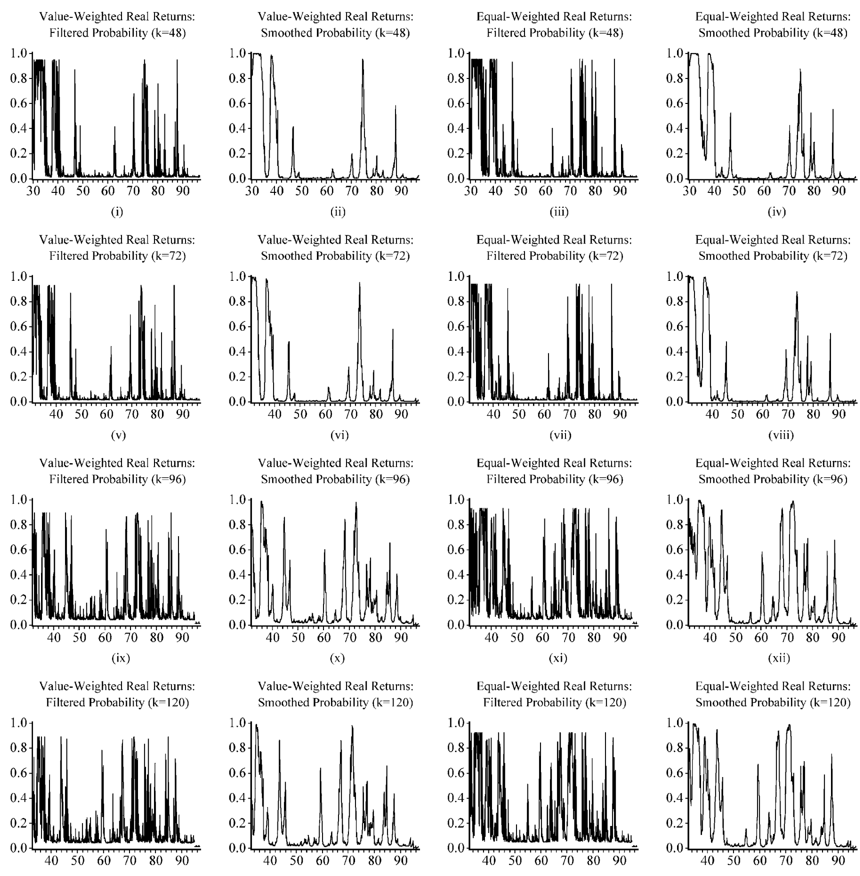

Fig. 1. Filtered and smoothed probabilities of a high volatility regime, 1926–1996.

shifts in the postwar period, they come at much less regular intervals. Given our finding that residual unexplained movements are largely unpredictable over long horizons, this historical pattern of regime changes suggests that volatility regime changes in 1930s and 1940s are responsible for the reported evidence of mean reversion. In particular, the timing of the tradeoff between risk and return in the 1930s and 1940s should produce negative sample autocorrelations at horizons that will be easiest to detect for $k=72$ and $k=96$.[10] Meanwhile, unlike with fads, there should be no systematic mean reversion in this scenario.[11] A dramatic implication, then, is that the apparent long-horizon predictability should disappear in the postwar period as the regime shifts become less regular. We use the constant mean and time-varying coefficient specification to test this implication and verify our explanation of the reported mean reversion evidence.

3.4. Constant mean and time-varying coefficient results

Table 3 reports maximum likelihood estimates for the constant mean and time-varying coefficient specification and holding periods of 48, 72, 96, and 120 months. The estimates for the variance of the time-varying parameter suggest changes in the apparent long-horizon predictability of stock returns. The likelihood ratio statistics for the null hypothesis of constant predictability H_0: $\sigma_v^2 = 0$ are as high as 4.3122 (p-value = 0.038) for the value-weighted portfolio ($k=72$) and 9.3183 (p-value = 0.002) for the equal-weighted portfolio ($k=72$).[12] These results hold in spite of the fact that the maximum likelihood estimate of a time-varying parameter variance has a point mass at zero when the true variance is small (Stock and Watson, 1998). Meanwhile, the estimates are generally consistent with the Wald statistics reported in Table 1, while they avoid the problems, discussed in Zivot and Andrews (1992) and Andrews (1993), associated with an assumption of a known breakpoint.

[10] As discussed previously, Jegadeesh tests with $k=72$ and $k=96$ are comparable to 3- and 4-year overlapping autoregressions (Richardson and Smith, 1994).

