A Weak Latent Trend Hides Strong Price Predictability: An Empirical Method For An Unrecognized Problem.

Eugenio S. A. Bobenrieth[†] Juan R. A. Bobenrieth[‡]
Brian D. Wright[♣] Ernesto A. Guerra[§]

[†] Pontificia Universidad Católica de Chile. Corresponding author. ebobenrieth@uc.cl
[‡]Departamento de Matemática, Universidad del Bío-Bío. jbobenri@ubiobio.cl
[♣] Department of Agricultural and Resource Economics, University of California, Berkeley. bwright@berkeley.edu
[§] Departamento de Economía, Universidad Católica de la Santísima Concepción. eguerra@ucsc.cl

Selected Paper prepared for presentation at the 2022 Agricultural & Applied Economics Association Annual Meeting, Anaheim, CA; July 31-August 2.

A Weak Latent Trend Hides Strong Price Predictability: An Empirical Method For An Unrecognized Problem.

Eugenio S. A. Bobenrieth[†] Juan R. A. Bobenrieth[‡]

Brian D. Wright[♣] Ernesto A. Guerra[§]

Abstract

A simple mapping of current annual commodity price to the subsequent price has been presented as evidence that rejects the standard nonlinear storage model and implies that current price is uninformative about expected price changes. We show how a weak trend can transform a price series with strongly predictable price crashes, so that price changes appear to be unpredictable. We rationalize such inference using a dynamic model including speculative inventories and random production with a weak productivity trend. The model has two regimes. In one, price changes follow a stochastic trend with a positive drift independent of the productivity trend. In the other, the expected price change is a predictable price-dependent jump to its conditional price expectation, a price target following a weak deterministic trend induced by the weak productivity trend. We illustrate our results using samples of commodity prices for which price autoregressions appear obviously linear and price changes appear to be non-predictable, as in simple unit root processes. For our empirical estimations we implement a novel asymptotically normal one-step estimator of the trend, the interest rate, and the price threshold of the nonstationary storage model, that solves a delicate problem of discontinuity of the regression in the limit. Our results show the relevance of discriminating the proper trend for each data point in the sample.

Keywords: predictability of returns, inventories, trends, jumps, estimation.

JEL codes: C51, C52, C53, D84.

[†] Pontificia Universidad Católica de Chile. Corresponding author. ebobenrieth@uc.cl

[‡]Departamento de Matemática, Universidad del Bío-Bío. jbobenri@ubiobio.cl

[♣] Department of Agricultural and Resource Economics, University of California, Berkeley. bwright@berkeley.edu

[§] Departamento de Economía, Universidad Católica de la Santísima Concepción. eguerra@ucsc.cl

I. Introduction

The volatility of prices of primary commodities makes their medium-term prediction an important issue. The return to public or private commodity projects often depends on their price prospects over the next half-decade or so, beyond the horizon of liquid futures or forward trading. For a commodity-dependent exporting country, the price prospects for a single commodity can strongly affect its international borrowing capacity; for some countries, the cost of energy imports is also highly relevant for their immediate future.

Predictions based on observed commodity price series receive great attention because, unlike inventories, they are accurately measured on liquid centralized markets, and widely available. Many commodity price series seem to exhibit occasional sharp price spikes followed by highly predictable reversions. But a simple "acid test" proposed by Deaton and Laroque (1996) and Deaton (2010) supports their claim that annual price changes are essentially unpredictable based on current price, falsifying the principal model of medium term price dynamics for markets for storable commodities, and discouraging pursuit of their own pioneering work on estimation of this model.

In this paper we contest this conclusion, showing that neglected small trends can linearize price autoregressions, inducing the illusion that forward price returns are non-predictable. Surprisingly unacknowledged in the literature, recognition of a small trend can reveal a non-linear relationship consistent with intertemporal arbitrage of inventories, implying episodes of strong predictability of price changes.

In a naive preliminary data exploration, next section we verify the non-predictability of annual price changes based on current price using illustrative samples of observed cotton and maize prices. However, adjustment of base prices for a deterministic trend estimated in a preliminary step can generate predictions of spreads qualitatively similar to those implied by the standard stationary storage model. This result motivates us to use a model of nonstationary storage arbitrage for empirical analysis of these samples.

We present in Section III such a model. The model includes a latent productivity trend that induces a partially latent trend in price. In this model, price has a partially latent trend induced by a latent productivity trend, consumption demand is stationary, and the real interest rate is constant. There are two price regimes. In one, inventories are positive, and in the absence of storage costs other than the financial cost, intertemporal storage arbitrage implies that the conditional expectation of price increases at the interest rate regardless of the downward long-run secular trend. In this regime price is a renewal process with a stochastic trend with positive drift equal to the interest rate, independent of the production trend. Storage decreases as current price rises, terminating when price reaches an endogenous trending threshold.

In the second regime, inventories are constrained at zero. In such regime the expected price follows a deterministic attractor which has an unbroken trend induced by the deterministic secular trend in production. The one-period discounted value of this trending attractor is the trending real price threshold that separates the regimes. This regime terminates when price passes below this trending threshold.

