Journal of Monetary Economics 19 (1987) 3–24. North-Holland

THE FISHER HYPOTHESIS AND THE FORECASTABILITY AND PERSISTENCE OF INFLATION

Robert B. BARSKY *

University of Michigan, Ann Arbor, MI 48109, USA
National Bureau of Economic Research, Cambridge, MA 02138, USA

For the period 1860 to 1939, the simple correlation of the U.S. commercial paper rate with the contemporaneous inflation rate is −0.17. The corresponding correlation for the period 1950 to 1979 is 0.71. This paper attributes this apparent change in the Fisher relation to differences in the stochastic process of inflation, rather than a change in any structural relationship between interest rates and expected inflation. Contrary to recent claims in the literature, there is little evidence of inflation non-neutrality in data from the pre-World War I period.

1. Introduction

The Fisher hypothesis, which states that nominal interest rates rise point-for-point with expected inflation, leaving the real rate unaffected, is one of the cornerstones of neoclassical monetary theory. Yet prior to World War II, there is essentially no evidence of the Fisher effect in data from Britain or the United States [see, e.g., Friedman and Schwartz (1976, 1982) and Summers (1983)]. For the period 1860 to 1939, the simple correlation of the U.S. three-month commercial paper rate with the *ex post* inflation rate over the horizon of the bill is −0.17. The corresponding correlation for the period 1950 to 1979 is 0.71.

This essay has three purposes. The first is to explain why data from the post-World War II period (particularly post-1960) look more 'Fisherian' (in the sense of displaying a higher correlation between short-term nominal interest rates and measured inflation or proxies for expected inflation) than do the pre-war data. I find that one can do better than the argument that the financial markets only gradually 'learned their Fisher' [Friedman and Schwartz (1976, 1982)]. This essay emphasizes, instead, the widely divergent serial correlation (in particular, persistence) properties of inflation under different monetary regimes. Inflation evolved from essentially a white noise process in the pre-World War I years, to a highly persistent, non-stationary ARIMA process in the post-1960 period. I argue that the appearance of an *ex post*

*I am grateful to Olivier Blanchard, Brad DeLong, Stanley Fischer, Phil Howrey, Greg Mankiw, Jeff Miron, Franco Modigliani, Charles Plosser, Larry Summers, and an anonymous referee for extremely helpful discussions and comments.

Fisher effect for the first time after 1960 reflects this change in the stochastic process of inflation, rather than a change in any structural relationship between nominal rates and expected inflation.

The second focus of this essay is the possible non-neutrality of inflation vis-a-vis the real interest rate. Especially prior to World War II, there is a drastic negative relationship between the realized inflation rate and the *ex post* real commercial paper rate. Does such a negative relationship carry over to the *ex ante* real rate? Mishkin (1981) reports that, in data from 1930 to 1980, high *lagged* inflation systematically predicts a significantly lower *ex post* real rate, which under rational expectations also implies that the *ex ante* real rate is negatively correlated with past inflation. Extending Mishkin's analysis to the pre-1930 period (especially the gold standard years prior to 1913), I find no such correlation in these earlier years. Interestingly, the gold standard years, which look by far the least Fisherian in regression studies of the relation between inflation and nominal interest rates [e.g., Summers (1983)], also show the least evidence of inflation non-neutrality. The post-war period, which shows much evidence of the Fisher effect [Friedman and Schwartz (1982) and Cagan (1984)], also indicates a strong negative relationship between inflation and the expected real return on short-term financial instruments. An important methodological conclusion is that attempting to estimate the response of nominal interest rates to inflation may not be the most reliable approach to studying whether expected inflation lowers real rates or whether financial markets exhibit inflation illusion.

I find little evidence of the non-neutrality of inflation in data from the gold standard period. This contradicts the conclusion of a frequently cited study by Summers (1983), which examines the relationship betwen interest rates and inflation using band spectral regression. My third goal in this essay is to understand the discrepancy between my inference and that of Summers (1983), and to challenge the validity of Summers' conclusions. To a large extent, this argument follows easily from the previous sections of the paper. Inflation in this early period was very nearly white noise, so that the variance of anticipated inflation was very small relative to the variance of realized inflation. This would lead to a massive errors-in-variables problem in ordinary least squares regression, causing the investigator to conclude incorrectly that interest rates failed to respond to expected inflation. However, the question remains whether or not band filtering the data, leaving only low-frequency components, alleviates the problem because 'low-frequency variations in the rate of inflation are almost completely forecastable, so the assumption that expected inflation can be proxied by actual inflation is warranted' [Summers (1983, p. 216)].[1]

[1]As I make clear in section 5, Summers' principal intent in choosing the band spectral technique was not to treat the errors-in-variables problem but instead to treat the endogeneity problem that arises if inflation and real rates are simultaneously determined in the short run.

McCallum (1984) shows that low-frequency estimation is not in general robust against misspecification of the distinction between anticipated and unanticipated movements in the regressors, but gives no indication of the empirical importance of his critique. I show that data on interest rates and inflation from the pre-World War I period in the United States represent a particularly unfortunate confrontation between McCallum's general problem and the stochastic environment. Deriving and implementing a frequency domain version of the Theil misspecification theorem, I find that the covariance between anticipated and *ex post* inflation displays no tendency to increase (relative to the variance of inflation) as the frequency is lowered.

The plan of this essay is as follows. Section 2 examines the serial correlation, persistence and forecastability of inflation in the several subperiods, and their implications for the Fisher effect under a limited information version of rational expectations. Section 3 applies the Mishkin (1981) approach to the study of real interest rates and inflation in an attempt to answer the question of whether periods of generally higher inflation rates were associated with a lower *ex ante* real interest rate. Section 4 tests (and does not reject) the hypothesis that, despite the markedly different (*ex post*) decadal mean inflation rates in the gold standard period, the population mean was constant across decades. Section 5 presents a misspecification analysis of the Summers (1983) study of the long-run Fisher hypothesis in the frequency domain. Section 6 contains a brief summary and some conclusions.

2. Inflation persistence and the appearance of an *ex post* Fisher effect

A number of authors have documented that the behavior of the bivariate stochastic process of short-term interest rates and inflation shows a marked change somewhere between the end of the classical gold standard and the beginning of the 1960's [see, in particular, Sargent (1973), Klein (1975), Shiller and Siegel (1977), Friedman and Schwartz (1976, 1982), Summers (1983), Barthold and Dougan (1985)]. The essence of this change is that interest rates displayed a zero (or slightly negative) correlation with contemporaneous inflation prior to 1930, while a strongly positive correlation has been observed since about 1960. I shall refer to a strong correlation between nominal interest rates and realized inflation as an *ex post* Fisher effect (as distinguished from the *ex ante* Fisher effect, an unobservable relationship between nominal rates and underlying expected inflation).