[11] Interestingly, given the Markov-switching specification, there should actually be some systematic mean aversion at shorter horizons. Heuristically, conditional on presently being in a high variance regime, past and future returns are likely to be somewhat above average. However, this effect dies out as horizons get longer and conditional expectations revert to unconditional levels.

[12] Although we report standard errors for all the parameters in the tables, we emphasize the likelihood ratio statistic for testing this hypothesis since Garbade (1977) shows that it has good finite sample properties in detecting a variety of forms of parameter instability. To be clear, the likelihood ratio test does not have the highest local asymptotic power against specific forms of parameter instability such as a random walk coefficient (Nyblom, 1989). However, our interest is in more general forms of instability, as well as in the actual time path of the coefficient through time, which the time-varying parameter approach provides. In addition, the asymptotic distribution of the likelihood ratio statistic is concentrated towards the origin under the null hypothesis, making the likelihood ratio test conservative in the sense that reported p-values understate the true level of significance (Garbade, 1977; also see Kendall and Stuart, 1973).

Table 3
Constant mean and time-varying coefficient specification: maximum likelihood estimates, 1926–1996

k	Adjusted sample period	σ_v	σ_ε	μ	Log-likelihood value
Value-weighted portfolio					
48	1934–1996	0.0000 (0.0004)	0.0456 (0.0012)	0.0064 (0.0012)	1261.47
72	1935–1996	0.0012 (0.0007)**	0.0451 (0.0012)	0.0064 (0.0013)	1246.38
96	1936–1996	0.0003 (0.0007)	0.0453 (0.0012)	0.0070 (0.0016)	1224.65
120	1937–1996	0.0000 (0.0001)	0.0453 (0.0012)	0.0043 (0.0021)	1203.76
Equal-weighted portfolio					
48	1934–1996	0.0000 (0.0002)	0.0591 (0.0015)	0.0090 (0.0014)	1064.53
72	1935–1996	0.0014 (0.0008)***	0.0575 (0.0015)	0.0097 (0.0011)	1066.00
96	1936–1996	0.0015 (0.0009)***	0.0572 (0.0015)	0.0095 (0.0008)	1051.18
120	1937–1996	0.0000 (0.0001)	0.0574 (0.0015)	0.0072 (0.0016)	1034.08

Estimates are calculated using the continuously compounded total monthly real return on the value-weighted portfolio and the equal-weighted portfolio of all NYSE-listed stocks. Data are available for the period of January 1926 to December 1996, with sample periods adjusted to account for lagged variables and starting up the Kalman filter. Asymptotic standard errors based upon second derivatives are reported in parentheses.

** Likelihood ratio statistic for H_0: $\sigma_v = 0$ is significant at 5% level.

*** Likelihood ratio statistic for H_0: $\sigma_v = 0$ is significant at 1% level.

Fig. 2 displays the filtered and smoothed inferences about the time-varying parameter. The point estimates of long-horizon predictability start out negative and potentially significant for both portfolios ($k = 72$ and $k = 96$), but are updated in the postwar period to reflect no apparent long-horizon predictability. The robustness of this basic historical pattern across both filtered and smoothed inferences suggests that it is not merely a consequence of initial uncertainty over time-varying parameter values when starting up the Kalman filter. Meanwhile, the filtered inferences suggest that the evidence against mean reversion became overwhelming in the mid-1950s for the value-weighted portfolio and in the mid-1970s for the equal-weighted portfolio. The upward spike in filtered inferences in the mid-1970s corresponds to a brief episode of apparent mean aversion as stock prices fell dramatically in 1974 following a long period of below-average returns. The filtered inferences suggest a possible return of mean reversion for both portfolios following the 1987 stock market crash. However, point estimates remain both statistically insignificant and economically small at the end of the sample.

It should be noted that we do not account for the effects of stock return heteroskedasticity for this specification. However, as shown in Morley (1999), allowing for a Markov-switching variance for ε_t in Eq. (4) only weakens the evidence of mean reversion by putting less weight on the volatile 1930s and 1940s data and enlarging the confidence bands around the point estimates, especially in the early part of the sample. In addition, other studies that account for het-

