The MIT Press

REFERENCES

Linked references are available on JSTOR for this article:
https://www.jstor.org/stable/27867542?seq=1&cid=pdf-reference#references_tab_contents
You may need to log in to JSTOR to access the linked references.

THE PREBISCH-SINGER HYPOTHESIS: FOUR CENTURIES OF EVIDENCE

David I. Harvey, Neil M. Kellard, Jakob B. Madsen, and Mark E. Wohar*

Abstract—We employ a unique data set and new time-series techniques to reexamine the existence of trends in relative primary commodity prices. The data set comprises 25 commodities and provides a new historical perspective, spanning the seventeenth to the twenty-first centuries. New tests for the trend function, robust to the order of integration of the series, are applied to the data. Results show that eleven price series present a significant and downward trend over all or some fraction of the sample period. In the very long run, a secular, deteriorating trend is a relevant phenomenon for a significant proportion of primary commodities.

I. Introduction

THIS paper reexamines the time-series properties of primary commodity prices relative to manufactures and, in particular, the Prebisch-Singer (PS) hypothesis that such prices present a downward secular trend. This is important because many developing countries rely on a small number of primary commodities to generate the majority of their export earnings. Overall, for the least developed countries, approximately 60% of export earnings are derived from primary commodities. However, for 40 countries, the production of three or fewer commodities explains all export earnings. This level of commodity dependency has profound policy implications conditional on the behavior of prices. Clearly, given strong evidence of a long-run downward trend in its relevant export commodities, a country might explore diversification of its export portfolio to include manufactures or services.[1]

Theoretical explanations that present, as a corollary, declining relative commodity prices include a low income elasticity of demand for primary commodities, lack of differentiation among commodity producers leading to highly competitive markets, productivity differentials between North (industrial) and South (commodity-producing) countries, and asymmetric market structure (where manufacturing industries capture oligopolistic rents relative to competitive firms earning zero economic profits and producing primary commodities). On the other hand, Lewis (1954) suggests a theoretical account of commodity price determination, which would imply a zero trend in relative prices of some primary commodities (see Deaton, 1999). Briefly, Lewis proposes that real wages will not grow in very poor countries because of unlimited supplies of labor at the subsistence wage. Therefore, the prices of tropical commodities like cocoa cannot, in the long run at least, exceed the costs of production in the lowest real wage region where the crop can be planted. Deaton (1999) subsequently comments, "There is no trend, because the poorest workers in the tropics remain as poor as ever. Prices will always eventually revert to base because, while short-run events can increase prices, sometimes for many years, long-run marginal cost is set by the poverty of the tropics and supply will eventually be forthcoming" (p. 30).

Early empirical evidence on the existence of a downward trend assumed that y_t (the logarithm of the relative commodity price) is generated by a trend-stationary (TS) process,

$$y_t = \alpha + \beta t + u_t, \ t = 1, \ldots, T, \tag{1}$$

where t is a linear trend and the random variable u_t is stationary with mean 0. The focal point of interest is the slope parameter β, which the PS hypothesis predicts will be less than 0. Estimations of equation (1) have typically found strong support for the PS hypothesis.[2] For example, Grilli and Yang (1988), using a data set of 24 annual commodity prices from 1900 to 1986, found a weighted aggregate index declined by 0.6% per annum, and most of the subsequent literature uses extended versions of the Grilli-Yang data set.

An alternative and commonly assumed generating process is represented by the difference-stationary (DS) model:

$$\Delta y_t = \beta + v_t, \ t = 2, \ldots, T, \tag{2}$$

where the generating process for v_t is stationary and invertible. Recent empirical studies estimating equation (2) have found evidence against the PS hypothesis. Notably, Kim et al. (2003) suggest that commodity prices exhibit unit root behavior, and modeling the 24 commodities that comprise the Grilli-Yang index as DS processes, it was found that just 5 had the negative trend predicted by the PS hypothesis.[3]

If y_t is truly generated by equation (2), the series contains a unit root, and standard tests of the null hypothesis H_0 : $\beta = 0$ based on equation (1) will suffer from severe size distortions, spuriously rejecting the null when no trend is

[1] Of course, economic decision making should account for costs as well as prices. For instance, it is quite possible that a long-run decline in prices is compensated by a long-run decline in marginal production costs. On the other hand, even given a positive long-run trend in prices, the substantial volatility of many commodity prices may hinder economic growth with difficulties in economic planning and disincentives to invest (see Blattman, Hwang, & Williamson, 2007).

[2] Spraos (1980), Sapsford (1985), Thirwall and Bergevin (1985), Grilli and Yang (1988), and Powell (1991) report results suggesting a deterioration in the terms of trade of commodity exporting developing countries, although not to the extent emphasized in Prebisch (1950) and Singer (1950).

[3] Similarly, Cuddington and Urzua (1989) found no deterioration in the terms of trade, but instead found that commodity prices fluctuated secularly around a stable mean (with a one-time break). For other studies on the long-run trends in commodity prices, see Powell (1991), Bleaney and Greenaway (1993), Labys (1993), Gafer (1995), Bloch and Sapsford (1997), Newbold and Vougas (1996), and Newbold, Rayner, and Kellard (2000). A good summary of this literature can be found in Greenaway and Morgan (1999) and Cuddington, Ludema, and Jayasuriya (2007).

present, even asymptotically. On the other hand, if the generating process is really equation (1), tests based on equation (2) are inefficient, lacking power relative to those based on equation (1). It is clear, therefore, that when the evidence as to the strength of the PS hypothesis is being assessed, the properties of the standard trend tests crucially depend on the integration properties of the commodity price series. Furthermore, if pretests for a unit root are first applied before adopting a trend test based on either equation (1) or (2), inference concerning the PS hypothesis is likely to be highly dependent on the results of the unit root pretests.

The situation is complicated because the true deterministic process underlying either equation (1) or (2) might also contain infrequent structural breaks in trend. For example, it is possible the correct generating process is a trend-stationary model with breaks:

$$y_t = \alpha + \beta t + \delta DU_t(\tau^*) + \gamma DT_t(\tau^*) + u_t,$$
$$t = 1, \ldots, T, \tag{3}$$

or, alternatively, a difference-stationary (about breaks) version:

$$\Delta y_t = \beta + \delta D_t(\tau^*) + \gamma DU_t(\tau^*) + \Delta u_t,$$
$$t = 2, \ldots, T, \tag{4}$$

where $DT_t(\tau^*) = 1(t > T^*)(t - T^*)$, $DU_t(\tau^*) = 1(t > T^*)$ and $D_t(\tau^*) = 1(t = T^* + 1)$, with $T^* = \lfloor \tau^* T \rfloor$ the (potential) break date with associated break fraction $\tau^* \in (0, 1)$, and where $1(.)$ denotes the indicator function and $\lfloor . \rfloor$ denotes the integer part of the argument. As in the case of testing for the presence of a linear trend, the properties of tests for the presence of a break in trend are also highly dependent on the order of integration of the series.