We use the model to show how a weak unbroken productivity trend might have no effect at all on the conditional expected price change for most sample observations. Yet that trend can have implications regarding the predictability of the returns, including the size, direction and predictability of price jumps, and trend breaks that are not modest at all.

We numerically calibrate the trending model with a parameterization for consumption demand and interest rate originally chosen by Wright and Williams (1982) to represent behavior typical of a stationary grain market, and used in subsequent papers including

Deaton and Laroque (1992) and Cafiero et al. (2011). We introduce a modest downward trend in the price attractor consistent with an exogenous upward trend in yield, as discussed in Section III. Simulation of this model in Section IV reproduces the qualitative behavior of returns discussed above, and replicates their highly nonlinear behavior observed after detrending the base prices.

Our simulations imply a unimodal distribution for detrended price, as seen in the maize price sample, detrended using our naive preliminary detrending. For cotton, detrended prices using preliminary detrending does not show a clear unimodal distribution. Does the apparent inconsistency of the detrended cotton price distribution imply rejection of this storage model, or a problem with the conventional approach of relying on preliminary detrending of the data prior to estimation of the model?

To answer this question we turn in Section V to a strategy for simultaneous estimation of three key parameters, the deterministic exponential trend, the interest rate and the real detrended price threshold, on nonstationary price data. Our estimation allows for interactions between the trend and other key parameters, and yields asymptotic standard errors of the estimates. But the fact that we are estimating other parameters along with the trend parameter implies a delicate problem of discontinuity of the regression in the limit, which prevents application of available proofs of consistency and asymptotic normality in this non-stationary context. To solve this problem we appeal to the novel proofs of consistency and asymptotic theory in Bobenrieth, Bobenrieth and Wright (2022).[1]

After a description of our illustrative samples of cotton and maize in Section VI, Section VII presents one-step estimates of the three key parameters for our samples of cotton and maize prices.Our results imply revisions of assumptions in the literature regarding conditional

[1] The proofs are in Bobenrieth, Bobenrieth and Wright (2022), available from the corresponding author.

price persistence, and the relative predictability of price changes during booms and slumps. For our nonstationary price samples, our approach not only allows calculation of standard errors of our estimates, but furnishes predictions of price changes superior to those obtained by a strategy employing the usual practice of detrending prior to estimation. Our estimate of the trending price attractor (a function of the trend, the interest rate and the price threshold), is key to prediction of price changes. It allows us to predict the dynamics of prices and to distinguish the regime-dependence of price volatility.

Our empirical model must account for the effect of the trend on arbitrage activity. Osborne (2004), Roberts and Schlenker (2013), Guerra et al. (2015), Gouel and Legrand (2017, 2021), and Jansen, Smith, and Carter (2018) recognize the existence of trends in commodity price data but they do not address their effects on storage arbitrage or on the consistency of their estimation procedures.

Bobenrieth et al. (2021) focus on the effect of trends on measured price autocorrelation. They detrend the price data in a preliminary step, and recognize the effects of trends on inventory decisions in a second step of their two-step estimation procedure. But they are not able to provide asymptotic standard errors for their estimates. We use the theoretical model in Bobenrieth et al. (2021) and focus on the linearization effects of trends, and the empirical relevance of one-step estimation. In contrast to their two-step empirical approach, here in a single step we estimate three key parameters of the Euler equation, the trend, the interest rate, and an endogenous threshold that separates two price regimes, using only the observed nonstationary real price data, and provide estimated standard errors.

We demonstrate the empirical importance of simultaneous estimation of the trend and other key variables by comparison of the histograms of price distributions detrended using trends estimated ex ante with those detrended using our new econometric procedure, which estimates trends simultaneously with the other parameters. Using the concept of entropy from information theory, the distribution of detrended cotton prices from our new trend estimate is more informative than that constructed from prior detrending.[2]

There is a large literature on estimation of trending time series with exogenous trend breaks, see for example Perron (1989), Perron and Zhu (2005), Kim and Perron (2009), Harvey and Leybourne (2015). Harvey et al. (2010) and Gouel and Legrand (2017, 2021) offer reviews related to commodity prices. Part of that literature discusses procedures for estimating trends with breaks that are denoted endogenous because the timing of breaks is determined by the data.

In contrast, we assume no breaks in the deterministic trend followed by the attractor. Observing large price jumps in our illustrative samples, some might infer exogenous breaks in trends, when in fact the jumps reflect the endogenous initiation or termination of the stochastic trend with positive drift equal to the interest rate, which is a renewal process in our model.

II. A naive preliminary

We begin by demonstrating the non-predictability of forward spreads based on observed prices using illustrative samples of real prices of cotton and maize shown in Figure 1.

Figure 1 here

[2]The distributions of detrended maize prices using preliminary detrending and using our one-step detrending are not significantly different, according to the Kolmogorov-Smirnov (KS) test for the distance of empirical distributions.

For both samples, Figure 2 verifies the negative result of the "acid test," showing that a mapping of current real annual price to the subsequent price realization appears strictly increasing and approximately linear, and with a slope close to the 45 degree line.[3] The conditional mean of price next year appears to be close to current price, whether or not the latter is high or low, similar to the behavior of a unit root process, with a modest drift. Figure 3 presents a different perspective that makes it clear that the current real price level is strikingly uninformative regarding the mean and dispersion of the one-year change in price, for both cotton and maize.