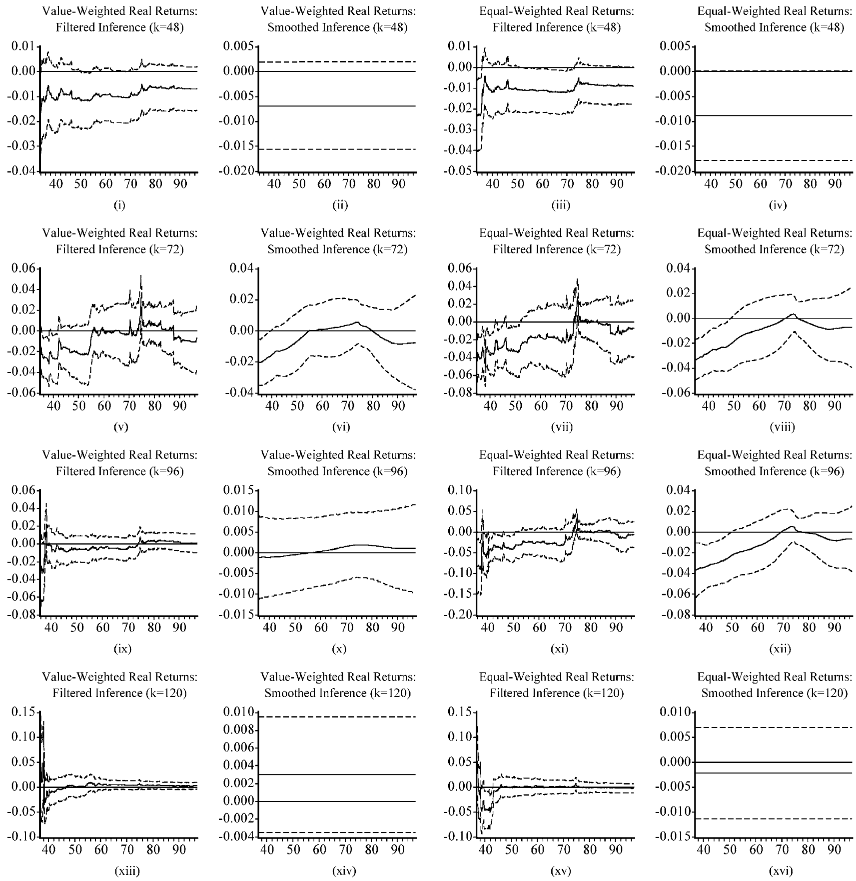

Fig. 2. Filtered and smoothed inferences about the time-varying parameter, 1926–1996 (dashed lines represent 95% confidence bands).

eroskedasticity, including Kim et al. (1991, 1998), McQueen (1992), and Kim and Nelson (1998), find weaker evidence of mean reversion. Thus, our findings would seem to unambiguously argue against any inherent mean reversion.

4. Conclusions

An intertemporal tradeoff between risk and return for the stock market as a whole explains the reported evidence of mean reversion in Fama and French (1988) and Poterba and Summers (1988). In particular, the historical timing of large changes in the level of stock market volatility produced changes in expected returns responsible for the apparent tendency of price movements to be offset over long horizons. Meanwhile, the absence of periodic changes in volatility during the postwar period corresponds to a disappearance of any apparent long-horizon predictability for the postwar data. We arrive at these conclusions in two ways. First, when we consider an empirical model of stock returns that captures volatility feedback in the presence of a positive tradeoff between market volatility and expected returns, we find that residual unexplained price movements are largely unpredictable over long horizons. Second, when we allow the apparent long-horizon predictability of stock returns to change over time with a time-varying parameter model, we find that postwar price movements are more consistent with the behaviour implicit in the historical tradeoff between risk and return than any systematic mean reversion.

We conclude this paper by noting that our results provide strong support for market efficiency. To be clear, we do not provide a decisive test of market efficiency. In particular, had we been unable to explain the reported mean reversion evidence, we could easily have argued that our results were a consequence of a misspecification of equilibrium expected returns, rather than, say, a failure of market efficiency due to a systematic overreaction by investors to news about fundamentals. However, given that our measure of the intertemporal tradeoff between risk and return does explain the reported mean reversion evidence, the argument for a failure of market efficiency due to investor overreaction is largely discredited. First, we find no evidence of opportunities for arbitrage over long horizons. Instead, the optimal forecast appears to be the same as the equilibrium expected return implied by the intertemporal tradeoff between risk and return. Second, the most recent estimates of long-horizon predictability actually support the random walk hypothesis. That is, contrary to the implication of a systematic overreaction by investors, postwar stock returns appear largely unpredictable over long horizons given past returns.