With regard to distinguishing between stationary and unit root behavior in a time series, neglecting a break in trend in an otherwise TS process can cause the spurious appearance of unit root behavior (see Perron, 1989), while a neglected trend break in a DS process can lead standard unit root tests to suggest an incorrect inference of stationarity (see Leybourne, Mills, & Newbold, 1998). Accounting for the former possibility, Leon and Soto (1997) and Kellard and Wohar (2006) apply unit root tests to relative commodity price series, allowing structural change under the TS alternative. For the 24 commodities of the Grilli-Yang index, and after allowance for (up to) one break, Leon and Soto classify twenty TS models for the 1900–1992 period, suggesting that shocks to commodity prices, in several cases, do not possess the permanent component suggested by Kim et al. (2003). Moreover, seventeen commodity prices report a negative trend and thus provide evidence in support of the PS hypothesis. Kellard and Wohar allow for (up to) two breaks, and although they find a similar number of TS processes, they point out that the negative trends reported by Leon and

Soto often exist only over some segment of the sample period, indicating less support for the PS hypothesis.

This paper contributes to the extant literature by examining evidence for the PS hypothesis using a new and much longer data set and seeks to ameliorate the effect of order of integration issues on the PS hypothesis testing procedure. First, we took the view that it would be both informative and interesting to use annual commodity price data from as far back as is sensibly possible. This resulted in the creation of a new unbalanced panel containing 25 relative commodity price series, 8 of which begin in 1650 (Beef, Coal, Gold, Lamb, Lead, Sugar, Wheat, Wool), 1 in 1670 (Cotton), 1 in 1673 (Tea), 2 in 1687 (Rice, Silver), 1 in 1709 (Coffee), 1 in 1741 (Tobacco), 1 in 1782 (Pig Iron), 3 in 1800 (Cocoa, Copper, and Hide), 1 in 1808 (Tin), 1 in 1840 (Nickel), 1 in 1853 (Zinc), 1 in 1859 (Oil), 1 in 1872 (Aluminum), and 2 in 1900 (Banana and Jute).[4] By contrast, the Grilli-Yang data set commences in 1900.[5] Second, powerful test procedures, robust to whether shocks are generated by an $I(0)$ or $I(1)$ process, for the presence of a linear trend (see Harvey, Leybourne, & Taylor, 2007) and a broken trend (see Harvey, Leybourne, & Taylor, 2009) are applied to the new data.

The remainder of the paper is organized as follows. Section II outlines the empirical methodology, and section III describes the new data. The empirical results and associated discussion are presented in section IV, and section V concludes.

II. Empirical Methodology

A. Testing for a Linear Trend

We are interested in testing the PS hypothesis where we have $H_0 : \beta = 0$ against a one-sided alternative $H_1 : \beta < 0$ but without assuming knowledge of whether u_t in equation (1) is $I(0)$ or $I(1)$. Harvey, Leybourne, and Taylor (2007; hereafter HLT, 2007), propose a relevant statistic based on taking a data-dependent weighted average of two trend statistics: one that is appropriate when the data are generated by an $I(0)$ process and a second when the data are $I(1)$. If, for example, it is known that u_t is $I(0)$, the appropriate trend statistic is the autocorrelation-robust t-ratio based on equation (1),

$$z_0 = \frac{\hat{\beta} - \beta_0}{s_0},$$
$$s_0 = \sqrt{\hat{\omega}_u^2 / \sum_{t=1}^{T} (t - \bar{t})^2}, \tag{5}$$

---

[4] Data are not available for some commodities from 1650 because they were not traded, not extracted, or not produced. Oil, for example, was not extracted for use in production before the mid-nineteenth century.

[5] Another good reason for the use of a longer historical sample is that the behavior of some commodity prices could perhaps be considered atypical over the twentieth century.

where $\hat{\alpha}$ and $\hat{\beta}$ denote the OLS estimators from equation (1) and $\hat{\omega}_u^2$ is the long-run variance estimator:

$$\hat{\omega}_u^2 = \hat{\gamma}_0 + 2 \sum_{j=1}^{T-1} h(j/\ell)\hat{\gamma}_j, \quad \hat{\gamma}_j = T^{-1} \sum_{t=j+1}^{T} \hat{u}_t \hat{u}_{t-j}, \qquad (6)$$

where $\hat{u}_t = y_t - \hat{\alpha} - \hat{\beta}t$, h(.) denotes the kernel function and ℓ the bandwidth. Conversely, if u_t is known to be $I(1)$, the appropriate trend statistic is the autocorrelation-robust t-ratio based on equation (2),

$$z_1 = \frac{\hat{\beta} - \beta_0}{s_1}, \qquad (7)$$

$$s_1 = \sqrt{\tilde{\omega}_v^2/(T-1)},$$

where $\tilde{\beta}$ is the OLS estimator of β in equation (2) and $\tilde{\omega}_v^2$ is the long-run variance estimator:

$$\tilde{\omega}_v^2 = \tilde{\gamma}_0 + 2 \sum_{j=1}^{T-2} h(j/\ell)\tilde{\gamma}_j, \qquad (8)$$

$$\tilde{\gamma}_j = (T-1)^{-1} \sum_{t=j+2}^{T} \tilde{v}_t \tilde{v}_{t-j},$$

where $\tilde{v}_t = \Delta y_t - \tilde{\beta}$. In both equations (6) and (8), following HLT (2007), we will use the quadratic spectral kernel with the Newey and West (1994) automatic bandwidth selection, adopting a nonstochastic prior bandwidth of $\lfloor 4(T/100)^{2/25} \rfloor$.

When it is not known a priori whether the series is $I(0)$ or $I(1)$, testing for a linear trend can be based on the weighted average of z_0 of equation (5) and z_1 of equation (7):

$$z_\lambda = \{1 - \lambda\}z_0 + \lambda z_1, \qquad (9)$$

where $\lambda \xrightarrow{p} 0$ when u_t is $I(0)$, while $\lambda \xrightarrow{p} 1$ when u_t is $I(1)$. HLT (2007) suggest the following exponential function for λ:

$$\lambda = \exp\left(-\left(\frac{U}{S}\right)^2\right), \qquad (10)$$

where U is a unit root statistic for testing the $I(1)$ null against the $I(0)$ alternative and S is a stationarity statistic for testing the $I(0)$ null against the $I(1)$ alternative. Given certain restrictions placed on U, S, and u_t, it can be shown that z_λ has a standard normal limiting distribution under the null.

With regard to the choices of U and S, HLT (2007) employ the local GLS-detrended augmented Dickey-Fuller t-test (DF-GLS$^\tau$) of Elliott, Rothenberg, and Stock (1996), while the KPSS test statistic ($\hat{\eta}_\tau$) of Kwiatkowski et al.

(1992) is chosen for S. Specifically, DF-GLS$^\tau$ is the usual t-ratio for testing $\rho = 0$ in the regression equation:

$$\Delta \tilde{u}_t = \rho \tilde{u}_{t-1} + \sum_{j=1}^{p} \phi_j \Delta \tilde{u}_{t-j} + \tilde{e}_t,$$

$$t = p + 2, \ldots, T, \qquad (11)$$

where $\tilde{u}_t$ are the local GLS detrended residuals obtained from the regression of $\mathbf{y}_{\bar{c}} = (y_1, y_2 - \bar{\rho}y_1, \ldots, y_T - \bar{\rho}y_{T-1})'$ on $\mathbf{Z}_{\bar{c}} = (z_1, z_2 - \bar{\rho}z_1, \ldots, z_T - \bar{\rho}z_{T-1})'$, where $z_t = (1, t)'$ and $\bar{\rho} = 1 - \bar{c}/T$ with $\bar{c} = -13.5$ (cf. Elliott et al., 1996). The number of lagged difference terms, p, included in equation (11) is determined by application of the autocorrelation-robust MAIC procedure of Ng and Perron (2001), setting the maximum lag length at $p_{\max} = \lfloor 12(T/100)^{1/4} \rfloor$. Notice that DF-GLS$^\tau$ is exact invariant to α and β. The KPSS statistic can be expressed as

$$\hat{\eta}_\tau = \frac{\sum_{t=1}^{T}(\sum_{i=1}^{t}\hat{u}_i)^2}{T^2 \hat{\omega}_u^2}, \qquad (12)$$

where the long-run variance estimator $\hat{\omega}_u^2$ is as defined in equation (6). Again, $\hat{\eta}_\tau$ is exact invariant to α and β. Finally, given these choices for U and S, HLT (2007) found the best finite sample performance was obtained by employing

$$\lambda = \exp\left(-0.00025\left(\frac{\text{DF-GLS}^\tau}{\hat{\eta}_\tau}\right)^2\right). \qquad (13)$$

Note that the constant 0.00025 does not affect the asymptotic properties of the z_λ test but gives rise to improved finite sample behavior.