Figure 2 and Figure 3 here

These figures do not negate the possibility of a secular trend. But any such trend, conceived as a constant contribution to price change each year is clearly modest relative to annual price variability, and has routinely been neglected in analyses of year-to-year price dynamics. As summarized by Cashin and McDermott (2002, p.175),

> "Although there is a downward trend in real commodity prices, this is of little practical policy relevance, since it is small and completely dominated by the variability of prices."

But is there actually a persistent stochastic or deterministic trend in these prices over the sample interval? One difficulty is that there might be alternating trends over long periods of time, resulting in apparent stationarity over the long run. As noted by Deaton and Laroque (2003, p. 291) for many commodity prices:

[3] We address in Section VI the two highest maize prices that might seem to be "outliers" in the lower panel of Figure 2.

> “it is possible to see upward or downward “trends” over prolonged periods. Over the whole period, the price is at least roughly consistent with its being a stationary time-series.”

The challenge is well illustrated in our price samples for cotton and maize, presented in Figure 1. For cotton, there is obviously a persistent upward movement starting at 1938, dramatically reversed after 1950. Only around 1970 did price return to its 1938 level. The maize series has a dramatic early jump followed by persistent upward and downward movements that result in little net change between the 1938 price and the prices in the 1960s, followed by a volatile but declining path at least through the 1990s.

Note that thus far our exploration has been entirely nonparametric. We now show that a modest secular trend might have no effect at all on the conditional expected price change for most sample observations. Yet that trend can have implications regarding the predictability of the returns reported in Figure 3, including the size, direction and predictability of price jumps, and indeed the number of trend breaks inferred from Figure 1, that are not modest at all.

We now use an estimated simple deterministic exponential trend for each price sample in Figure 1, and create a new “detrended” base price series. Reordering the same observed returns shown in Figure 3 using these new base prices, detrended using trends estimated ex ante, reveals evidence, shown in Figure 4, of nonlinear non-linear relationships. At low detrended base prices, mean returns (calculated on the trending data) still appear to be on average positive but highly variable. But at high “detrended” base prices, returns are predominantly negative jumps. The correspondences in Figures 2 and 3 show little evidence of predictability of price changes because, as discussed below, a neglected secular trend scrambles signals conveyed by the base prices. Thus, the secular trend creates the illusion of unpredictability of price changes.

This reordering transforms the empirical histograms, plotted at the horizontal axes in Figure 3 to those depicted in Figure 4.

Figure 4 here

Notice that the histogram for detrended maize prices now appears more consistent with a unimodal distribution with thinner tails. For cotton, the transformation of the histogram after detrending is less clear.

Up to this point we have made no assumptions about the determinants of the price data that we have discussed, apart from prior estimation of an exponential secular trend. Note, however, that the apparent tendency of the price changes in Figure 4 to be negative at high detrended prices is suggestive of the behavior of changes from high price "spikes" in the model of arbitrage of inventories with occasionally binding non-negativity constraints in the tradition of Gustafson (1958), Samuelson (1971), Wright and Williams (1982), Scheinkman and Schechtman (1983), and Deaton and Laroque (1992, 1996). Can inclusion of a weak secular trend in a commodity storage model induce an apparently linear price correspondence that implies that price is unpredictable, in this class of models?

Motivated by this question, next section we present a nonstationary version of such a model, modified to include a latent productivity trend that induces a partially latent trend in price.

III. A nonstationary storage model

We generalize the standard commodity storage model to the case where stochastic production includes a deteministic trend.

There is an endogenous demand for storage generated by profit-maximizing speculators with rational expectations. We further assume that the functional form of the latent consumption demand and the nature of the production process are such that their interaction jointly implies a price attractor following a log-linear trend. This attractor is active as the conditional price expectation when arbitrage is not active.

We use here the model introduced in Bobenrieth et al. (2021). All agents are competitive and have rational expectations. There is a stationary, strictly decreasing inverse consumption demand function F, such that price at period t is given by $P_t = F(C_t)$, where C_t denotes consumption at time t.

For many important commodities, long run global price declines are commonly attributed to the effects of persistently increasing productivity. Accordingly, we specify stochastic production H_t to follow a trend which for simplicity we assume to be exogenous and deterministic. Consider first the case in which storage is not feasible. In this case consumption C_t equals production H_t, which is equal to total available supply Z_t, and price at t is $P_t = F(C_t) = F(H_t)$. If production can be expressed as $H_t = F^{-1}(\lambda^t F(h_t))$, where $\lambda > 0$ and h_t is an *i.i.d.* shock, then $P_t = \lambda^t F(h_t)$.