Acknowledgements

We have received helpful comments from Charles Engel, Dick Startz, Eric Zivot, participants in seminars at the Federal Reserve Bank of Dallas, the Federal Reserve Bank of St. Louis, Washington University, and the University of Washington, and two anonymous referees, but responsibility for any errors is entirely our own. Support from the National Science Foundation under grants SBR-9711301 and SES-9818789 and the Ford and Louisa Van Voorhis endowment at the University of Washington is gratefully acknowledged. This paper is based in part on Morley's (1999) doctoral dissertation at the University of Washington. Earlier drafts of the paper circulated under the title "Do Changes in the Market Risk Premium Explain the Empirical Evidence of Mean Reversion in Stock Prices?"

A.1. Estimation of the time-varying mean and fixed coefficient specification

For the specification presented in Eqs. (2) and (4), an extended version of the filter discussed in Hamilton (1989) is given by the following three steps.

Step 1a: Calculate the joint probability of S_t, S_{t-1}, S_{t-k}, and $S_t^* \equiv \sum_{j=2}^{k-1} S_{t-j}$ and solve for $\Pr[S_t=1 \mid r_{t-1}, r_{t-2}, \ldots]$ by summing across all possible values of S_{t-1}, S_{t-k}, and S_t^*:

$$\begin{aligned}&\Pr[S_t=j, S_{t-1}=i, S_{t-k}=h, S_t^*=m \mid r_{t-1}, r_{t-2}, \ldots]\\&\quad=\Pr[S_t=j \mid S_{t-1}=i]\cdot\Pr[S_{t-1}=i, S_{t-k}=h, S_t^*=m \mid r_{t-1}, r_{t-2}, \ldots],\end{aligned}\tag{A.1}$$

$$\begin{aligned}&\Pr[S_t=1 \mid r_{t-1}, r_{t-2}, \ldots]\\&\quad=\sum_{i=0}^{1}\sum_{h=0}^{1}\sum_{m=0}^{k-2}\Pr[S_t=1, S_{t-1}=i, S_{t-k}=h, S_t^*=m \mid r_{t-1}, r_{t-2}, \ldots].\end{aligned}\tag{A.2}$$

Step 1b: Calculate the conditional density of r_t:

$$\begin{aligned}&f(r_t \mid S_t=j, S_{t-1}=i, S_{t-k}=h, S_t^*=m, r_{t-1}, r_{t-2}, \ldots)\\&\quad=\frac{1}{\sqrt{2\pi\sigma_{S_t}^2}}\exp\left\{-\frac{1}{2\sigma_{S_t}^2}\left(r_{S_t}^* - \beta(k)\sum_{j=1}^{k} r_{S_{t-j}}^*\right)^2\right\},\end{aligned}\tag{A.3}$$

where $r_{s_t}^* \equiv r_t - \mu_0 - \mu_1 \Pr[S_t=1 \mid r_{t-1}, r_{t-2}, \ldots] - \delta(S_t - \Pr[S_t=1 \mid r_{t-1}, r_{t-2}, \ldots])$. Note that given data up to time t, S_t, S_{t-1}, S_{t-k}, S_t^*, and particular values for the parameters, we observe all the elements of Eq. (A.3) since

$$\begin{aligned}\sum_{j=1}^{k} r_{S_{t-j}}^* &= \sum_{j=1}^{k} r_{t-j} - \mu_0\cdot k\\&\quad-(\mu_1-\delta)\cdot\sum_{j=1}^{k}\Pr[S_{t-j}=1 \mid r_{t-j-1}, r_{t-j-2}, \ldots] - \delta\sum_{j=1}^{k} S_{t-j}\end{aligned}$$

and

$$\sum_{j=1}^{k} S_{t-j} = S_{t-1} + S_t^* + S_{t-k}.$$

Step 2: Calculate the joint density of r_t, S_t, S_{t-1}, S_{t-k}, and S_t^* and collapse across all possible states to find the marginal density of r_t:

$$\begin{aligned} & f(r_t,S_t=j,S_{t-1}=i,S_{t-k}=h,S_t^*=m \mid r_{t-1},r_{t-2},\ldots) \\ & \quad = f(r_t \mid S_t=j,S_{t-1}=i,S_{t-k}=h,S_t^*=m,r_{t-1},r_{t-2},\ldots) \\ & \qquad \times \Pr[S_t=j,S_{t-1}=i,S_{t-k}=h,S_t^*=m \mid r_{t-1},r_{t-2},\ldots], \end{aligned} \tag{A.4}$$

$$\begin{aligned} & f(r_t \mid r_{t-1},r_{t-2},\ldots) \\ & \quad = \sum_{j=0}^{1}\sum_{i=0}^{1}\sum_{h=0}^{1}\sum_{m=0}^{k-2} f(r_t,S_t=j,S_{t-1}=i,S_{t-k}=h,S_t^*=m \mid r_{t-1},r_{t-2},\ldots). \end{aligned} \tag{A.5}$$