B. Testing for a Broken Trend

Previous work has suggested that relative commodity prices may not be best represented by a single, secular trend but by some segmented alternative (see Kellard & Wohar, 2006). When examining the case for a breaking trend, this literature has, as in the single-trend context, relied on procedures that require pretesting for a unit root. In particular, Kellard and Wohar (2006) employed the test developed by Lumsdaine and Papell (1997), a procedure that allows shifts in the intercept and trend terms under the TS alternative hypothesis. Specifically, the structural breaks are endogenously chosen, using a search procedure, to maximize the chance of rejecting the unit root with drift null. As such, these unit root tests are not tests for structural change. Additionally, they do not allow the possibility of structural change under the null.

To circumvent the issues surrounding a pretest and to assess directly whether a trend contains a break, we require a test of $H_0 : \gamma = 0$ against a two-sided alternative $H_1 : \gamma \neq 0$ but without assuming knowledge of whether u_t in equation (3) is $I(0)$ or $I(1)$. Harvey et al. (2009; hereafter

HLT, 2009), in the methodological spirit of the previous subsection, propose such a test, based on a data-dependent weighted average of two individual statistics, one of which is appropriate when the stochastic component of the series is $I(0)$ and the other when it is $I(1)$. First, consider the case where only a break in trend occurs (model A in the notation of HLT, 2009). The appropriate model in the $I(0)$ case is equation (3) with $\delta = 0$:

$$y_t = \alpha + \beta t + \gamma DT_t(\tau^*) + u_t, \quad t = 1, \ldots, T, \qquad (14)$$

while the corresponding model in the $I(1)$ case is equation (4) with $\delta = 0$:

$$\Delta y_t = \beta + \gamma DU_t(\tau^*) + \Delta u_t, \quad t = 2, \ldots, T. \qquad (15)$$

Denote by $t_0(\tau^*)$ and $t_1(\tau^*)$ the autocorrelation-robust t-ratios for testing $\gamma = 0$ based on equations (14) and (15), respectively. For the implicit long-run variance estimators, we follow HLT (2009) and employ the Bartlett kernel $h(j/\ell) = 1 - j/(\ell + 1)$, with bandwidth parameter $\ell = \lfloor 4(T/100)^{1/4} \rfloor$. Now our interest is in the case where the true break fraction τ^* cannot be considered known a priori. We therefore need to search for the appropriate break point. HLT (2009) follow the approach of Andrews (1993) and consider statistics based on the maxima of the sequences of statistics $\{|t_0(\tau)|, \tau \in \Lambda\}$ and $\{|t_1(\tau)|, \tau \in \Lambda\}$, where $\Lambda = [\tau_L, \tau_U]$, with $0 < \tau_L < \tau_U < 1$, where the quantities τ_L and τ_U will be referred to as the trimming parameters, and where it is assumed throughout that $\tau^* \in \Lambda$. Defining $\Lambda^* = \{\lfloor \tau_L T \rfloor, \ldots, \lfloor \tau_U T \rfloor\}$, these statistics are given by

$$t_0^* = \sup_{s \in \Lambda^*} |t_0(s/T)|, \qquad (16)$$

and

$$t_1^* = \sup_{s \in \Lambda^*} |t_1(s/T)|, \qquad (17)$$

with the associated break-point estimators of τ^* given by $\hat{\tau} = \arg \sup_{s \in \Lambda^*} |t_0(s/T)|$ and $\tilde{\tau} = \arg \sup_{s \in \Lambda^*} |t_1(s/T)|$, respectively, such that $t_0^* \equiv |t_0(\hat{\tau})|$ and $t_1^* \equiv |t_1(\tilde{\tau})|$.

Given a lack of knowledge concerning the order of integration of the series, HLT (2009) then propose a test statistic based on the data-dependent weighted average of the supremum statistics for a broken trend under $I(0)$ and $I(1)$ shocks:

$$\begin{aligned} t_\lambda = &\{\lambda(S_0(\hat{\tau}), S_1(\tilde{\tau})) \times t_0^*\} \\ &+ m_\xi\{[1 - \lambda(S_0(\hat{\tau}), S_1(\tilde{\tau}))] \times t_1^*\}, \end{aligned} \qquad (18)$$

where m_ξ is positive finite constant and $S_0(\hat{\tau})$ and $S_1(\tilde{\tau})$ are auxiliary statistics chosen such that as the sample size diverges to positive infinity, the weight function $\lambda(.\,,\,.)$ converges to unity when u_t is $I(0)$ and to 0 when u_t is $I(1)$, such that t_λ will collapse to t_0^* when u_t is $I(0)$, and to t_1^* when u_t is $I(1)$. For $S_0(\hat{\tau})$ and $S_1(\tilde{\tau})$, HLT (2009) adopt the

stationarity test statistics of KPSS calculated from the relevant residuals of equations (14) and (15), respectively, estimated using the respective break date estimates $\hat{\tau}$ and $\tilde{\tau}$. The long-run variance estimators employed in the computation of the KPSS statistics again use the Bartlett kernel with bandwidth parameter $\ell = \lfloor 4(T/100)^{1/4} \rfloor$. Finally, HLT (2009) posit a weight function,

$$\lambda(S_0(\hat{\tau}), S_1(\tilde{\tau})) = \exp[-\{gS_0(\hat{\tau}), S_1(\tilde{\tau})\}^2], \qquad (19)$$

where g is a positive constant, since this will clearly converge to unity when u_t is $I(0)$ and to 0 when u_t is $I(1)$, as required. HLT (2009) show that the asymptotic null distribution of the weighted statistic t_λ of equation (18) differs as to whether u_t is $I(0)$ or $I(1)$; moreover, in neither case is this distribution standard normal. However, the constant m_ξ in equation (18) can be chosen such that for the selected significance level, the asymptotic null critical value of t_λ is the same regardless of whether u_t is $I(0)$ or $I(1)$. For the trend break tests to be operational, we also need to specify the constant g in equation (19). After Monte Carlo simulation of the finite sample size and power of the tests for a range of possible settings, HLT (2009) recommend the choice $g = 500$, giving rise to a t_λ test with both acceptable size and decent power across the range of simulation experiments considered.

HLT (2009) also consider a second model (model B), which extends model A by allowing for the possibility of a break in the level occurring simultaneously with the break in trend. The t_λ test is specified in exactly the same way as for model A, except now the appropriate models on which to base $t_0(.)$ and $t_1(.)$, and also the KPSS statistics, are equations (3) and (4), respectively.[6] Asymptotic critical values for the t_λ tests for both models A and B are provided in table 1 of HLT (2009), along with the corresponding values of m_ξ.[7] Following HLT (2009), we use 10% trimming, such that $\tau_L = 0.1$ and $\tau_U = 0.9$.