The introduction of the possibility of storage changes the model dramatically. Suppose now that non-negative storage $Z_t - C_t$ is feasible, that competitive storers are risk neutral, and that the (fixed) interest rate $r > 0$ is the only storage cost. If inventories are positive, then competitive storage arbitrage implies that expected prices appreciate at the interest rate r, independent of the supply trend. That is,

$$E_t P_{t+1} = (1 + r) P_t,$$

where E_t denotes expectation conditional on information at time t. If the percentage change of expected price is less than the interest rate, then there is no incentive for storage, current consumption C_t equals current total available supply Z_t, and price maps from current market supply, that is, $P_t = F(C_t) = F(Z_t)$. Accordingly, price is characterized by the following condition:

$$F(C_t) \;=\; \max\Big[F(Z_t),\; \frac{1}{1+r} E_t F(C_{t+1})\Big], \tag{1}$$

s.t.

$$Z_{t+1} \equiv Z_t - C_t + H_{t+1},\; \forall\, t \in \mathbb{N}. \tag{2}$$

F is assumed to be the derivative of a Hyperbolic Absolute Risk Aversion (HARA) function. That is,

$$\frac{F''F}{(F')^2} = \kappa, \tag{3}$$

where κ is a constant. This set of functions is quite general; it includes linear, log-linear, and iso-elastic consumption demand functions.[4]

The model admits a latent time-invariant representation, with normalized available supply z_t, normalized consumption c_t, and normalized production h_t, given by $z_t \equiv F^{-1}[\lambda^{-t}F(Z_t)]$, $c_t \equiv F^{-1}[\lambda^{-t}F(C_t)]$, and $h_t \equiv F^{-1}[\lambda^{-t}F(H_t)]$, respectively. Equations (1) and (2) have the following latent representation:[5]

[4]For example $F(C) = (A+BC)^{\frac{1}{1-\kappa}}$ if $\kappa \neq 1$, and $F(C) = e^{A+BC}$ if $\kappa = 1$, where A and B are constants.

[5]To rule out bubble models, we assume that $EF(h) < \infty$, where E denotes expectation with respect to the random variable h.

$$F(c_t) \;=\; \max\Big[F(z_t),\ \frac{\lambda}{1+r}E_tF(c_{t+1})\Big], \tag{4}$$

s.t.

$$z_{t+1} \equiv \lambda^{\kappa-1}(z_t - c_t) + h_{t+1}. \tag{5}$$

Denote the "detrended price," which is not independent of the trend, to be normalized latent price $p_t \equiv F(c_t)$. If $\lambda < 1+r$, standard arguments imply that there exists a stationary rational expectations equilibrium (SREE) price function $p : [\underline{h}, \infty) \to \mathbb{R}$ for detrended price:

$$p_t = p(z_t) \;=\; \max\Big[F(z_t),\ \frac{\lambda}{1+r}E_t\ p(z_{t+1})\Big]. \tag{6}$$

Furthermore, p is non-negative, continuous, and strictly decreasing whenever is strictly positive. The following complementary inequalities hold:

$$\begin{aligned} p(z) = F(z), &\quad \text{for} \quad z \le F^{-1}(p^*),\\ p(z) > F(z), &\quad \text{for} \quad z > F^{-1}(p^*), \end{aligned}$$

where $p^* \equiv \dfrac{\lambda}{1+r}Ep(h) \in \mathbb{R}$.

The Euler equation (6) implies the following autoregression for detrended prices:

$$E_tp_{t+1} \;=\; \left(\frac{1+r}{\lambda}\right)\min[p^*,\ p_t]. \tag{7}$$

In terms of observable prices, we have the equivalent autoregression:

$$E_t P_{t+1} \;=\; (1+r)\min[\lambda^t p^*, P_t]. \tag{8}$$

Remark 1. *The trend in the price attractor* $(1+r)\lambda^t p^*$ *is induced by the interaction of the latent production trend and the latent consumer demand of the HARA class. The definition of* c_t *implies that* $F(C_t) = \lambda^t F(c_t)$, *thus price* $P_t = F(C_t)$ *and detrended price* $p_t = F(c_t)$ *are such that* $P_t \;=\; \lambda^t p_t$.

Remark 2. *The trend equation* $P_t = \lambda^t p_t$ *is not a dynamic equation. Indeed,* $P_t \;=\; \lambda^t p_t$ *is only a representation of observed price* P_t *in terms of time* t, *the trend parameter* λ, *and the detrended price* p_t. *Thus the dynamics of observed price* P_t *are a reflection of the dynamics of detrended price* p_t, *which is latent. For periods when arbitrage is active, expected detrended price increases at a rate higher than* r *to compensate for the trend, as evident in equation (7).*

Although the secular trend in production is reflected in the evolution of consumption and price in the long run, the trend is absent in the expected return to stock-holding. If stocks are held, then $\dfrac{E_t P_{t+1} - P_t}{P_t}$, the spread between the conditional expectation $E_t P_{t+1}$ and the current trending price P_t, equals the interest rate r independent of the trend. In the alternate regime, if stocks are not held, then the spread is:

$$\frac{E_t P_{t+1} - P_t}{P_t} = \frac{(1+r)\lambda^t p^* - P_t}{P_t}.$$

The secular trend parameter λ affects expected price changes only when inventories are zero, and only through its effect on the location of the trending attractor relative to current price.[6]

[6]The value of p^* itself depends upon the values of the paramaters r and λ.