Step 3: Update the joint probability of S_t, S_{t-1}, S_{t-k}, and S_t^* given r_t and solve for the joint probability of S_t, S_{t-k+1}, and S_{t+1}^*:

$$\begin{aligned} & \Pr[S_t=j,S_{t-1}=i,S_{t-k}=h,S_t^*=m \mid r_t,r_{t-1},\ldots] \\ & \quad = \frac{f(r_t,S_t=j,S_{t-1}=i,S_{t-k}=h,S_t^*=m \mid r_{t-1},r_{t-2},\ldots)}{f(r_t \mid r_{t-1},r_{t-2},\ldots)}, \end{aligned} \tag{A.6}$$

$$\begin{aligned} & \Pr[S_t=j,S_{t-k+1}=g,S_{t+1}^*=n \mid r_t,r_{t-1},\ldots] \\ & \quad = \sum_{i=0}^{1}\sum_{h=0}^{1} \Pr[S_t=j,S_{t-1}=i,S_{t-k}=h,S_{t-k+1}=g,S_{t+1}^*=n \mid r_t,r_{t-1},\ldots], \end{aligned} \tag{A.7}$$

where

$$\begin{aligned} & \Pr[S_t=j,S_{t-1}=i,S_{t-k}=h,S_{t-k+1}=g,S_t^*=m \mid r_t,r_{t-1},\ldots] \\ & \quad = \Pr[S_{t-k}=h \mid S_{t-k+1}=g] \\ & \qquad \times \Pr[S_t=j,S_{t-1}=i,S_{t-k}=h,S_t^*=m \mid r_t,r_{t-1},\ldots] \\ & \quad \Rightarrow \Pr[S_t=j,S_{t-1}=i,S_{t-k}=h,S_{t-k+1}=g,S_{t+1}^*=n \mid r_t,r_{t-1},\ldots], \end{aligned}$$

since $S_{t+1}^* = S_{t-1} + S_t^* - S_{t-k+1}$.

Then, given $\Pr[S_{-i^*}=1]$ for $i^*=0,\ldots,k-1$ and $\Pr[S_0=i,S_{-k+1}=h,S_1^*=m]$ for $i=0,1$, $h=0,1$, and $m=0,\ldots,k-1$, we iterate through Eqs. (A.1)–(A.7) for $t=1,\dots,T$ to obtain the filtered probability $\Pr[S_t=1 \mid r_{t-1},r_{t-2},\dots]$. We use the unconditional probabilities for $\Pr[S_{-i^*}=1]$:

$$\Pr[S_{-i^*}=0]=\frac{1-p}{2-p-q}, \tag{A.8}$$

$$\Pr[S_{-i^*}=1]=\frac{1-q}{2-p-q}. \tag{A.9}$$

As for $\Pr[S_0=i,S_{-k+1}=h,S_1^*=m]$, deriving unconditional joint probabilities in terms of q and p is impractical for large values of k. Instead, we treat these initial probabilities, denoted $\Pi=(\pi_1,\dots,\pi_{4(k-1)})$, as $4\times(k-1)$ additional parameters to be estimated.

We use the marginal density of r_t given in Eq. (A.5) to find maximum likelihood estimates of the parameters as follows:

$$\max_{\theta} l(\theta)=\sum_{t=1}^{T}\ln\left(f(r_t \mid r_{t-1},r_{t-2},\dots)\right), \tag{A.10}$$

where $\theta=(\beta,\sigma_{\varepsilon 0},\sigma_{\varepsilon 1},q,p,\mu_0,\mu_1,\delta,\Pi)$.[13]

In addition to making the above inferences, we also obtain the smoothed probability $\Pr[S_t=1 \mid r_T,r_{T-1},\dots]$ by employing Kim's (1994) smoothing algorithm. Specifically, given the filtered probability $\Pr[S_t=j \mid r_t,r_{t-1},\dots]$, which can be found by collapsing across states for Eq. (A.7), and the conditional probability $\Pr[S_t=j \mid r_{t-1},r_{t-2},\dots]$, given in Eq. (A.2), we iterate backwards through the following two equations (conditional on $S_t=j$ and $S_{t+1}=l$, where $j=0,1$ and $l=0,1$):

$$\begin{aligned}&\Pr[S_{t+1}=l,S_{t+1}=j \mid r_T,r_{T-1},\dots]\\ &\quad=\frac{\Pr[S_{t+1}=l \mid r_T,r_{T-1},\dots]\cdot\Pr[S_t=j \mid r_t,r_{t-1},\dots]\cdot\Pr[S_{t+1}=l \mid S_t=j]}{\Pr[S_{t+1}=l \mid r_t,r_{t-1},\dots]}\end{aligned} \tag{A.11}$$

$$\Pr[S_{t+1}=j \mid r_T,r_{T-1},\dots]=\sum_{l=0}^{1}\Pr[S_{t+1}=l,S_{t+1}=j \mid r_T,r_{T-1},\dots]. \tag{A.12}$$