III. Data

The often employed Grilli-Yang (GY) data set comprises 24 internationally traded, nonfuel commodities.[8] Each annual nominal commodity price (in U.S. dollars) series is deflated by the United Nations Manufacturers Unit Value (MUV) index, the MUV series reflecting the unit values of manufacturing exports from a number of industrial countries. As noted in the introduction, the Grilli-Yang data set begins in 1900, primarily because this is the starting date for the MUV series; however, commodity and manufacturing price data can be sampled backward well before this time. Given the extensive interest in modeling and analyzing the

---

[6] The null hypothesis H_0 must now be restated as $H_0 : \gamma = \delta = 0$, in order to obtain a pivotal limiting null distribution for the test statistic.

[7] The choice of $g = 500$ also applies to model B.

[8] The commodities are Aluminum, Banana, Beef, Cocoa, Coffee, Copper, Cotton, Hide, Jute, Lamb, Lead, Maize, Palm Oil, Rice, Rubber, Silver, Sugar, Tea, Timber, Tin, Tobacco, Wheat, Wool, and Zinc.

long-run trends of relative commodity prices, it would appear important to utilize as many of the existing data as is sensibly possible.

Creating a large and representative data set of relative commodity prices prior to 1900 is not a trivial task. A number of disparate historical sources exist, covering different prices and sample periods.[9] A potential problem associated with construction of commodity price indexes is that movements in commodity import prices are not always synchronized across nations because some of the data include import duties and transport costs that vary across nations and over time. Furthermore, not all commodities are traded in markets for which spot or future price quotations for specified grades and quantities exist. An example is the oil market for which, at least until recently, the free market has been small. The cross-country variations in the growth of tariffs, import quotas, and transport costs were particularly large during World War I. A severe example is oil, for which prices were significantly lower in the United States than Europe due to the risk associated with sea transportation through the Atlantic Ocean and export embargoes.

The seriousness of these problems has been addressed in two papers. Based on very long historical data on commodities for the UK and Netherlands, Froot, Kim, and Rogoff (1995) find that the volatility and persistence of deviations from the law of one price have been stable over time, which suggests, at least for the UK and Netherlands, high comovements of commodity prices across nations. In the more recent study of Pesaran et al. (2006), some evidence of purchasing power parity among OECD countries is found for consumer goods such as meat, bread, tobacco, clothing, footwear, fruits, and other consumables covered in the consumer price index. These observations suggest that commodity prices change at different rates across nations; however, the difference is not significant. To get the most representative price for commodities, the average price for a specified commodity across nations, in common currency, is calculated for most commodities.

Pooling the various sources, a data set of 25 nominal primary commodity prices (in GBP) can be formed, but as a result of employing all available data, the series are of unequal lengths.[10] Specifically, twelve series begin in the seventeenth century (Beef, Coal, Cotton, Gold, Lamb, Lead, Rice, Silver, Sugar, Tea, Wheat, Wool), three series begin in the eighteenth century (Coffee, Tobacco, Pig Iron), eight series begin in the nineteenth century (Aluminum, Cocoa, Copper, Hide, Nickel, Oil, Tin, Zinc), and two start from 1900 (Banana and Jute). Twenty of these commodities are also found in the GY data set, and 23 are nonfuel.

Constructing a historical price index of manufactures (HPIM), stretching back to 1650, presented similar challenges to that of building the commodity price series; specifically, numerous sources and definitions of prices exist.[11] For example, while the important studies of Grilli and Yang (1988) and Lewis (1952) use manufacturing export unit indexes for selected industrialized countries and interpolate the data through the world war periods, we use the manufacturing value-added price deflator in the post-1870 period and various deflators for manufacturing products before then. The value-added price deflator has three advantages over export unit values: (a) it omits the influence of intermediate products; (b) it allows compositional changes; and (c) technological progress is to some extent reflected in the deflator. By contrast, export unit values, which are the measured manufacturing export value divided by the weight of export, fail to allow compositional changes in exports and innovation-induced price reductions. A value-preserving shift from exports of heavy manufactured metals to exports of electronics, for example, will artificially increase export unit values, even if prices have remained unaltered. Kravis and Lipsey (1984), for example, have advocated strongly against using export unit indexes because they exaggerate the long-term growth in manufacturing prices. However, manufacturing value-added prices are not free from measurement problems, especially because they do not fully allow for technological progress (see, for example, Griliches, 1979).

Unlike the production of services, manufacturing products are tradable and, as such, manufacturing price data are of relatively good quality (Griliches, 1979). However, the quality of the manufacturing value-added price deflators is likely to deteriorate as we go back in history. For the panel of countries included in the value-added deflator, data were not comprehensively available before 1870. Therefore, for our earliest periods, we use a composite index of prices of the most important manufacturing products. For example, from 1650 to 1784, the Dutch index is an unweighted average of textiles, soap, and paper, while the British index is composed of prices of various items such as leather backs, tallow candles, broadcloth (which is used for clothing and upholstering and is sold in large quantities in the world), beverages, linen, bread, oats, and stockings, among other products.

Given the above discussion, our manufacturing value-added price index is therefore spliced from subperiod series over the following periods: 1950–2005, 1870–1950, 1784–1870, and 1650–1784. Major industrialized countries are included in the index, and the data are converted into a common currency. The period 1950–2005 uses an unweighted average of manufacturing value-added price deflators for 22 industrialized countries (Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany,

---

[9] See the data appendix for a fuller description of the source of each price, available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest.2010.12184.

[10] GBP and not USD is used because the United States did not have its own currency before independence in 1776.

[11] Although it is possible to get data for commodity prices from before 1650, we could find no reliable source of manufacturing prices.

TABLE 1.—ONE-SIDED TESTS FOR A NEGATIVE TREND

	z_λ	Growth Rate (%)	90% c.i.	95% c.i.	99% c.i.
Aluminum	−2.133**	−3.908	±3.014	±3.591	±4.720
Banana	−0.991				
Beef	1.062				
Coal	1.110				
Cocoa	−0.928				
Coffee	−2.037**	−0.774	±0.625	±0.745	±0.979
Copper	−0.606				
Cotton	−0.659				
Gold	−0.231				
Hide	−0.659				
Jute	−1.531*	−1.479	±1.589	±1.893	±2.488
Lamb	0.743				
Lead	−0.571				
Nickel	−0.464				
Oil	−0.846				
Pig Iron	−0.493				
Rice	−1.224				
Silver	−1.516*	−0.823	±0.893	±1.065	±1.399
Sugar	−3.024***	−1.195	±0.650	±0.774	±1.018
Tea	−3.485***	−1.399	±0.660	±0.787	±1.034
Tin	0.287				
Tobacco	1.259				
Wheat	−1.265				
Wool	−1.318*	−0.653	±0.815	±0.972	±1.277
Zinc	−2.463***	−0.922	±0.616	±0.733	±0.964

Notes: *, **, and *** denote rejection at the 10%, 5%, and 1% significance levels, respectively. For tests that reject using the z_λ test, growth rates and two-sided confidence intervals are reported. These are obtained using equations (8) and (9) of HLT (2007). Testing against a two-sided alternative using z_λ does not lead to any more rejections of the no trend null. The test statistics are identical to those in the table, with critical values obtained from the standard normal distribution. None of the series with positive trends have significant coefficients. Two of the series (Jute and Wool) that reject at the 10% level against a one-sided alternative no longer reject when two-sided critical values are used.