In this section I develop relations linking the correlation between interest rates and inflation to the serial correlation, or persistence, properties of inflation, and to the percentage of forecastable variation in inflation. I go on to show, based on comparative analysis of the pre- and post-1913 inflation processes, that the simple considerations developed here explain at least the gross features of the dramatic historical change in the relationship between interest rates and inflation.

2.1. Implications of forecastability and persistence

Following McCallum (1983), suppose that the underlying model generating nominal interest rates is $i_t = \rho + E_t[\pi_{t+1}] + \varepsilon_t$. For now, let ε_t be a white noise uncorrelated with π_t, π_{t+1}, and $E_t[\pi_{t+1}]$; this will be relaxed somewhat in section 5, where we consider estimation in the frequency domain. A regression of the nominal rate i_t on π_t, the inflation rate just realized, yields an estimated coefficient with probability limit $\text{cov}(i_t, \pi_t)/\text{var}(\pi)$. This probability limit is in turn just the first-order serial correlation coefficient of inflation, regardless of the (stationary) process generating inflation, since (under our assumptions) $\text{cov}(i_t, \pi_t) = \text{cov}(\pi_{t+1}, \pi_t)$.[2] If inflation is nearly serially uncorrelated, an investigator regressing i_t on π_t would obtain a coefficient close to zero, and conclude that there is 'no Fisher effect', even though Fisher's theory is built into the underlying model. Thus, under this scenario, regression of the nominal interest rate on inflation measures the persistence of inflation rather than the response of interest rates to inflationary expectations.

Alternatively, consider the regression of the nominal interest rate on the *ex post* inflation rate realized at the maturity of the bill. Under our hypothesised underlying model, this corresponds to the case of classical errors in variables, and the probability limit of the estimated coefficient is $\text{var}(E_t[\pi_{t+1}])/\text{var}(\pi)$. This probability limit corresponds to $\text{plim}(R^2)$ from a regression of π_{t+1} on all of the variables in the agents' information sets that would have been relevant to one-step-ahead forecasting of inflation. The estimated 'Fisher coefficient' from this regression will have probability limit zero if inflation is a martingale difference with respect to all of the relevant information. A regression of the nominal rate on *ex post* inflation measures the percentage of forecastable variation in inflation, and (once again) is not a test of Fisher's hypothesis about nominal rates and expected inflation.

In summary, regression of i_t on π_t tells us something about the persistence of inflation, while regression of i_t on π_{t+1} is likely to convey information about the forecastability of inflation, as measured by the percentage of total variation in inflation that agents were able to forecast on the basis of information including, but not limited to, inflation's own past. In the highly simplified model given above, these are the *only* considerations reflected in the regressions, since the Fisher hypothesis of full adjustment for expected inflation was built into the structure. More generally, as discussed in Summers (1983) and in section 5 of this paper, the real rate will exhibit systematic (rather than merely white noise) variation, and regressions of i_t on π_t or π_{t+1} will reflect this fact as well.[3] This in no way undermines the point that we have made: a full underlying, or *ex ante*, Fisher effect is consistent with any

[2] If inflation follows a non-stationary process with a unit root, the regression of i_t on π_t would yield a coefficient of unity asymptotically, even if π_t is not a random walk.

[3] For this reason, I do not test a formal restriction that the coefficient of inflation in the interest rate equation should equal the first-order serial correlation of inflation.

observed correlation between i_t and π_t or π_{t+1}. Even when other conditions are favorable (i.e., variation in the real rate is uncorrelated with inflation), a strong *ex post* Fisher effect will appear in the data only when inflation is highly persistent and/or forecastable.

One modification of the Fisherian model with ε_t white noise will prove useful. If inflation is a martingale difference, as I am about to show was apparently the case prior to World War I, the above model predicts that nominal interest rates will also be white noise. In fact, nominal rates prior to 1913 followed highly persistent processes, although they were not a random walk [Mankiw and Miron (1986)]. One can easily account for the coexistence of interest rate persistence and serially uncorrelated inflation by letting ε_t, which represents real rate variation, follow an ARMA process. It is useful to think of a version of Fisher's underlying theory obtaining as long as ε_t is uncorrelated with current, lagged, and future π_t; i.e., real rate variation has a life of its own, independent of the inflation process. Our essential results from above continue to hold; interest rates will display no correlation with π_t if inflation is serially uncorrelated (or with π_{t+1} if inflation is unforecastable from other information as well), although the underlying model incorporates full adjustment of interest rates for expected inflation. The variation in nominal rates under this formulation is indicative entirely of real rate variation.

2.2. *Persistence and forecastability of inflation from the gold standard to the present: Empirical analysis*

2.2.1. *Data*

I examine data from both the United States and Britain. The U.S. data are quarterly. I use the Warren–Pearson all-items wholesale price index prior to 1919, and the Bureau of Labor Statistics consumer price index thereafter. Prices are taken from the third month of each quarter to limit the problem of time aggregation (which distorts the serial correlation properties of the resulting inflation rates), and to match as closely as possible (with respect to timing) the three-month commercial paper rate taken from the first month of the quarter.

The British data are annual, and were constructed by linking the Elizabeth Schumpeter consumer price series with the Gayer, Rostow, and Schwartz index and the Sauerbeck–Statist series [all found in Mitchell and Deane (1962)], as in Shiller and Siegel (1977) and Barsky and Summers (1985).

2.2.2. *Persistence*

Table 1 presents estimated autocorrelation functions for various subperiods of the quarterly U.S. and annual British data. The contrasting nature of the

Table 1
Autocorrelations of inflation.