[13] Since we are not particularly interested in making inferences about the startup probabilities, we do not report their estimates. Also, for practical reasons, we treat their values as known for calculation of asymptotic standard errors based upon second derivatives. We consider this approach reasonable since inferences about the other parameters are virtually identical for other startup probabilities such as an equal probability for each initial state.

A.2. Estimation of the constant mean and time-varying coefficient specification

For specification presented in Eqs. (4) and (5), the Kalman filter is given by the following six equations (let $\beta_t \equiv \beta_t(k)$, $y_t \equiv r_t - \mu$ and $x_t \equiv \sum_{j=1}^{k} y_{t-j}$):

Prediction

$$\beta_{t\,|\,t-1} = \beta_{t-1\,|\,t-1}, \tag{A.13}$$

$$P_{t\,|\,t-1} = P_{t-1\,|\,t-1} + \sigma_v^2, \tag{A.14}$$

$$\eta_{t\,|\,t-1} \equiv y_t - y_{t\,|\,t-1} = y_t - \beta_{t\,|\,t-1} x_t, \tag{A.15}$$

$$f_{t\,|\,t-1} = x_t P_{t\,|\,t-1} x'_t + \sigma_\varepsilon^2, \tag{A.16}$$

Updating

$$\beta_{t\,|\,t} = \beta_{t\,|\,t-1} + K_t \eta_{t\,|\,t-1}, \tag{A.17}$$

$$P_{t\,|\,t} = P_{t\,|\,t-1} - K_t x_t P_{t\,|\,t-1}, \tag{A.18}$$

where $\beta_{t\,|\,t-1} \equiv E[\beta_t \,|\, r_{t-1}, r_{t-2}, \dots]$, for example, is the conditional expectation of β_t; $P_{t\,|\,t-1}$ is the variance of $\beta_{t\,|\,t-1}$; $f_{t\,|\,t-1}$ is the variance of $\eta_{t\,|\,t-1}$; and $K_t \equiv P_{t\,|\,t-1}\ x'_t\ f_{t\,|\,t-1}^{-1}$ is the Kalman gain.[14]

Given some initial values $\beta_{0\,|\,0}$ and $P_{0\,|\,0}$, we iterate through Eqs. (A.13)–(A.18) for $t = 1, \dots, T$ to obtain filtered inferences about β_t conditional on information up to time t. Also, as a by-product of this procedure, we obtain $\eta_{t\,|\,t-1}$ and $f_{t\,|\,t-1}$, which based on the prediction error decomposition (Harvey, 1990) can be used to find maximum likelihood estimates of the hyper-parameters as follows:

$$\max_{\theta} l(\theta) = -\frac{1}{2}\sum_{t=\tau+1}^{T} \ln(2\pi f_{t\,|\,t-1}) - \frac{1}{2}\sum_{t=\tau+1}^{T} \eta_{t\,|\,t-1} f_{t\,|\,t-1}^{-1} \eta_{t\,|\,t-1}, \tag{A.19}$$

where $\theta = (\mu, \sigma_\varepsilon, \sigma_v)$.

Note that we ignore the first τ observations in calculating the likelihood function. Since we do not observe β_0, and it has no unconditional expectation under the random walk specification given in Eq. (5), we must make an arbitrary guess as to its value and assign our guess an extremely large variance (e.g., $\beta_{0\,|\,0} = 0$ and $P_{0\,|\,0} \gg 0$). We then use the first τ observations to determine $\beta_{\tau\,|\,\tau}$ and $P_{\tau\,|\,\tau}$, which we treat as the initial values in the Kalman filter for the purposes

[14] For a more general discussion of the Kalman filter and time-varying parameter models, as well as details on the derivation of the Kalman gain, refer to Hamilton (1994a,b) and Kim and Nelson (1999).

of maximum likelihood estimation.[15] In practice, there is no exact rule as to what value of τ to use. Roughly speaking, we choose τ such that the effects of our arbitrary initial guess are minimized subject to including as much data in estimation as possible. The adjusted samples given in Table 3 reflect our choices for τ. The reported estimates appear to be robust to larger values of τ.