Greece, Ireland, Italy, Japan, Luxembourg, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, UK, and United States. For the period 1870–1950, an unweighted average of manufacturing value-added price deflators for ten countries is used (Australia, Canada, Denmark, Finland, France, Germany, Japan, Sweden, UK, and United States). For the period 1784–1870, we use an unweighted average of various manufacturing price series for five major industrialized countries (France, Germany, Netherlands, UK, and United States). In the period 1650–1784, the index is constructed as an unweighted average of manufacturing prices in the Netherlands and the UK. These two nations were major trading nations during that period. Furthermore, based on recent data, Jurado and Vega (1994) find that the law of one price holds for manufacturing products, which suggests that the potential country selection bias is likely to be small, provided that the recent data are representative of the historical data.

It is of interest to consider how our historical price index of manufactures (HPIM) compares with the MUV index for the period since 1900, over which the MUV index is available. Over the entire twentieth century, our index increased by 3,384% while the MUV index increased 3,310%. In absolute terms, the difference is not large and thus is reflected in a very high correlation coefficient of 0.995. However, in relative terms, there are a few significant differences, most notably during the period 1914–1945, when the MUV index is often 25% below our index. This result suggests that export unit values are potentially biased measures of price movements, particularly when long data series are considered.

Finally, deflating the nominal commodity series with our manufacturing value-added price index resulted in a data set of relative commodity prices covering a 356-year period from 1650 to 2005.[12]

IV. Empirical Results

A. Trend Function Analysis

Table 1 shows the results of applying the order of integration robust trend tests to the new relative commodity price data set outlined in the previous section. In particular, column 2 gives the one-sided z_λ test statistic in equation (9) for each individual series. Notably, for eight commodities (Aluminum, Coffee, Jute, Silver, Sugar, Tea, Wool, and Zinc), the null of no trend is rejected (at least at the 10% level) in favor of the alternative of a negative trend. This seems a remarkable result considering the sample length of the commodities. The tea series, for example, commences in 1673 and has declined at an annual average rate of 1.40% (see column 3).

Of course, previous literature has detected structural breaks in the trend of relative commodity prices. Therefore, table 2 gives the results of applying the order of integration robust trend break tests to all 25 series based on employing model A. Specifically, columns 2, 3, and 4 give the two-sided t_λ test statistic in equation (18) at the 10%, 5%, and

[12] It is useful to note that we undertake the analysis using data only in GBP since the relative prices are invariant to currency denomination, that is, the conversion factor, is eliminated when we divide one series in GBP by another GBP-denominated series.

TABLE 2.—TWO-SIDED MODEL A TESTS FOR A BREAK IN TREND

	t_λ 10%	t_λ 5%	t_λ 1%	Estimated Break Date
Aluminum	1.987	2.030	2.118	
Banana	1.619	1.647	1.704	
Beef	1.202	1.228	1.281	
Coal	1.150	1.174	1.225	
Cocoa	1.291	1.319	1.376	
Coffee	1.357	1.386	1.446	
Copper	1.073	1.096	1.144	
Cotton	2.067	2.112	2.203	
Gold	0.896	0.916	0.955	
Hide	12.805*	12.819**	12.849***	1905
Jute	2.015	2.040	2.091	
Lamb	1.146	1.171	1.222	
Lead	1.562	1.591	1.650	
Nickel	1.517	1.546	1.605	
Oil	2.356*	2.397	2.481	1875
Pig Iron	1.233	1.260	1.314	
Rice	1.742	1.763	1.807	
Silver	1.744	1.782	1.859	
Sugar	3.655*	3.682**	3.736***	1951
Tea	1.323	1.352	1.410	
Tin	1.548	1.581	1.650	
Tobacco	13.125*	13.134**	13.153***	1951
Wheat	5.808*	5.832**	5.883***	1938
Wool	1.914	1.955	2.040	
Zinc	1.556	1.579	1.627	

Notes: *, **, and *** denote rejection at the 10%, 5%, and 1% significance levels respectively. t_λ denotes the model A test recommended by HLT (2009), using 10% trimming as in that paper. For tests that reject in favor of a break in trend, the estimated break date is reported. This is obtained using a weighted average of the implied levels and first difference break date estimators, as is done in the empirical application of HLT (2009).

1% levels, respectively. Only five of the series show evidence of a broken trend (Hide in 1905, Oil in 1875, Sugar in 1951, Tobacco in 1951, and Wheat in 1938). Of course, while model A allows a possible break in trend, it does not allow a simultaneous break in the level. To allow this possibility, the commodities were retested employing model B; however, a rejection of the no break null was obtained for only one further commodity (Jute in 1960), as reported in table 3.

In total, six price series now show a break in the trend function over the sample period. Although it is typically difficult to ascertain the causal factors, the preponderance of breaks located (close to or) in the twentieth century would perhaps suggest the effect of strong technological progress providing downward price pressure on many commodities including Oil (Castaneda, 2003), Wheat (Evenson & Kislev, 1973), Sugar (Swerling & Timoshenko, 1957), and Tobacco (Johnson, 1984). On the other hand, the 1960 break for Jute would seem to have its root in competition from petroleum-based synthetics, entering the market, and competing with jute for practically all of its uses (Grilli, 1975; Heitzman & Worden, 1989). Finally, the 1905 break in Hide is plausibly associated with increased meat consumption, leading to an increased supply of hides (Mack, 1956). To assess the

TABLE 3.—TWO-SIDED MODEL B TESTS FOR A BREAK IN TREND

	t_λ 10%	t_λ 5%	t_λ 1%	Estimated Break Date
Jute	2.908*	2.896	2.879	1960

Notes: * denotes rejection at the 10% significance level.
t_λ denotes the model B test recommended by HLT (2009), using 10% trimming as in that paper.

postbreak direction of the trend, the z_λ test was reapplied to the relevant commodities to the postbreak period of each of these six series; table 4 shows the results.

Notably, for five commodities, the null of no trend is rejected (at least at the 5% level) in favor of the alternative of a downward trend. In three cases (Hide, Tobacco, and Wheat), this is new evidence of a downward trend, revealed only when structural breaks are accounted for. Only for Oil is the null of no trend not rejected.

To summarize, eleven commodities (Aluminum, Coffee, Hide, Jute, Silver, Sugar, Tea, Tobacco, Wheat, Wool, and Zinc) are found to present a negative trend over all or some later fraction of their sample period; thus, the PS hypothesis appears to hold for a significant proportion of the primary commodities considered. The remainder (Banana, Beef, Cocoa, Copper, Cotton, Lamb, Lead, Rice, Tin, Pig Iron, Coal, Nickel, Gold, and Oil) all reveal a zero trend.

At this point, a useful comparison exercise is to apply the methodology employed above to the GY data set typically used in the literature.[13] Currently the latest incarnation of the GY data set contains annual data on real commodity prices from 1900 to 2003 (see Pfaffenzeller, Newbold, & Rayner, 2007). Table 5 provides results of application of the z_λ test in equation (9) to individual commodity prices in the GY data set, alongside those of the new data provided in our current paper (hereafter referred to as HKMW).[14]

Column 2 gives the z_λ results for the GY data, and for only three series (Aluminum, Rice, and Sugar) is the null of no trend rejected in favor of a negative trend. In contrast, column 4, which reports the z_λ results for the HKMW data, shows that thirteen commodities (Aluminum, Banana, Cotton, Hide, Jute, Lead, Rice, Silver, Sugar, Tea, Wheat, Wool, and Zinc) present a declining trend over the 1900–2003 period.