Thus the intuitive inference that annual trend movements in production are roughly proportional to annual trend movements in observed price, year by year,, we is wrong for storable commodities. When discretionary stocks are held, the trend is not only dominated by price variability, as observed by Cashin and McDermott (2002), but indeed absent from the one-period forward price spread.

Price P_t is a mapping from calendar time t and normalized price, $P_t = \lambda^t p_t$. In turn, the dynamic behavior of the latent state, represented by the counterpart latent price p_t, is nonlinear, characterized by autoregression (7). The absence of a one-to-one correspondence between the observed price P_t and the counterpart latent price p_t implies that the price autoregression hides the nonlinear dynamics of the state that are key to understanding the economics of commodity prices. Indeed, given the negative trend in prices, any given observed price level at which there is positive storage is a price level at which there stocks are zero at a sufficiently later time period.

Under appropriate conditions, the following theorem establishes the ergodicity of the latent model. A detailed description of such conditions and the proof of the theorem are in Bobenrieth et al. (2021, Theorem 1 and Theorem 2).

Theorem 1. *The Markov process of normalized available supply $\Phi \equiv \{z_t\}_{t \geq 0}$ is uniformly ergodic, that is, it has a unique invariant probability measure ν_∞ which is a global attractor, and there exists constants $k > 1$ and $R < \infty$ such that for all initial $z_0 \in D$, we have:*

$$||\nu_t - \nu_\infty|| \leq Rk^{-t},$$

where $|| \cdot ||$ denotes the total variation norm, and ν_t is the distribution on z_t conditional on z_0.

IV. Trends and predictability of price changes

To explore the effects of a negative price trend, we conduct Monte Carlo experiments with a heuristic model with the specification introduced in Wright and Williams (1982), and re-used in Williams and Wright (1991, pp. 58-62), Deaton and Laroque (1992, p. 11), Cafiero et al. (2011, pp. 45-46), and Bobenrieth et al. (2021). Inverse consumption demand is $F(C) = 600 - 5C$, harvest has a normal distribution with mean 100 and standard deviation 10, and the only cost of storage is the interest rate, $r = 0.05$.

We alter this heuristic model only by specifying a price trend of minus two percent per year (of the order of magnitude of the trends we estimate for cotton and maize), and implement a numerical simulation of the model described in the previous section, taking account of the effect of the trend on arbitrage incentives. We truncate the normal distribution at plus and minus 5 standard deviations from the mean, as in Deaton and Laroque (1992) and Cafiero et al. (2011), and numerically generate a long sequence of realized prices. From this sequence of prices we then take 75 samples of $T = 72$ prices (the same sample size of our illustrative samples of cotton and maize), starting at $t = 1,\ t = 72$, etc., with each initial value normalized to its detrended level.

Figure 5 shows forward proportional differences, $\dfrac{P_{t+1} - P_t}{P_t}$, mapped from trending prices (top panel) and detrended price (lower panel).[7] The linearity and the lack of dependence on current price suggested by the dark patch in the top panel of Figure 5 are strongly reminiscent of the correspondences in Figures 2 and 3. They imply a linear upsloping autoregression correspondence like Figure 2, and are consistent with empirical claims that commodity prices

[7] We discard the forward proportional difference from the maximum trending price in the sample since it is negative by construction. Remember that Figure 5 reflects 75 samples each with the 72 realizations in Figures 3 and 4. So trending prices above 125 and detrended prices above 175 would have low probability of appearing in one random sample of 72

are highly persistent even during price booms.

The vertical axis in the lower panel in Figure 5 shows the same one-period proportional differences of trending prices shown in the vertical axis in the top panel, but now mapped from detrended prices in the horizontal axis. The dark patch of observations indicates that most of the mass in the invariant distribution is centered on approximately the midpoint of the support of detrended prices at which arbitrage is active, to the left of the detrended threshold p^*. We can use Figure 5, and what we know about the underlying dynamic stochastic model, to refine inferences from observed prices and returns, as presented for example in Figures 3 and 4 above.

The downslope of the dark patch in the lower panel of Figure 5 might suggest that mean conditional returns are in general decreasing in detrended base price, an inference that might well be reinforced by inspection of both panels of Figure 4. However this inference is incorrect for base prices below the threshold detrended price p^*. In the lower panel of Figure 5, at every detrended price to the left of the threshold price p^*, stocks are positive, and the expected return to storage with trending prices is constant at the interest rate, 0.05. At the lowest detrended price in the sample, price changes in this heuristic model are roughly symmetric around the positive expected change. At higher detrended prices, the distribution of these changes in prices becomes increasingly skewed as detrended price rises toward the threshold p^* and the stocks available to cushion the effects of a bad harvest decline. To the left of the threshold detrended price p^*, the concentrated mass in the lower panel of Figure 5 must trend down to maintain the constant expected price change, given this skewness.

The light vertical flare above the horizontal line and to the left of the vertical dashed line p^* indicates the prevalence of price jumps from prices close to, but below the stockout threshold. Given t, for price realizations higher than the stockout threshold (to the right of the vertical dashed line p^*), the expectation of the subsequent price is $(1+r)\lambda^t p^*$, independent of current price P_t. Conditional on this expectation, the expected price drop increases one-for-one with current price P_t. That is why the correspondence exhibits a downward slope to the right of the threshold in the lower panel of Figure 5.