Finally, given $\beta_{T|T}$ and $P_{T|T}$ from the last iteration of the Kalman filter, we iterate backwards through the following two equations in order to obtain smoothed inferences about β_t conditional on information up to time T:

Smoothing

$$\beta_{t|T} = \beta_{t|t} + P_{t|t} P_{t+1|t}^{-1} (\beta_{t+1|T} - \beta_{t+1|t}), \tag{A.20}$$

$$P_{t|T} = P_{t|t} + P_{t|t} P_{t+1|t}^{-1} (P_{t+1|T} - P_{t+1|t}) P_{t+1|t}^{\prime -1} P_{t|t}^{\prime}. \tag{A.21}$$

A.3. Calculation of mean reversion following a price shock

We measure the economic significance of parameter estimates by calculating the implied mean reversion of a price shock over a 4-year horizon. That is, we calculate the cumulative effect of a shock on j-period-ahead return forecasts, where $j = 1, \dots, 48$ months. Construction of a given j-period-ahead forecast is somewhat complicated for the Jegadeesh regression equation. Specifically, following Doan et al. (1984), we need to employ an iterative procedure to calculate multi-period forecasts given a 1-month dependent variable. For the first specification, the law of iterated expectations implies that the resulting forecast represents $E[\sum_{i=1}^{j} r_{t+i} | r_t, r_{t-1}, \dots]$. However, it should be noted that the law of iterated expectations does not apply for the extensions since multi-period forecasts are nonlinear functions of past information.

The iterative approach to calculating economic significance works as follows. First, at time t, there is a one-unit shock. Then, for $k \geq 48$ and $j \leq 48$, the j-period-ahead expected demeaned return is calculated recursively for $j = 1, \dots, 48$ months by $r_{t+j|t} - \hat{\mu} = \hat{\beta}(k) R_{t+j-1|t}$, where $R_t = 1$ and, more generally, $R_{t+j-1} \equiv 1 + \sum_{i=1}^{j-1} (r_{t+i|t} - \hat{\mu})$ is the cumulative effect of a one-unit shock over a $j-1$ period horizon, with $\hat{\mu}$ and $\hat{\beta}(k)$ representing point estimates of the parameters. Finally, the 4-year mean reversion following the initial shock is given by $R_{t+48|t} - 1$.

[15] Alternatively, we could treat the initial value as a hyper-parameter to be estimated by maximum likelihood estimation. Inferences are very similar in both cases. However, since the hyper-parameters are treated as known, the standard error bands surrounding the inferences in this alternative case would dramatically understate the true degree of uncertainty during the early part of the sample. This is precisely when our uncertainty should be greatest.

References

Andrews, D.W.K., 1993. Tests for parameter instability and structural change with unknown change point. Econometrica 61, 821–856.

Campbell, J.Y., Hentschel, L., 1992. No news is good news: an asymmetric model of changing volatility in stock returns. Journal of Financial Economics 31, 281–318.

Campbell, J.Y., Shiller, R., 1988. The dividend–price ratio and expectations of future dividends and discount factors. Review of Financial Studies 1, 195–227.

Cecchetti, S.G., Lam, P., Mark, N.C., 1990. Mean reversion in equilibrium asset prices. American Economic Review 80, 398–418.

Doan, T., Litterman, R.B., Sims, C.A., 1984. Forecasting and conditional projection using realistic prior distributions. Econometric Reviews 3, 1–100.

Engel, C., Hamilton, J.D., 1990. Long swings in the dollar: are they in the data and do markets know it? American Economic Review 80, 689–713.

Engle, R.F., Watson, M., 1987. Kalman filter: applications to forecasting and rational-expectation models. Advances in Econometrics, Fifth World Congress, vol. 1, pp. 245–281.

Fama, E.F., French, K.R., 1988. Permanent and temporary components of stock prices. Journal of Political Economy 96, 246–273.

French, K.R., Schwert, G.W., Stambaugh, R.F., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–29.

Garbade, K., 1977. Two methods for examining the stability of regression coefficients. Journal of the American Statistical Association 72, 54–63.

Garcia, R., 1998. Asymptotic null distribution of the likelihood ratio test in Markov switching models. International Economic Review 39, 763–788.

Hamilton, J.D., 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica 57, 357–384.

Hamilton, J.D., 1994a. Time Series Analysis. Princeton Univ. Press.