On the basis of the results in table 5, it would appear as if the GY data set is predisposed toward rejection of the PS hypothesis. To observe if this conclusion is maintained after allowing structural breaks, we applied the same procedure as before to the GY and HKMW data over the period 1900–2003: namely testing for a break in trend first using t_λ based on model A, and then t_λ based on model B. For series where breaks in trend were detected, we applied the z_λ test to the postbreak period. The outcome of this procedure was that negative trends were detected over the latter portion of the series for four further commodities when using the GY data (Banana, Coffee, Jute, and Lead) and three further commodities (Beef, Coffee, and Tin) when using the HKMW data. Thus, even after allowing the presence of structural breaks, the GY data reject the null of no trend far less frequently than does the HKMW data set. Overall, the GY data suggest that downward trends are present in seven

[13] We thank an anonymous referee for this helpful suggestion.
[14] Note that in this comparison section of the paper, the HKMW data set is curtailed to run only from 1900 to 2003, so as to be directly comparable to the GY data set.

TABLE 4.—ONE-SIDED TESTS FOR A NEGATIVE TREND (POSTBREAK)

	z_λ	Growth Rate (%)	90% c.i.	95% c.i.	99% c.i.
Hide	−17.101***	−2.214	±0.213	±0.254	±0.333
Jute	−3.635***	−4.466	±2.021	±2.408	±3.165
Oil	−0.084				
Sugar	−2.191**	−2.939	±2.207	±2.630	±3.456
Tobacco	−2.221**	−1.856	±1.375	±1.638	±2.153
Wheat	−1.771**	−1.987	±1.846	±2.200	±2.891

Notes: ** and *** denote rejection at the 5% and 1% significance levels respectively. Testing against a two-sided alternative does not lead to any more or fewer rejections of the no trend null, although, as would be expected, the significance levels are different in some cases.

TABLE 5.—ONE-SIDED TESTS FOR A NEGATIVE TREND USING GRILLI-YANG (GY) AND HKMW DATA, 1900–2003

	GY Data		HKMW Data	
	z_λ	Growth Rate (%)	z_λ	Growth Rate (%)
Aluminum	−1.958**	−1.654	−2.517***	−2.331
Banana	−0.204		−1.516*	−1.276
Beef	2.168		0.070	
Cocoa	−0.337		−1.079	
Coffee	−0.171		−1.000	
Copper	−1.035		−0.862	
Cotton	−1.028		−1.942**	−1.992
Hide	−0.652		−6.021***	−1.991
Jute	−0.727		−1.650**	−1.659
Lamb	3.956		1.023	
Lead	−1.060		−2.193**	−1.684
Rice	−1.938**	−1.350	−2.054**	−1.790
Silver	−0.181		−1.790**	−2.075
Sugar	−2.535***	−1.223	−4.611***	−2.449
Tea	−1.111		−1.859**	−1.956
Tin	−0.167		−1.172	
Tobacco	1.100		−0.352	
Wheat	−1.191		−2.870***	−1.813
Wool	−0.806		−2.524***	−1.885
Zinc	−0.276		−2.165**	−1.338

Note: *, **, and *** denote rejection at the 10, 5%, and 1% significance levels, respectively.

relative commodity price series, over either the sample period 1900–2003 or some postbreak subsample of this time span, while the HKMW data give rise to evidence in favor of negative trends for sixteen commodities.

B. Analysis of Cyclical Components

Finally, it is important to note that although this paper is primarily concerned with the issue of long-term trends, the identification of cycles in commodity prices has also been a popular theme in the literature.[15] In the seminal work of Kondratieff (1935), long waves or cycles of 45 to 60 years were posited. More recently, Cuddington and Jerrett (2008) employed the asymmetric bandpass filter of Christiano and Fitzgerald (2003) to decompose metals prices into three components:

$$y_t = T_t + LC_t + SC_t, \ t = 1, \ldots, T, \tag{20}$$

where T_t is a trend component, LC_t is a long-term cyclical component, and SC_t is a short-term cyclical component. Christiano and Fitzgerald propose a finite sample approxi-

[15] Again, we thank an anonymous referee for this helpful point.

mation for the bandpass filter that is nearly optimal for typical economic time series. The filter extends the bandpass approach developed by Baxter and King (1999) and allows computation of cyclical components for the full span of the time series without trimming at the end points. It can be applied to series regardless of whether they contain a unit root or are (trend) stationary and has recently been employed by a number of authors (see, for example, Fuhrer and Rudebusch, 2004; Fisher, 2006; and Fernald, 2007). Following Cuddington and Jerrett (2008) we define long-term cycles as those lasting from twenty to seventy years, thus nesting those proposed by Kondratieff (1935). Additionally, the trend is defined as all cyclical components lasting seventy years or longer and short-term cycles as those ranging from two to twenty years. Denoting by $CF(c_1, c_2)$ the Christiano and Fitzgerald bandpass filter that passes cycles between c_1 and c_2 in length, we can define T_t, LC_t, and SC_t as follows:

$$T_t = CF(70, \infty) \tag{21}$$

$$LC_t = CF(20, 70) \tag{22}$$

$$SC_t = CF(2, 20). \tag{23}$$

The first step of the Christiano and Fitzgerald filter is to detrend the data under the assumption that a unit root is present in the data; specifically, the estimate of the trend coefficient is the OLS estimate of β in equation (2). For the series where breaks in trend were detected, we modify this first step by including the break in the deterministic component used for the detrending. The estimates of the trend and break parameters are again obtained under an assumption that a unit root is present in the stochastic component of the series—from OLS estimation of either equation (15) for model A or equation (4) for model B.

Table 6 provides some summary measures of the cyclical components when the Christiano and Fitzgerald filter is applied to individual commodity price series of the full (1650–2005) HKMW data set.

Column 2 shows the standard deviation of the long-term cyclical component (LC_t) and column 3, the ratio of the standard deviations of LC_t and the total nontrend cyclical component ($SC_t + LC_t$). The range of this ratio extends from 0.48 (Wheat) to 0.90 (Aluminum), clearly indicating the primacy of the long-term component in cyclical com-

TABLE 6.—SUMMARY MEASURES OF TWENTY- TO SEVENTY-YEAR CYCLICAL
COMPONENTS (LC_T)

	s.d. (LC_t)	$\dfrac{s.d.(LC_t)}{s.d.(SC_t + LC_t)}$	Mean Periodicity	AR(1) Parameter
Aluminum	0.354	0.900	25.750	0.982
Banana	0.095	0.684	25.333	0.978
Beef	0.167	0.759	25.727	0.981
Coal	0.096	0.670	25.917	0.980
Cocoa	0.261	0.783	23.714	0.975
Coffee	0.185	0.680	28.111	0.980
Copper	0.177	0.749	43.333	0.991
Cotton	0.195	0.640	26.545	0.978
Gold	0.172	0.878	27.455	0.985
Hide	0.121	0.581	30.500	0.985
Jute	0.135	0.531	25.000	0.985
Lamb	0.196	0.812	24.500	0.980
Lead	0.133	0.689	25.417	0.985
Nickel	0.198	0.711	25.800	0.980
Oil	0.367	0.726	26.250	0.974
Pig Iron	0.161	0.749	30.200	0.980
Rice	0.180	0.676	23.545	0.984
Silver	0.150	0.730	32.000	0.994
Sugar	0.146	0.593	26.909	0.980
Tea	0.137	0.659	25.000	0.986
Tin	0.188	0.728	28.800	0.988
Tobacco	0.131	0.709	37.167	0.984
Wheat	0.103	0.481	29.800	0.986
Wool	0.106	0.531	26.727	0.982
Zinc	0.175	0.725	35.667	0.982