Sample period	Price series	Standard error	Lags	Autocorrelations								
1870–1913	U.S. wholesale (quarterly)	0.08	1–8	0.02	−0.03	0.10	−0.03	−0.09	−0.08	−0.02	0.07	
			9–16	0.08	0.06	0.00	0.14	−0.06	−0.21	−0.10	0.08	
	Ljung–Box Q-test $Q(24) = 32.16$	(Signif. level) 0.11										
1729–1913	British wholesale (annual)	0.07	1–5	0.14	−0.17	−0.22	−0.16	0.10				
			6–10	0.10	0.05	0.08	0.00	−0.07				
	Ljung–Box Q-test $Q(10) = 30.37$	(Signif. level) 0.00										
1919–1938	U.S. CPI (quarterly)	0.11	1–8	0.50	0.37	0.12	0.04	−0.15	−0.18	−0.17	−0.04	
			9–16	−0.15	−0.12	−0.12	0.05	−0.02	−0.02	−0.04	0.00	
	Ljung–Box Q-test $Q(16) = 46.45$	(Signif. level) 0.00										
1947–1959	U.S. CPI (quarterly)	0.14	1–8	0.32	0.11	0.28	0.00	−0.24	−0.26	−0.24	−0.23	
			9–16	−0.17	−0.00	0.06	0.10	0.20	0.10	0.10	−0.05	0.13
	Ljung–Box Q-test $Q(16) = 33.75$	(Signif. level) 0.01										
1960–1979	U.S. CPI (quarterly)	0.11	1–8	0.81	0.77	0.72	0.70	0.60	0.48	0.45	0.36	
			9–16	0.36	0.36	0.42	0.35	0.41	0.44	0.53	0.60	
	Ljung–Box Q-test $Q(16) = 66$	(Signif. level) 0.00										

data from the different monetary regimes is striking. Pre-World War I inflation shows no sign of persistence at all, while the data from 1914 to 1959 show non-trivial, though moderate positive serial correlation. Finally, post-1960 inflation shows very great persistence.

Prior to World War I, inflation in both the United States and Britain was very nearly white noise, as indicated by the substantively small estimated autocorrelations, most of which are also (individually) statistically insignificant.[4] Note, in particular, that the first-order autocorrelations, which are the 'predicted' regression coefficients for the regression of i_t on π_t, are essentially zero. The Q-statistics, which test the joint hypothesis that the first n autocorrelations are all zero for specified n, do not reject the white noise hypothesis for the U.S. although for the British annual data that hypothesis is formally rejected at the one percent significance level. These rejections reflect the estimates for lags three through five, all of which are negative. The predominance of negative serial correlation in inflation from the gold standard period was noted earlier by Klein (1975).[5]

It is striking that the data from the pre-World War I gold standard years show *no* sign of positive serial correlation. To the extent that there is some negative serial correlation, this could be indicative of 'business cycles', as suggested by Sargent (1973), or adjustments peculiar to a gold standard [see, e.g., Rockoff (1984) and Barsky and Summers (1985)]. One might note, however, that plausible moving average measurement error in the underlying price data would also lead to a predominance of negative values in the correlogram of the inflation series. Thus the true inflation data may have been even closer to white noise than the measured series.

An alternative (though isomorphic) characterization of the behavior of inflation over the past two and a half centuries makes use of frequency domain techniques. Instead of estimating spectral densities – which involves an important judgemental aspect, since smoothing is required to obtain consistent estimates – I present the results of the Durbin cumulative periodogram test [see, e.g., Malinvaud (1980)]. This compares the sum of the periodogram values (beginning at the low-frequency end) with a 45 degree line through the origin, the integrated spectral density of theoretical white noise. The 45 degree line (not shown explicitly) is flanked by confidence bounds derived under the null hypothesis of white noise. Note that no ad hoc decisions need to be made in order to implement this procedure.

[4] Very similar results obtain beginning in 1880, or splitting the sample at 1896 and examining the two halves separately.

[5] Although his emphasis is somewhat different, Klein (1975), in noting the implications of the serial correlation of inflation for interest rate determination, anticipated some of the present analysis. Ibrahim and Williams (1978) suggested that the Fisher hypothesis may not be testable using older data because of the absence of positive serial correlation in inflation rates, a point which also anticipates the argument of this paper.

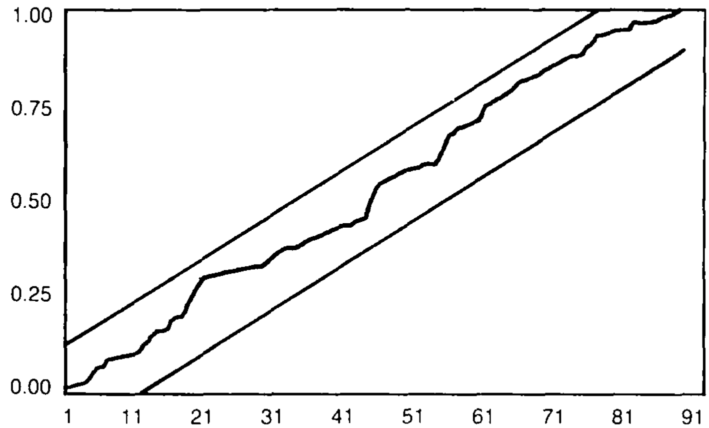

Fig. 1a. Cumulative periodogram of U.S. inflation, 1870–1913 (quarterly).

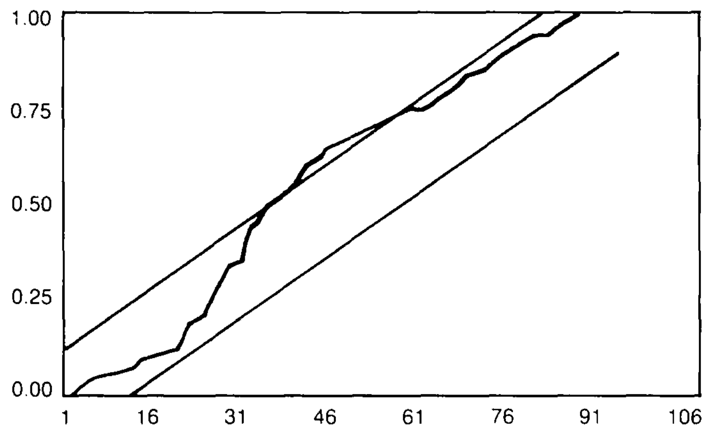

Fig. 1b. Cumulative periodogram of British inflation, 1729–1913 (annual).

Figs. 1a, 1b, and 1c show the cumulative periodogram with confidence bounds for the U.S. from 1870 to 1913, Britain from 1729 to 1913, and the U.S. from 1940 to 1979, respectively. The results here simply restate in terms of frequencies our previous conclusions. Since the cumulative periodogram in fig. 1a stays well within the region bounded by the two parallel diagonals, the early U.S. data show no departure at all from random behavior. The long British series subtly rejects pure randomness in light of the accentuated middle

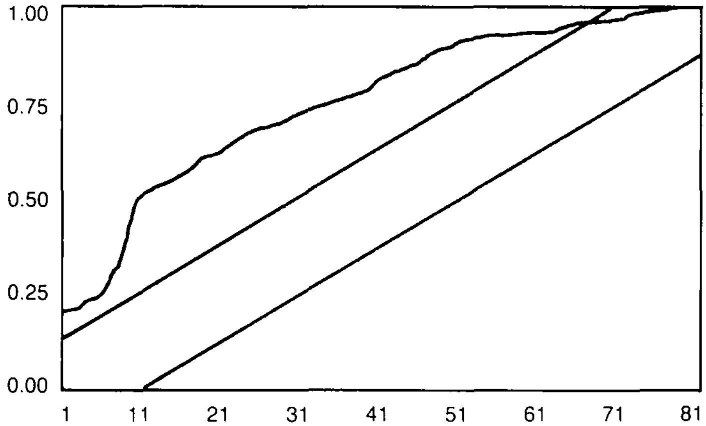

Fig. 1c. Cumulative periodogram of U.S. inflation, 1940–1979 (quarterly).