Hamilton, J.D., 1994b. State-space models. In: Engle, R.F., McFadden, D.L. (Eds.), 1994b. Handbook of Econometrics, vol. 4, North-Holland, Amsterdam, pp. 3014–3077.

Hamilton, J.D., Susmel, R., 1994. Autoregressive conditional heteroskedasticity and changes in regime. Journal of Econometrics 64, 307–333.

Hansen, B.E., 1992. The likelihood ratio test under nonstandard conditions: testing the Markov switching model of GNP. Journal of Applied Econometrics 7, S61–S82.

Harvey, A.C., 1990. The Econometric Analysis of Time Series. 2nd edn. MIT Press, Cambridge.

Jarque, C., Bera, A., 1980. Efficient tests for normality, heteroskedasticity, and serial independence of regression residuals. Economics Letters 6, 255–259.

Jegadeesh, N., 1991. Seasonality in stock price mean reversion: evidence from the U.S. and the U.K. Journal of Finance 46, 1427–1444.

Kendall, M., Stuart, A., 1973. 3rd edn. The Advanced Theory of Statistics, vol. 2, Hafner Publishing, New York.

Kim, C.-J., 1994. Dynamic linear models with Markov-switching. Journal of Econometrics 60, 1–22.

Kim, C.-J., Nelson, C.R., 1998. Testing for mean reversion in heteroskedastic data II: autoregression tests based on Gibbs-sampling-augmented randomization. Journal of Empirical Finance 5, 385–396.

Kim, C.-J., Nelson, C.R., 1999. State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications. MIT Press, Cambridge.

Kim, C.-J., Nelson, C.R., Startz, R., 1998. Testing for mean reversion in heteroskedastic data based on Gibbs-sampling-augmented randomization. Journal of Empirical Finance 5, 131–154.

Kim, C.-J., Morley, J.C., Nelson, C.R., 2001, Is there a positive relationship between stock market volatility and the equity premium? Manuscript, Washington University.

Kim, M.-J., Nelson, C.R., Startz, R., 1991. Mean reversion in stock prices? A reappraisal of the empirical evidence. Review of Economic Studies 48, 515–528.

Maheu, J.M., McCurdy, T.H., 2000. Identifying bull and bear markets in stock returns. Journal of Business and Economic Statistics 18, 100–112.

Mayfield, E.S., 1999. Estimating the market risk premium, Manuscript, Harvard Business School.

McQueen, G., 1992. Long-horizon mean-reverting stock prices revisited. Journal of Financial and Quantitative Analysis 27, 1–18.

Morley, J.C., 1999. Essays in empirical finance, PhD dissertation, University of Washington.

Nyblom, J., 1989. Testing for the constancy of parameters over time. Journal of the American Statistical Association 84, 223–230.

Officer, R.R., 1976. The variability of the market factor of the New York Stock Exchange. Journal of Business 46, 434–453.

Poterba, J.M., Summers, L.H., 1988. Mean reversion in stock prices: evidence and implications. Journal of Financial Economics 22, 27–59.

Richardson, M., 1993. Temporary components of stock prices: a skeptic's view. Journal of Business and Economic Statistics 11, 199–207.

Richardson, M., Smith, T., 1994. A unified approach to testing for serial correlation in stock returns. Journal of Business 67, 371–399.

Richardson, M., Stock, J.H., 1989. Drawing inferences from statistics based on multiyear asset returns. Journal of Financial Economics 25, 323–348.

Schaller, H., van Norden, S., 1997. Regime switching in stock market returns. Applied Financial Economics 7, 177–191.

Schwert, G.W., 1989a. Business cycles, financial crises, and stock volatility. Carnegie-Rochester Conference Series on Public Policy 31, 83–126.

Schwert, G.W., 1989b. Why does stock market volatility change over time? Journal of Finance 44, 1115–1153.

Stock, J.H., Watson, M.W., 1998. Median unbiased estimation of coefficient variance in a time-varying parameter model. Journal of the American Statistical Association 93, 349–358.

Summers, L.H., 1986. Does the stock market rationally reflect fundamental values? Journal of Finance 41, 591–601.

Turner, C.M., Startz, R., Nelson, C.R., 1989. A Markov model of heteroskedasticity, risk, and learning in the stock market. Journal of Financial Economics 25, 3–22.

White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48, 817–838.

Zivot, E., Andrews, D.W.K., 1992. Further evidence on the great crash, the oil price shock, and the unit root hypothesis. Journal of Business and Economic Statistics 10, 251–270.