modity price movements. Such results strongly underline the relevance of Kondratieff's identification of long-term cycles in commodity prices, although it should be noted that the mean periodicity of the cycles in LC_t (see Table 6, column 4) ranges from 23.5 years (Rice) to 43.3 years (Copper), clearly lower than the 45- to 60-year interval consonant with Kondratieff cycles. In any case, the existence of cyclical components lasting longer than 20 years suggests that common policy initiatives to smooth either commodity prices themselves or producer or consumer incomes around a trend may require economic planning over an extremely long time horizon. Finally, column 5 of table 6 reports the first-order autocorrelation measure of persistence in the LC_t component of each series—that is, the parameter estimate $\hat{\phi}$ in the AR(1) regression $LC_t = \mu + \phi LC_{t-1} + \varepsilon_t$ estimated by least squares. The estimates are all in excess of 0.97, showing a very high degree of persistence in the long-term cyclical components of the relative primary commodity price series.

The presence of persistent, long-term cycles in real commodity prices over the past four centuries informs a number of contemporary economic issues. From a research perspective, it would suggest that the hundred years or so of data employed in previous studies of commodity prices is probably too short to adequately assess cyclical behavior. Furthermore, although obvious, it is useful to highlight that the theory behind the much investigated PS hypothesis is silent with respect to any cyclical determination. More theoretical and empirical work in the area of real price commodity price cycles is encouraged after our long-run data analysis.

From a policymaker perspective, we might posit that the conventional view of real commodity prices as presenting (a) a negative or zero trend and (b) relatively high volatility around that trend, can be augmented by (c) cycles of lower and higher prices over long horizons. Although short-term movements in prices can be hedged using options or futures, no such market-based instruments currently exist for such longer-term movements. Clearly, during periods of rising (the recent context) or falling commodity prices, policymakers need to be keenly aware that prices may not return to equilibrium for many years.

V. Conclusion

Primary commodity production contributes a significant fraction of the export volume of many developing countries. Given this context, the time-series properties of such prices, relative to manufactured goods, have important policy implications. A negative trend, for example, in the relative price of a country's main export commodity indicates the need to consider diversifying the export mix.

The literature presents the consensual position that the typically large variance of relative commodity prices makes it difficult to ascertain the existence of a trend (see Deaton, 1999, and Cashin & McDermott, 2002). However, this has not inhibited academic study on the issue and, commencing with the seminal work of Prebisch (1950) and Singer (1950), debate has raged as to whether relative commodity prices actually suffer from long-run secular decline.

Given the subjugation of the trend of prices by the variance of prices, the empirical results with respect to trend existence and direction are unsurprisingly mixed. In particular, the results are often conditional on the assumed order of integration of the relative price processes. Since the properties of standard tests for trends are highly dependent on whether the series in question contains a unit root, it is difficult to draw unambiguous conclusions regarding the presence of trends using standard trend tests alone. This situation is further problematized by the possibility that structural breaks may occur in the underlying mean or trend function. If such breaks occur but are not adequately modeled, further errors of inference can arise.

This paper makes a number of new contributions to the literature while assessing the evidence for a long-run trend in primary commodity prices. An entirely new data set of 25 major primary commodity prices, relative to manufactures, is assembled by pooling, for the first time, numerous historical data sources. This can be compared with the commonly employed Grilli-Yang data set, which contains 24 commodities, at an annual frequency, stretching back to 1900. After an exhaustive search of the available sources, the new data set contains data from 1650, providing not only historical interest but many more degrees of freedom with which to disentangle any trend component from its variance. Specifically, twelve series begin in the seventeenth century

(Beef, Coal, Cotton, Gold, Lamb, Lead, Rice, Silver, Sugar, Tea, Wheat, Wool), three series begin in the eighteenth century (Coffee, Tobacco, Pig Iron), eight series begin in the nineteenth century (Aluminum, Cocoa, Copper, Hide, Nickel, Oil, Tin, Zinc), and two start from 1900 (Banana and Jute).

In addition to the new data, this paper also applies new time-series techniques by Harvey et al. (2007, 2009) to assess the trend function and the existence of any possible structural breaks. The tests are based on a data-dependent weighted average of the relevant statistics when the stochastic component of the time series is assumed to be $I(0)$ and $I(1)$, respectively. In this manner, the tests employed are robust to the order of integration issues that have plagued the literature. Our empirical methodology thus requires no unit root pretests and simply tests the trend function directly. The empirical results are informative. Initially, we examined the trend function of each commodity price series without allowing for the possibility of structural breaks. It was found that eight commodities (Aluminum, Coffee, Jute, Silver, Sugar, Tea, Wool, and Zinc) present a secular downward trend. As a specific example, consider that the relative price of an important commodity like coffee has been declining at an annual rate of 0.77% for approximately 300 years! Second, we tested each price series for a break in the trend function, and for the series that rejected in favor of a break, the postbreak period was reanalyzed for the presence of a downward trend. This resulted in a further three commodities being found to present a negative trend (Hide, Tobacco, and Wheat).

Overall, eleven major commodities show new and robust evidence of a long-run decline in their relative price. In our opinion, this provides much more robust support that the Prebisch-Singer hypothesis is relevant for commodity prices. For the remaining fourteen commodities, no positive and significant trends could be detected over all or some fraction of the sample period. These zero-trending commodities suggest that the Lewis hypothesis may also play a part in explaining the behavior of certain commodity prices; conversely, however, in the very long run, there is simply no statistical evidence that relative commodity prices have ever trended upward.

REFERENCES

Andrews, Donald W. K., "Tests for Parameter Instability and Structural Change with Unknown Change Point," *Econometrica* 61 (1993), 821–856.

Baxter, M., and R. G. King, "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," this REVIEW 81 (1999), 575–593.

Blattman, C., J. Hwang, and J. G. Williamson, "Winners and Losers in the Commodity Lottery: The Impact of Terms of Trade Growth and Volatility in the Periphery 1870–1939," *Journal of Development Economics* 82 (2007), 156–179.

Bleaney, M., and D. Greenaway, "Long-Run Trends in the Relative Price of Primary Commodities and in the Terms of Trade of Developing Countries," *Oxford Economic Papers* 45 (1993), 349–363.

Bloch, H., and D. Sapsford, "Some Estimates of Prebisch and Singer Effects on the Terms of Trade between Primary Producers and Manufacturers," *World Development* 25 (1997), 1873–1884.

Cashin, P., and C. J. McDermott, "The Long-Run Behaviour of Commodity Prices: Small Trends and Big Variability," *IMF Staff Papers* 49 (2002), 175–199.

Castaneda, C., "Oil Industry: Technological Change" (pp. 126–129), in J. Mokyr (Ed.), *The Oxford Encyclopedia of Economic History* (New York: Oxford University Press, 2003).

Christiano, L., and T. J. Fitzgerald, "The Bandpass Filter," *International Economic Review* 44 (2003), 435–465.

Cuddington, J., and D. Jerrett, "Super Cycles in Metals Prices?" *IMF Staff Papers* 55:4 (2008), 541–565.