(i.e., 'business cycle') frequencies. The 1940 to 1979 U.S. sample is sufficiently concentrated in the low frequencies (i.e., persistent) that the cumulative periodogram exceeds the upper bound from the beginning. Note that the slope is steeper than that of the 45 degree line until the upper third or so of the frequency band, where it becomes flatter, confirming that the higher frequencies contribute proportionately less to the variance of this series than to a white noise series.

Table 2 summarizes my Box–Jenkins identifications for United States inflation over the past 140 years. I experimented with ARMA models even for the 1870 to 1913 sample, but asymptotic F-tests of the hypothesis that all of

Table 2

ARIMA models for inflation, 1870–1979, quarterly U.S. data.

Sample period	Identification	Fitted model
1870–1913	Essentially white noise	$x_t = \varepsilon_t$
1919–1938	AR(2)	$x_t = -\underset{(0.84)}{0.30} + \underset{(0.11)}{0.42}\, x_{t-1} + \underset{(0.12)}{0.17}\, x_{t-2} + \varepsilon_t$
1947–1959	AR(1)	$x_t = \underset{(0.65)}{1.63} + \underset{(0.14)}{0.33}\, x_{t-1} + \varepsilon_t$
1960–1979	IMA(1,1)	$x_t = x_{t-1} - \underset{(0.09)}{0.46}\, \varepsilon_{t-1} + \varepsilon_t$

the coefficients are zero did not reject.[6] For subperiods of 1913 to 1979, I obtained parameter estimates which yielded approximately white noise residuals. These are shown in the column at right. Note that I identify the 1960 to 1979 inflation process as a non-stationary IMA (1, 1). The estimated fraction of each period's innovation that is permanent is about 0.5.

All of the results presented here accentuate the marked change that inflation underwent over the last seventy years. Inflations prior to 1913 displayed no tendency to persist. If anything, there was a slight tendency for inflations to be followed by deflation two to four years later. From 1913 to 1979 the persistence of inflation became progressively greater. The great momentum of inflation in the last twenty years or so is reflected in an ARIMA representation with a unit root. As noted by Granger and Newbold (1977), forecasting the *level* of an integrated series is a relatively easy task. This is reflected in the strong correlation between interest rates and inflation seen after 1960.

The greater the persistence of inflation, the more expected inflation will resemble current inflation, and hence the stronger will be the appearance of an *ex post* Fisher effect. In particular, this regularity will hold over a class of models in which $\mathrm{d}i/\mathrm{d}\pi^e$, the response of nominal interest rates to true expected inflation, is identically unity, as required for superneutrality. Thus the correlation between nominal interest rates and *realized* inflation may say more about the stochastic process followed by inflation than about the truth of Fisher's theory that expected real interest rates are orthogonal to inflation. In section 5, I show that this continues to be true even at the low frequencies.

2.2.3. Forecastability

As noted above, it is impossible to say with complete confidence how forecastable inflation was in a particular historical setting, since we have access to at most a subset of the agents' information sets. In this section, I focus on the percentage of total variation explained by inflation's own past and by lagged gold production to obtain a 'limited information' metric of forecastability. I use the 'corrected' squared correlation coefficient

$$\bar{R}^2 = 1 - \frac{\dfrac{1}{T-K}\sum_{i=1}^{T}\varepsilon_i^2}{\dfrac{1}{T-1}\sum_{i=1}^{T}(\pi_i - \bar{\pi})^2},$$

[6]Summers (1984) reports that regressions of interest rates on ARIMA forecasts of inflation yielded results virtually identical to those from regressions of interest on *ex post* inflation for the pre-1930 period. Summers argues that the regressions using ARIMA forecasts are robust against the charge of specification error made by McCallum (1984). Our asymptotic F-tests do not support Summers' claim. The ARIMA models have no explanatory power. However, the coefficients estimated from finite samples will not be identically zero, and therefore the forecasts display considerable spurious variation.

Table 3

$\bar{R}^2$'s from regressions of U.S. inflation on lagged information.

Sample period	Information set	$\bar{R}^2$
1870–1913 (quarterly)	4 lags of inflation	−0.00
	12 lags of inflation	0.01
1870–1913 (annual)	1 lag of percentage change in monetary gold	0.11
	5 lags of percentage change in monetary gold	0.16
1919–1938 (quarterly)	4 lags of inflation	0.23
	12 lags of inflation	0.27
1947–1959 (quarterly)	4 lags of inflation	0.21
	12 lags of inflation	0.29
1960–1979 (quarterly)	4 lags of inflation	0.70
	12 lags of inflation	0.75

so that results involving different numbers of lagged regressors will be comparable. Note that as the sample size increases, $\bar{R}^2$ converges in probability to the ratio of the variance of the 'forecastable component' to the total variance.

Table 3 presents $\bar{R}^2$'s from regressions of U.S. inflation on lagged information from the various subperiods. In addition to univariate autoregressions, regressions using the growth of the world monetary gold stock from Kitchin (1930) are shown. The gold stock regressions necessitate the use of annual data. The choice of the growth rate of world gold stocks as a forecasting variable was suggested by what I call the 'traditional view' of price movements during the gold standard period, to be discussed further in section 3. Briefly, this view attributes the major swings in prices prior to 1913 to changes in the rate of gold production, particularly as a result of discoveries of new gold and improved methods of extraction.

The results in table 3 suggest that the forecastability of inflation underwent the same evolution that characterized inflation persistence. Past inflation explained none of the variance of current inflation prior to 1913. Lagged gold production was somewhat more useful. However, it is somewhat doubtful whether agents had access to all of the information about gold production reflected in the data used here, which were not assembled until 1930. Inflation shows a moderate degree of forecastability between 1919 and 1959, with $\bar{R}^2$'s based on past inflation alone on the order of 20 percent. Finally, inflation from 1960 to 1979 was highly forecastable. Three quarters of the variation in inflation from this period could have been foreseen from the past behavior of inflation alone. The increased forecastability and persistence of inflation since 1913 (and especially since 1960) account impressively for the gross features of the historical evolution of the relationship between interest rates and inflation. In particular, the complete absence of an *ex post* Fisher effect prior to 1913 and the strong emergence of such an effect after 1960 are fully rationalized.