Figure 5 here

The two panels of Figure 5 comprise a dramatic illustration of the empirical mischief that can be made by a neglected modest trend in price, even in a relatively small sample. The trend transforms the complex nonlinear relation between detrended price and the distribution of immediately subsequent return realizations, illustrated by the correspondences in the lower panel of Figure 5 into a picture that suggests an approximately constant distribution of return realizations conditional on any trending price, shown in the top panel of Figure 5. In this heuristic model, high detrended prices are likely to be fleeting. Comparison of the two right-hand quadrants in the lower panel of Figure 5 confirms that, in the stockout region, prices are relatively unlikely to rise. The difference between the top and lower panels in Figure 5 is reflected in the distinct histograms for trending price and detrended price in the panels of Figure 5.

In the context of intertemporal accumulation, detrending reorders the base prices. The correspondence between forward returns and the detrended base prices reveals the nonlinear relation implied by constrained storage arbitrage. Neglect of a small trend - or more generally, neglect of any similarly persistent influence on price - can hide the fact that when the trending price is above the trending threshold $\lambda^t p^*$, the spread reflects an expected price jump equal to the difference between current price and $(1+r)$ times that threshold; below this threshold, the spread equals the expected return on storage, r. Since trending price P_t is a strictly increasing function of detrended price p_t, and a strictly decreasing function of time ($P_t = \lambda^t p_t$), the same value of trending price can map to quite different distributions of returns.

Our numerical results make a case for developing an econometric procedure that does not impose the assumption that the trend is exactly zero. However, without this assumption the question of consistency and accuracy of the estimators naturally arises. Following Bobenrieth, Bobenrieth and Wright (2022), in the next section we present asymptotic theory of a remarkably simple one-step least squares estimator for the trend parameter λ, the detrended threshold p^*, and the interest rate r.

V. Estimation of key parameters

Our empirical model is based on the threshold nonlinear price autoregression (8). The presence of an exponential trend in the threshold affects the regression predictor; our empirical model violates key assumptions of continuous threshold models (see Hansen 2017 for a survey), including continuity of the regression in the limit. Bobenrieth, Bobenrieth and Wright (2022) provide asymptotic foundations for hypothesis testing using one-step least squares estimators of nonlinear nonstationary stochastic models. Given data on prices only, in this section we present the results of strong consistency and asymptotic normality for one-step nonlinear least squares estimation of three key parameters of the model: the trend parameter λ, the detrended threshold p^*, and the interest rate r.[8]

Our estimation procedure is robust to the latent specification of the consumption demand in the HARA set and to the distribution of net supply; they are drawn from an infinite

[8]The proofs of our asymptotic results are in Bobenrieth, Bobenrieth, and Wright (2022), available upon request from the authors.

combination of couples (F, H_t) that are consistent with the exponential trend in prices.

For the consistency result, it is assumed that the distribution of detrended output h_t is absolutely continuous with strictly positive derivative on the interior of its support, which is assumed to be compact. We assume that the invariant distribution for the detrended price process p_t has support $[\underline{p}, \overline{p}]$, with $0 < \underline{p} < p^* < \overline{p} < \infty$.

Equation (8) implies:

$$P_{t+1} \;=\; (1+r)\min\Big\{\lambda^t p^*,\; P_t\Big\} \;+\; e_{t+1}, \text{ where } E_t(e_{t+1}) = 0 \tag{9}$$

For clarity of exposition, in the remainder of the paper we write a subscript "0" to denote the true parameter values. $\theta_0 = (\lambda_0, p_0^*, \gamma_0)$, where $\gamma_0 \equiv 1 + r_0$. We assume that the parameter space Θ is compact.

Our objective is to estimate θ_0 using least squares. To avoid an identification problem, we divide the regression model (9) by P_t:

$$\frac{P_{t+1}}{P_t} \;=\; \gamma_0 \min\left\{\frac{\lambda_0^t p_0^*}{P_t},\; 1\right\} \;+\; \epsilon_{t+1} \tag{10}$$

where $E_t(\epsilon_{t+1}) = 0$.

To simplify the notation we use $f_t(\theta)$ to denote the predictor, that is,

$$f_t(\theta) \equiv \;\gamma \min\left\{\frac{\lambda^t p^*}{P_t},\; 1\right\} \;=\; \gamma \min\left\{\left(\frac{\lambda}{\lambda_0}\right)^t \frac{p^*}{p_t},\; 1\right\}.$$

Therefore, the normalized regression model is:

$$Y_{t+1} \;=\; f_t(\theta_0) \;+\; \epsilon_{t+1}, \quad \text{where} \quad Y_{t+1} = \frac{P_{t+1}}{P_t} \tag{11}$$

Given the presence of the term $(\lambda/\lambda_0)^t$, the objective function in the least squares minimization does not converge uniformly in the parameter space; it does not satisfy the uniform convergence condition of Jennrich (1969). The regression model (11) does not satisfy the Lipschitz condition for consistency in Wu (1981), nor the Lipschitz condition of Andrews (1987), the continuity-type smoothness conditions of Pötscher and Prucha (1989, 1994), nor the Lipschitz L_1-conditions of Skouras (2000) or Lai (1994). The errors $\{\epsilon_{t+1}\}_{t\in\mathbb{N}}$ are not independent, unlike the cases of Wu (1981), van de Geer (1990), and Pollard and Radchenko (2006).