Cuddington, J., R. Ludema, and S. Jayasuriya, "Prebisch-Singer Redux," in D. Lederman and W. Maloney (Eds.), *Natural Resources: Neither Curse nor Destiny* (Stanford, CA: Stanford University Press, 2007).

Cuddington, J., and C. Urzua, "Trends and Cycles in the Net Barter Terms of Trade: A New Approach," *Economic Journal* 99 (1989), 426–442.

Deaton, A., "Commodity Prices and Growth in Africa," *Journal of Economic Perspectives* 13 (1999), 23–40.

Elliott, G., T. J. Rothenberg, and J. H. Stock, "Efficient Tests for an Autoregressive Unit Root," *Econometrica* 64 (1996), 813–836.

Evenson, R. E., and Y. Kislev, "Research and Productivity in Wheat and Maize," *Journal of Political Economy* 81 (1973), 1309–1329.

Fernald, J. G., "Trend Breaks, Long-Run Restrictions, and Contractionary Technology Improvements," *Journal of Monetary Economics* 54 (2007), 2467–2485.

Fisher, J. D. M., "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy* 114 (2006), 413–451.

Froot, K., M. Kim, and K. Rogoff, "The Law of One Price over 700 Years," NBER working paper no. 5132 (1995).

Fuhrer, J. C., and G. D. Rudebusch, "Estimating the Euler Equation for Output," *Journal of Monetary Economics* 51 (2004), 1133–1153.

Gafer, J., "Recent Trends in the Terms of Trade of Jamaica: 1955–86," *Applied Economics* 27 (1995), 161–165.

Greenaway, D., and C. W. Morgan (Eds.), *The Economics of Commodity Markets* (Aldershot, UK: Edward Elgar, 1999).

Griliches, Z., "Issues in Assessing the Contribution of Research and Development to Productivity Growth," *Bell Journal of Economics* 10 (1979), 92–116.

Grilli, E., *The Future for Hard Fibers and Competition from Synthetics* (Baltimore, MD: Johns Hopkins University Press, 1975).

Grilli, R. E., and M. C. Yang, "Commodity Prices, Manufactured Goods Prices, and the Terms of Trade of Developing Countries," *World Bank Economic Review* 2 (1988), 1–48.

Harvey, D. I., S. J. Leybourne, and A. M. R. Taylor, "A Simple, Robust and Powerful Test of the Trend Hypothesis," *Journal of Econometrics* 141 (2007), 1302–1330.

—— "Simple, Robust and Powerful Tests of the Breaking Trend Hypothesis," *Econometric Theory* 25 (2009), 995–1029.

Heitzman, J., and L. Worden, "Bangladesh: A Country Study" (Washington, DC: Federal Research Division, Library of Congress, 1989).

Johnson, P., *The Economics of the Tobacco Industry* (New York: Praeger, 1984).

Jurado, M., and J. Vega, "Purchasing Power Parity: An Empirical Analysis," *Investigaciones Económicas* 18 (1994), 539–556.

Kellard, N., and M. E. Wohar, "On the Prevalence of Trends in Primary Commodity Prices," *Journal of Development Economics* 79 (2006), 146–167.

Kim, T., S. Pfaffenzeller, A. Rayner, and P. Newbold, "Testing for Linear Trend, with Application to Relative Primary Commodity Prices," *Journal of Time Series Analysis* 24 (2003), 539–551.

Kondratieff, N. D., "The Long Waves in Economic Life," this REVIEW 17 (1935), 105–115.

Kravis, I. B., and R. E. Lipsey, "Prices and Terms of Trade for Developed Country Exports of Manufactured Goods," in B. Csikos-Nagy, D. Hague, and G. Hall (Eds.), *The Economics of Relative Prices* (London: Macmillan, 1984).

Kwiatkowski, D., P. C. B. Phillips, P. Schmidt, and Y. Shin, "Testing the Null Hypothesis of Stationarity against the Alternative of a Unit

Root: How Sure Are We That Economic Time Series Have a Unit Root?" *Journal of Econometrics* 54 (1992), 159–178.

Labys, W. C., "The Changing Severity and the Impact of Commodity Price Instability," in M. Nissanke, and A. Hewitt (Eds.), *Economic Crises in Developing Countries* (London: Pinter Publishers, 1993).

Leon, J., and R. Soto, "Structural Breaks and Long-Run Trends in Commodity Prices," *Journal of International Development* 9 (1997), 347–366.

Lewis, W. A., "World Production, Prices, and Trade, 1870–1960," *Manchester School of Economics and Social Studies* 20 (1952), 105–138.

―――― "Economic Development with Unlimited Supplies of Labour," *Manchester School of Economics and Social Studies* 22 (1954), 139–191.

Leybourne, S. J., T. C. Mills, and P. Newbold, "Spurious Rejections by Dickey-Fuller Tests in the Presence of a Break under the Null," *Journal of Econometrics* 87 (1998), 191–203.

Lumsdaine, R. L., and D. Papell, "Multiple Trend Breaks and the Unit Root Hypothesis," this REVIEW 79 (1997), 212–218.

Mack, R., *Consumption and Business Fluctuations: A Case Study of the Shoe, Leather, Hide Sequence* (Cambridge, MA: NBER Books, 1956).

Newbold, P., A. J. Rayner, and N. Kellard, "Long-Run Drift, Comovement, and Persistence in Real Wheat and Maize Prices," *Journal of Agricultural Economics* 51 (2000), 106–121.

Newbold, P., and D. Vougas, "Drift in the Relative Prices of Primary Commodity Prices: A Case Where We Care about Unit Roots," *Applied Economics* 28 (1996), 653–661.

Newey, W. K., and K. D. West, "Automatic Lag Selection in Covariance Matrix Estimation," *Review of Economic Studies* 61 (1994), 631–653.

Ng, S., and P. Perron, "Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power," *Econometrica* 69 (2001), 1519–1554.

Perron, P., "The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis," *Econometrica* 57 (1989), 1361–1401.

Pesaran, H., R. Smith, T. Yamagata, and L. Hvozdyk, "Pairwise Tests of Purchasing Power Parity Using Aggregate and Disaggregate Price Measures," CWPE working paper no. 0634 (2006).

Pfaffenzeller, S., P. Newbold, and A. Rayner, "A Short Note on Updating the Grilli and Yang Commodity Price Index," *World Bank Economic Review* 21 (2007), 151–163.

Powell, A., "Commodity and Developing Countries Terms of Trade: What Does the Long-Run Show?" *Economic Journal* 101 (1991), 1485–1496.

Prebisch, R., "The Economic Development of Latin America and Its Principal Problems," *Economic Bulletin for Latin America* 7 (1950), 1–12.

Sapsford, D., "The Statistical Debate on the Net Terms of Trade between Primary Commodities and Manufacturers: A Comment and Some Additional Evidence," *Economic Journal* 95 (1985), 781–788.

Singer, H., "The Distribution of Gains between Investing and Borrowing Countries," *American Economic Review, Papers and Proceedings* 40 (1950), 473–485.

Spraos, J., "The Statistical Debate on the Net Barter Terms of Trade," *Economic Journal* 90 (1980), 107–128.

Swerling, V., and V. Timoshenko, *The World's Sugar: Progress and Policy* (Stanford, CA: Stanford University Press, 1957).

Thirwall, A. P., and J. Bergevin, "Trends, Cycles, and Asymmetry in Terms of Trade," *World Development* 13 (1985), 805–817.