Some puzzles remain. The rather abrupt appearance of mild inflation persistence afer 1914 was not met even by a small increase in the correlation between interest rates and inflation, in contradiction to the implications of the illustrative model presented in this section. The period between 1914 and 1953 was characterized by so many special circumstances – two world wars, the Great Depression, price controls, and a massive interest rate pegging program – however, that it is not clear what to make of this failure.

3. Real interest rates and inflation non-neutrality: Evidence from the pre-1914 period

The essence of Fisher's hypothesis is that nominal interest rates are set so that *forecastable* inflations do not systematically lower real interest rates. Thus an alternative approach to testing the Fisher proposition focuses on its implication that the *ex post* real rate should bear no systematic relationship to lagged inflation. Mishkin (1981) notes that if $X_t\beta$ is the projection of the *ex ante* real rate on any information set known at time t (represented by the vector X_t), the projection of the *ex post* real rate on that same information set is also $X_t\beta$. This follows from the fact that, under rational expectations, the projection of the expectational error on X_t must equal zero. As noted by Mishkin (1981), these regressions have no structural interpretation. However, if the paradigm of Summers (1983), in which 'high-inflation' decades alternated with periods of deflation or low inflation, and in which nominal rates failed to adjust for these regime changes, really characterized the pre-1940 period, regression of *ex post* real commercial paper rates on a number of lags of inflation should yield a significantly negative estimate of the sum of the coefficients.

Table 4 presents regression results involving three-month U.S. commercial paper rates and four, eight, and twenty quarters of lagged inflation. F-tests of the hypothesis that the lagged inflation rates do not help predict the *ex post* real rate are shown in the final column. For 1870 to 1913, the hypothesis that lagged inflation rates are irrelevant for predicting the *ex post* real rate is not rejected for any choice of the number of lags. For the post-1930 period, that hypothesis is soundly rejected. After 1930, real rates were on average lower

Table 4
Regression of U.S. real commercial paper rates on lagged information.

Sample period	Number of lags	Sum of lag coefficients (standard errors in parentheses)		$\bar{R}^2$	F-test of significance of regression
		(A) Lagged inflation			
1870–1913	4	−0.07 (0.15)		−0.01	$F(4,171) = 0.69$
	8	−0.02 (0.22)		−0.01	$F(8,167) = 0.84$
	20	−0.32 (0.31)		0.06	$F(20,155) = 1.55$
1930–1979	4	−0.61 (0.08)		0.27	$F(4,195) = 19.45$
	8	−0.55 (0.09)		0.29	$F(8,191) = 11.27$
	20	−0.57 (0.11)		0.30	$F(20,179) = 5.17$
		(B) Lagged inflation, Nominal rate			
1870–1913	4	−0.11 (0.10)	2.3 (0.6)	0.06	$F(5,170) = 3.4$
		(C) Lagged growth of monetary gold			
1870–1913 (annual data)	1	−2.4 (0.80)		0.16	$F(1,41) = 9.0$
	5	−2.3 (0.86)		0.22	$F(5,37) = 3.35$

during inflationary periods. This conclusion continues to hold when the 1930's and 1940's are omitted (not shown here).

These results distinctly fail to lend support to the view that non-neutrality of inflation characterized the period prior to 1913. They do suggest that inflation non-neutrality may have characterized the post-1930 and post-1950 periods. The earlier years are usually thought of as the least 'Fisherian' because of the absence of close co-movement of nominal rates and *ex post* inflation, while the post-1950 period is characterized by a high correlation of realized inflation and nominal rates. Yet, when the Fisher theory is restated in such a way that the expectational error is in the left-hand-side variable, it is the pre-1913 period which provides the least evidence against Fisher's neutrality theory.

Since the nominal rate is known at time t, it is appropriate to add it to the information set, alongside lagged inflation. This is done in the middle set of entries in table 4. The coefficient of the nominal rate and its standard error are shown, in addition to the sum of the coefficients on lagged inflation and its standard error. The results are quite striking. Lagged inflation continues to be irrelevant for predicting the *ex post* real rate. The nominal rate is a statistically significant predictor of the real rate, although the R^2 is quite low. The estimates are consistent with the view that changes in nominal rates indicated changes in real rates on a one-for-one basis prior to 1913, a hypothesis put forth by Shiller and Siegel (1977).

In the previous section of this paper, we found that inflation during the gold standard years was nearly white noise, implying that the price level was essentially a random walk. A more traditional view [see Delong (1985) and Cagan (1984) for recent restatements] divides the 1870 to 1913 period into two segments: an era of generally declining prices from 1870 until 1896 and an inflationary period after 1896. To the extent that this 'traditional view' is based on an examination of plots of the price series alone, it requires little discussion; the appearance *ex post* of spurious trends in random walks is a well known phenomenon. A more sophisticated version of the traditional view emphasizes changes in the trend growth of world gold stocks. Prior to 1896 the argument goes, a dearth of gold discoveries combined with a high growth rate of potential GNP lead to a deflationary regime; after 1896, new sources in South Africa and Australia combined with the discovery of the cyanide process lead to more rapid increases in the world stock of gold, and hence inflation. Furthermore, it is claimed [Delong (1985)], agents living at the time were aware of a switch from a deflationary to an inflationary regime.

Underlying the Summers (1983, 1984) view that the gold standard provides substantial evidence against the Fisher hypothesis is the notion that the period included important swings in the trend growth of prices. Since the univariate inflation process is white noise, it shows no evidence of such variation. The possibility remains that the alleged 'regime change', along with a failure of nominal interest rates to adjust, would be reflected in a negative relationship between real rates and recent rates of gold production. Thus, the final entries in table 4 show the results of regressions of real interest rates on lagged growth rates of the world gold stock from Kitchin (1930). This necessitated a move to annual data. Real interest rates do show a stronger negative relationship to lagged gold production than to lagged inflation, which corresponds to our finding in section 2 that inflation was somewhat forecastable from lagged gold production numbers. If there is an argument in defense of Summers (1983), it is likely to be based on the forecastability of inflation from lagged gold growth, rather than any features of the univariate inflation process. However, there are alternative scenarios other than inflation non-neutrality which could also have lead to a correlation between real rates and the growth of the gold stock.[7]

[7] For example, consider the model used by Barsky and Summers (1985) to study Gibson's Paradox. An announcement that real rates would be lower in the future would cause an immediate rise in the relative price of gold. This would lead to increased gold production now, leading to a negative correlation between current gold production and future real rates.