Given $\mu \neq \theta_0$, and a ball $B(\mu)$ centered at μ which does not contain θ_0, let $A_T \equiv \inf_{\theta\in B(\mu)} \sum_{t=1}^{T} \{f_t(\theta) - f_t(\theta_0)\}^2$. In order to be able to identify θ_0, it is necessary that $A_T \to \infty$, as $T \to \infty$.[9] For regression (11) it is possible to prove that there exists $N \in \mathbb{N}$ and a strictly finite positive constant b such that $A_T \geq b\,T$, for $T \geq N$, thus identifying θ_0.

Define $\hat{\theta}_T$ to be the least squares estimator of θ_0. Using the facts that $\{\epsilon_{t+1} f_t(\theta)\}_{t\in\mathbb{N}}$ is a uniformly bounded martingale difference sequence, and that the gradient of the predictor $\nabla f_t(\theta)$ has a polynomial order, Bobenrieth, Bobenrieth, and Wright (2022) establish the convergence of $\sup_{\theta\in B(\mu)} \frac{1}{T}\left|\sum_{t=1}^{T} \epsilon_{t+1}\left(f_t(\theta) - f_t(\theta_0)\right)\right|$ to zero, as $T \to \infty$, almost surely. Based on the classical approach of Wald (1949), this convergence implies that $\lim_{T\to\infty} ||\hat{\theta}_T - \theta_0|| = 0$, almost surely.

Furthermore, Bobenrieth, Bobenrieth, and Wright (2022) prove that

$$\left\{T^{3/2}(\widehat{\lambda}_T - \lambda_0), T^{1/2}(\widehat{p^*}_T - p_0^*), T^{1/2}(\widehat{\gamma}_T - \gamma_0)\right\}_{T\in\mathbb{N}}$$

[9] See Wu (1981, p. 502).

converges in distribution to a normal random vector with mean zero and covariance matrix $\Sigma^{-1}\Lambda\Sigma^{-1}$, where

$$\Lambda \equiv 2\begin{pmatrix} \frac{2Ap_0^{*2}\gamma_0^2}{3\lambda_0^2} & \frac{Ap_0^*\gamma_0^2}{\lambda_0} & \frac{Ap_0^{*2}\gamma_0}{\lambda_0} \\ \frac{Ap_0^*\gamma_0^2}{\lambda_0} & 2A\gamma_0^2 & 2Ap_0^*\gamma_0 \\ \frac{Ap_0^{*2}\gamma_0}{\lambda_0} & 2Ap_0^*\gamma_0 & 2(B+Ap_0^{*2}) \end{pmatrix}, \qquad \Sigma \equiv \begin{pmatrix} \frac{2Cp_0^{*2}\gamma_0^2}{3\lambda_0^2} & \frac{Cp_0^*\gamma_0^2}{\lambda_0} & \frac{Cp_0^{*2}\gamma_0}{\lambda_0} \\ \frac{Cp_0^*\gamma_0^2}{\lambda_0} & 2C\gamma_0^2 & 2Cp_0^*\gamma_0 \\ \frac{Cp_0^{*2}\gamma_0}{\lambda_0} & 2Cp_0^*\gamma_0 & 2D \end{pmatrix},$$

where $A \equiv \lim\limits_{t\to\infty} E\left(\dfrac{\epsilon_{t+1}}{p_t}\mathbb{1}_{\{p_t>p_0^*\}}\right)^2$, $B \equiv \lim\limits_{t\to\infty} E\left(\epsilon_{t+1}\mathbb{1}_{\{p_t\leq p_0^*\}}\right)^2$,
$C \equiv \lim_{t\to\infty} E\left(\frac{1}{p_t}\mathbb{1}_{\{p_t>p_0^*\}}\right)^2$, and $D \equiv \lim\limits_{t\to\infty} E\left(\min\left\{\dfrac{p_0^*}{p_t},\ 1\right\}\right)^2$.

VI. Data

Our two illustrative price series are samples of annual cotton and maize prices. We choose cotton because it is one of the industrial commodities included in Cashin and McDermott (2002), and is the commodity referenced by Deaton and Laroque (1996 pp. 913 and 914) as exemplifying the high persistence of price at all price levels in the sample interval 1900-1987. We include maize because it is a major traded agricultural commodity with clearly trending price and yield, and plausibly little interaction with the global cotton market.

The sample interval (1936-2007) excludes the major effects of the persistent and unprecedented boost in maize demand related to United States mandates and subsidies for biofuels, with implementation starting around 2007, and of recently fast-rising Chinese imports of maize and soybeans.