4. Variation in inflation within and between decades

Summers (1983) prints decadal mean inflation rates for 1860 to 1940, and argues that their substantially varying values indicate important changes in trend inflation that should not have been ignored by bondholders. To test the validity of this claim, we might ask whether the data provide evidence that the *population* means in fact varied across decades. This is the problem of 'one way analysis of variance'. The null hypothesis that the decadal means were all equal can be tested with the statistic

$$\frac{\sum_{j=1}^{k}\sum_{i=1}^{n}\left[(x_{ij}-\bar{\bar{x}})^2-(x_{ij}-\bar{x}_j)^2\right]\Big/k-1}{\sum_{j=1}^{k}\sum_{i=1}^{n}(x_{ij}-\bar{x}_j)^2\Big/T-k},$$

where $\bar{\bar{x}}$ is the 'pooled' sample mean and $\bar{x}_j$ is the mean for decade j. The statistic is distributed $F_{T-k,k-1}$ [Rao (1973)]. The numerator represents the variation *between* decades, while the denominator measures the *within* variation. Only if the 'between' variation is of sufficient relative magnitude is the hypothesis of equality rejected.

Table 5 presents the analysis of variance for the subperiods 1870 to 1909, 1860 to 1939, and 1940 to 1979. Only for the post-1940 decades is there strong ground for rejection. For the earlier periods, the F-statistics are surprisingly small, reflecting the large amount of within-decade variation in the inflation rates. The observed variation in sample means is consistent with a constant population mean up to 1940.

It is true that the foregoing tests assume constancy of the population variance, and that this can sometimes be rejected by the data. For instance, the

Table 5
Test of equality of decadal mean inflation rates.

Sample period	Source of variation	Degrees of freedom	Sum of squares	Mean square	F	Marginal significance level
1870–1909	between	3	0.0205	0.0068	1.41	0.26
	within	36	0.1745	0.0049		
	total	39	0.195			
1860–1939	between	7	0.105	0.015	1.69	0.13
	within	72	0.64198	0.0089		
	total	79	0.747			
1940–1979	between	3	0.0166	0.00553	4.65	0.01
	within	36	0.04278	0.00119		
	total	39	0.05938			

variability of inflation was especially great during the Civil War decade of the 1860's and the World War I period 1910 to 1919. However, Box (1954) shows that, as long as the number of data points in each cell is the same, the bias in the test from falsely maintaining a constant population variance is always in favor of rejection. Thus failure to reject in the early decades cannot be attributed to bias from the maintained hypothesis of a constant population variance.

The results in this section cast serious doubt on the view that there were 'high inflation' and 'low inflation' (or deflationary) decades during the gold standard period. When one allows for the possibility of sampling variation, the likelihood emerges that the appearance of changes in the trend growth of prices so central to the arguments of Summers (1983), Cagan (1984), and Delong (1985) is spurious.[8] Our findings thus further undermine the view that data from the gold standard years can provide substantive evidence against inflation neutrality.

5. Specification analysis of Summers' spectral estimates of the Fisher equation

Summers (1983) uses band spectral regression to filter out high-frequency components of the inflation and interest rate data, and estimate a 'long-run' Fisher equation. Summers argues for this procedure on the grounds that: (1) economic theory suggests approximate super-neutrality of money only as a steady state proposition; and (2) low-frequency estimation is robust against the errors-in-variables problem that arises from the use of actual inflation as a proxy for anticipated inflation.

Summers (1983) did not emphasize the errors-in-variables issue, preferring to focus on point (1). The choice of the band spectral technique to deal with the short-run joint endogeneity of inflation and real rates was an innovative one. Appropriate inference, however, depends also on the truth of the second claim, on which this section focuses. I first show that it is not correct in general, a point previously made in a slightly different way by McCallum (1984). I then go on to show that for the 1870 to 1913 period in particular it leads to extremely misleading conclusions.

Suppose that the 'true' population model is $i_t = \rho + \mathrm{E}_t[\pi_{t+1}] + \varepsilon_t$, so that rationally expected inflation appears with a coefficient of unity, ρ is the unconditional expectation of the real rate, and ε reflects variation in the *ex ante* real rate. In general, ε_t will be correlated with $\mathrm{E}_t[\pi_{t+1}]$, as well as serially correlated. Thus, even if it were possible to observe inflationary expectations exactly, OLS estimation of the above relation would not be meaningful.

[8] My calculations agree with DeLong (1985) that if one splits the sample around the 1896 turning point, a 't-test' rejects equality of the two means. But such an *ex post* choice for the breaking point hardly seems legitimate.

Let us assume, however [with Summers (1983)], that the steady state superneutrality result properly implies the absence of *low-frequency* correlation between ε_t and $E_t[\pi_{t+1}]$. In other words, we assume that if we could observe true expected inflation, the low-frequency band spectral estimate of its coefficient would have probability limit equal to unity. In fact, in order to isolate the effect of proxying expected inflation with π_t, let us go a step further and proceed as if ε_t is uncorrelated with both $E_t[\pi_{t+1}]$ and π_t at all frequencies. If regression of i_t on π_t does not make sense in this most favorable of circumstances, it will be no better when endogeneity (of π_t or $E_t[\pi_{t+1}]$) is readmitted.

We can now ask what the probability limit of the band spectral estimator B over specified frequencies will be. This is given by $\operatorname{plim} B = \int S_{i,\pi}(\omega) / \int S_{\pi,\pi}(\omega)$, where $S_{i,\pi}(\omega)$ is the (population) cross-spectrum of interest and inflation, $S_{\pi,\pi}(\omega)$ is the spectrum of inflation, and the integral is taken over the specified frequency band $(-\underline{\omega}, \underline{\omega})$. Under our assumptions, $\operatorname{plim} B = \int S_{E_t[\pi_{t+1}],\pi_t} / \int S_{\pi,\pi}$ and is thus equal to the band spectral regression (over the same frequency band) of expected inflation at $t+1$ on actual inflation at t.

If the integrals are taken over the entire interval $(-\pi, \pi)$, the probability limit of the OLS estimator is obtained, and this is seen to be $\operatorname{cov}(\pi_t, E_t[\pi_{t+1}])/\operatorname{var}(\pi_t)$. This coincides with the standard Theil (1957) specification error result: the estimated coefficient in a regression with an erroneous explanatory variable differs (for large samples) multiplicatively from the true coefficient by the regression of the 'correct' explanatory variable on the erroneous, included one. Letting ρ be the correlation between π_t and $E_t[\pi_{t+1}]$, note that $\operatorname{cov}(\pi_t, E_t[\pi_{t+1}])/\operatorname{var}(\pi_t) = \rho\sigma(E_t[\pi_{t+1}])/\sigma(\pi_t) < 1$, since $\rho < 1$ and a rational forecast varies less than the series being forecast. Thus the OLS estimate of the response of i to expected inflation must be biased downward as long as inflation is a stationary series. The extent of this bias depends on the stochastic process followed by inflation. In the limit as π approaches white noise behavior, $\operatorname{cov}(\pi_t, E_t[\pi_t + {}_1])/\operatorname{var}(\pi_t)$ approaches zero, and regression of i on π will yield a zero coefficient even though a full response of i to *expected* inflation obtains by hypothesis.

Does the situation improve as we focus on the relation at lower frequencies? Suppose that π has an autoregressive representation $\pi_{t+1} = A(L)\pi_t + \varepsilon_t$, and that expectations are based on the univariate process followed by inflation. Then $\operatorname{plim} B = \int A(e^{-i\omega})$, the integration once again taking place over $(-\underline{\omega}, \underline{\omega})$. The probability limit of the limiting zero frequency estimator (the limit of the integral as $\underline{\omega} \to 0$) is just the sum of the coefficients in the autoregressive representation of π. Only if these coefficients sum to unity does low-frequency estimation circumvent (in the limit) the specification bias from using π_t in place of $E_t[\pi_{t+1}]$. Alternatively, the required condition is that the process generating inflation have a unit root [Box and Jenkins (1976, p. 102)]. The IMA (1,1) process is a noteworthy special case. It should not be surprising if low-frequency regression results from a period in which inflation followed an

Table 6
'Auxilliary regressions' of $E_t(\pi_{t+1})$ on π_t.

Sample period	Coefficient of π_t	$\bar{R}^2$
1870–1913	0.02 (0.02)	−0.00
1930–1979	0.34 (0.03)	0.32
1960–1979	0.58 (0.05)	0.60

IMA (1,1) process appear more favorable to the Fisher effect than results from a period in which inflation was nearly white noise.

It is not hard to implement the above formulae under the assumption that inflationary expectations do not differ too much from the predictions of a univariate ARIMA model. The procedure is to generate ARIMA forecasts and then to compute the 'auxilliary regression' [Theil (1957)] of the one-step-ahead inflation forecast on inflation at t. Corresponding to each quarter from 1860 to 1980, I estimated an ARMA (1,1) model using the previous eighty observations (and hence only information that would have been available to agents at the time of the forecast), and computed the one-step-ahead predictions.

Table 6 presents OLS regressions of $E_t[\pi_{t+1}]$ on π_t for the subperiods 1870 to 1913, 1930 to 1979, and 1960 to 1979. Figs. 2a through 2c show the gain $|S_{E_t[\pi_{t+1}],\pi_t}(\omega)|/S_{\pi,\pi}(\omega)$ as a function of frequency. The gain (or transfer function) indicates the magnitude of the relationship between the two series at each frequency. For the 1870 to 1913 period, the OLS relation between inflation at t and the one-step-ahead inflation forecast is negligible, and there is no tendency for this relation to strengthen at lower frequencies. In the later periods, on the other hand, a substantial OLS relationship appears, and the downward-sloping gain function shows that the coefficient at lower frequencies is larger than that at high frequencies, and larger than the OLS regression coefficient. The 1960 to 1979 results are particularly striking in this regard. The coefficient approaches one closely for long cycles. This is precisely the period in which we identified inflation as an IMA (1,1) process.

We thus can account for the coefficient estimates obtained by Summers solely in terms of the stochastic properties of inflation, i.e., in a model in which the adjustment of nominal interest rates to rationally expected inflation is always one-for-one. The more 'Fisherian' results for the post-1930 and particularly the post-1960 period appear to reflect an increase in the extent to which (a smoothed version of) actual inflation proxies for expected inflation rather than a change in the structural relationship between interest rates and *expected* inflation. Even more surprisingly, the estimated gain functions sug-

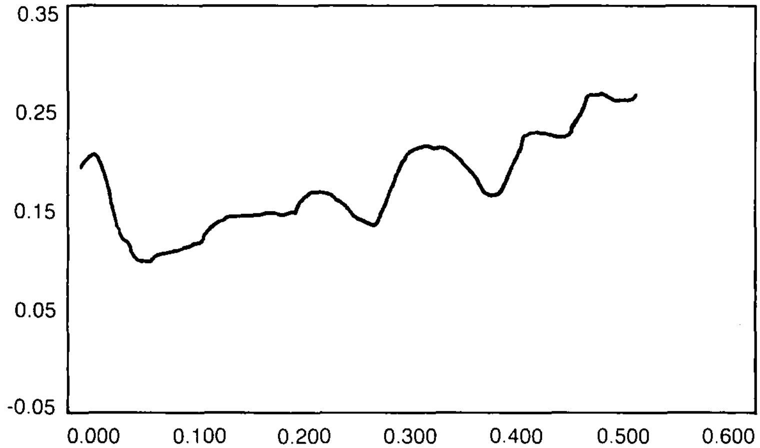

Fig. 2a. Gain function, U.S. inflation, 1870–1913.

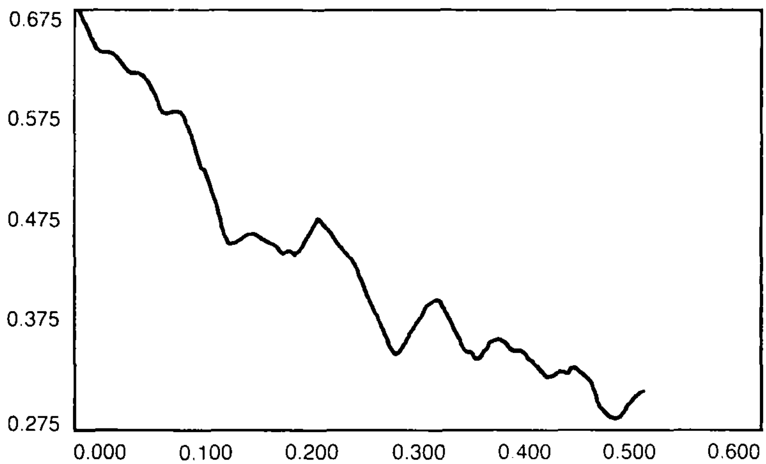

Fig. 2b. Gain function, U.S. inflation, 1930–1979.

gest that the apparent strengthening of the Fisher relation at lower frequencies in the post-1930 data may also reflect properties of the inflation process rather than short-run vs. long-run adjustment to expected inflation. Recall that the gain functions shown above are estimates of what the band spectral regression coefficient of i on *actual* inflation 'should' be under the maintained hypotheses of full adjustment to expected inflation and (limited information) rational expectations. The closeness of these hypothetical coefficients to those actually obtained in Summers (1983) is striking.

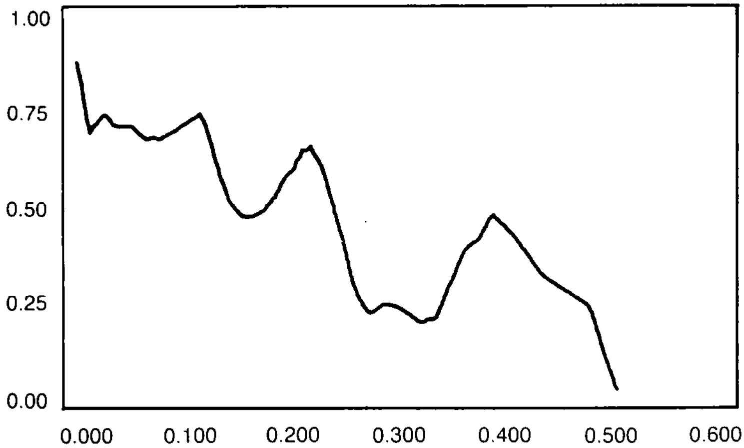

Fig. 2c. Gain function, U.S. inflation, 1960–1979.

6. Summary and conclusions

In assessing empirical evidence on the Fisher hypothesis, it is important to distinguish between two quite separate issues. The first is the degree of correlation between nominal interest rates and realized inflation. The second is the relationship between expected inflation and expected real rates.

This paper concurs fully with Summers (1983) in the conclusion that essentially none of the variance of nominal interest rates prior to 1930 is accounted for by inflation, actual or expected. However, I do not conclude that the data, at least those prior to 1913, show much evidence of inflation illusion or non-neutrality of expected inflation vis-a-vis real rates. The time series properties of inflation prior to 1913 suggest that the relationship between actual inflation and expected inflation in this period was negligible (although there are some caveats regarding lagged gold production), and was no stronger at low frequencies than at high ones. The estimated non-response of nominal rates to inflation in these data is as likely a reflection of this phenomenon as it is evidence of a non-adjustment of nominal rates to *expected* inflation. Indeed, when the Fisher equation is 'turned around' so that the expectational error is associated with the left-hand-side variable rather than the regressor, no significantly negative relationship between real rates and inflation appears.

I conclude with an important caveat. One way of restating the main empirical result of this paper is that there was probably little variation in expected inflation prior to 1913. Thus, although this period cannot provide significant evidence against Fisher's hypothesis, neither can it tell us what would have happened if expected inflation had varied widely. It would be wrong to conclude that the early period provides positive evidence in favor of Fisher. However, when the data from the gold standard years are understood within the framework of this paper, much of the apparently overwhelming case against the Fisher hypothesis disappears.

References

Barsky, R. and L. Summers, 1985, Gibson's paradox and the gold standard, National Bureau of Economic Research working paper no. 1680.

Barthold, T. and W. Dougan, 1985, The Fisher hypothesis under different monetary regimes, Dartmouth College working paper.

Box, G.E.P., 1954, Some theorems on quadratic forms applied in the study of analysis of variance problems, Annals of Mathematical Statistics 25, 290–302.

Box, G.E.P. and G.M. Jenkins, 1976, Time series analysis: Forecasting and control (Holden-Day, San Francisco, CA).

Cagan, P., 1984, Comment on Benjamin and Kochin, in: M.D. Bordo and A.J. Schwartz, eds., A retrospective on the classical gold standard (University of Chicago Press for N.B.E.R., Chicago, IL).

DeLong, B., 1985, Observations on the nonadjustment of nominal interest rates (Center for Research in Economic History, Harvard University, Cambridge, MA).

Engle, R., 1974, Band spectrum regression, International Economic Review 15, 1–11.

Fisher, I., 1930, The theory of interest (Macmillan, New York).

Friedman, M. and A. Schwartz, 1976, From Gibson to Fisher, Explorations in Economic Research 3, 288–289.

Friedman, M. and A. Schwartz, 1982, Monetary trends in the United States and the United Kingdom (University of Chicago Press for N.B.E.R., Chicago, IL).

Ibrahim, I.B. and R.M. Williams, 1978, The Fisher relationship under different monetary standards, Journal of Money, Credit, and Banking 10, 363–370.

Granger, C. and P. Newbold, 1977, Forecasting economic time series (Academic Press, New York).

Kitchin, J., 1930, Gold production, in: The international gold problem (Royal Institute of Economic Affairs, Oxford University Press, London).

Klein, B., 1975, Our new monetary standard: Measurement and effects of price uncertainty, 1880 to 1973, Economic Inquiry 13, 461–484.

Malinvaud, E., 1980, Statistical methods of econometrics (North-Holland, Amsterdam).

Mankiw, N.G. and J. Miron, 1986, The changing behavior of the term structure of interest rates, Quarterly Journal of Economics 101, 211–228.

McCallum, B., 1984, On low-frequency estimates of long-run relationships in macroeconomics, Journal of Monetary Economics 14, 3–14.

Mishkin, F., 1981, The real rate: An empirical investigation, Carnegie–Rochester Conference Series on Public Policy 12, 519–541.

Muth, J., 1960, Optimal properties of exponentially weighted forecasts, Journal of the American Statistical Association 55, no. 290, 299–306.

Mitchell, B.R. and P. Deane, 1964, Abstract of British historical statistics (Cambridge University Press, Cambridge).

Rao, C.R., 1974, Linear statistical inference and applications (Wiley, New York).

Rockoff, H., 1984, Some evidence on the real price of gold, its cost of production, and commodity prices in: M.D. Bordo and A.J. Schwartz, eds., A retrospective on the classical gold standard (University of Chicago Press, Chicago, IL).

Sargent, T., 1973, Interest and prices in the long run: A study of the Gibson paradox, Journal of Money, Credit, and Banking 5, 385–449.

Shiller, P. and J. Siegel, 1977, The Gibson paradox and historical movements in real interest rates, Journal of Political Economy 85, 891–907.

Summers, L., 1983, The nonadjustment of nominal interest rates: A study of the Fisher effect, in: J. Tobin, ed., Macroeconomic prices and quantities: Essays in memory of Arthus Okun (Brookings Institution, Washington, DC).
Summers, L., 1984, Estimating the long-run relationship between interest rates and inflation: A response to McCallum, National Bureau of Economic Research working paper no. 1448.
Theil, H., 1957, Specification errors and the estimation of economic relationships, Review of the International Statistical Institute 25, 41–51.
Warren and Pearson, 1933, Prices.