The nominal prices are from Pfaffenzeller, Newbold, and Rayner (2007) for 1936-2003, extended to 2004-2007 using Pfaffenzeller (2013).[10] We take their annual calendar averages, and then normalize them by the 1977-79 average, following the description given in Pfaffenzeller, Newbold, and Rayner (2007). The deflator for the series is, in line with many time series papers using commodity prices, the Manufactures Unit Value (MUV) Index. As a trade-weighted index of traded commodities, the MUV is probably more relevant to agricultural producers facing a world market for output and for the material inputs that constitute a high share of production costs, and it avoids the need to explicitly account for exchange rate movements.

VII. Empirical results

This is the first paper to implement the empirical estimator for which we present the asymptotic theory described above in Section IV for one-step estimation. We use the annual price data for cotton and maize discussed above. Implementing our one-step nonlinear least squares procedure, we estimate $(1 + r_0)$, p_0^*, and the trend parameter λ_0, and present the results in Table 1.

[10] Nominal prices correspond to Cotton (Outlook "CotlookA index"), middling 1-3/32 inch, traded in Far East, C/F beginning 2006; previously Northern Europe, c.i.f.; Maize (US), no. 2, yellow, f.o.b. US Gulf ports.

Table 1 here

As is well known (see for example Deaton and Laroque, 1992), the interest rate is difficult to estimate in dynamic stochastic models, and our estimates are no exception, as shown in the high values for the standard errors of the estimates of $(1+r_0)$. However, note the negative relationship between r_0 and p_0^* implied by the theory of storage; an increase in the rate of interest implies a higher storage cost, consistent with a lower level of the threshold parameter p_0^*. The negative covariances of our estimates of $(1+r_0)$ and p_0^* reflect this negative correlation, which in turn explains the lower standard errors for our estimates of the attractor $(1+r_0)p_0^*$, generated using the delta method, relative to the standard errors of the estimates for r_0 and p_0^* (see Table 1).

Bobenrieth et al. (2021) implement a two-step econometric procedure for the same storage model we use here. They detrend observed prices in a preliminary step. In their procedure each data point is fitted to the same trend, regardless of its regime. In contrast, our new one-step estimator discriminates the trend regime for each data point in the single estimation step. The estimated trending price thresholds using preliminary detrending (implying annual price trends of -2.08% and -1.82% for cotton and maize respectively) and using our one-step approach (implying annual price trends -2.63% and -1.64% for cotton and maize respectively) are both shown in Figure 6. For the sample of cotton prices, the detrending method changes the sample splitting for 18% of the prices.

Figure 6 here

By construction, our one-step procedure yields a lower sum of square errors for the forecast of within sample price changes, compared to the two-step estimation of the model which involves preliminary detrending. But does it make a difference for the state, as represented by the detrended price?

In our data exploration in Figure 4 we used preliminary detrending of prices, the same procedure of Bobenrieth et al. (2021). Figure 7 shows the same vertical axis as Figure 4, but now with the detrended price in the horizontal axis estimated using our new one-step approach. Histograms of the corresponding detrended prices are shown in the horizontal axes.

Figure 7 here

We applied a two-sample KS test of the hypothesis that the price series used in Figure 7, detrended by the one-step approach, are from the same distribution as the prices detrended using a trend estimated in a preliminary step, used in Figure 4. For cotton, the KS D statistic is 0.222, significant at the five percent level (p-value 0.048). For maize, the difference is insignificant (D statistic 0.097, p-value 0.868). Accordingly, we now focus on the comparison of the histograms of detrended prices for cotton in Figure 4 and Figure 7. Our results show that our one-step procedure is able to recover the unimodality of the price distributions, which is a standard feature of prices from the stationary storage model.

We use Shannon's (1948) entropy index to compare the histograms of detrended cotton prices with our one-step approach versus preliminary detrending,

$$I \equiv -\sum_{j=1}^{n} p(j)[\ln(p(j))],$$

where $p(j)$ denotes the frequency of data in bin j. This index I measures the degree to which the histogram is concentrated or dispersed; it is zero if all of the probability mass is concentrated at a single bin. The more unpredictable the detrended price in the histogram, the greater the entropy index.

The I index for one-step detrending and preliminary detrending are 2.5969 and 2.7624 respectively.[11] The values for the I indices suggest that our one-step estimation provides more informative detrended price histograms than preliminary detrending.

Our focus in this paper is on the estimation of expected relative price changes, $\dfrac{E_t P_{t+1} - P_t}{P_t}$. Equation (8) implies that the expected relative price change, the spread, is constant to the left of the threshold detrended price p_0^* but it slopes down to the right of the estimate of p_0^* as suggested in the correspondence generated in the heuristic model, illustrated in Figure 5 above:

$$\frac{E_t P_{t+1} - P_t}{P_t} = \begin{cases} r_0 & \text{if } p_t \leq p_0^* \\ \frac{(1+r_0)p_0^*}{p_t} - 1 & \text{if } p_t > p_0^*. \end{cases}$$

Both panels in Figure 8 reproduce on the vertical axis the first differences of trending price of cotton and maize seen in Figures 3, 4 and 6, and price on the horizontal axis is detrended using our one-step estimates, as in Figure 6. But in Figure 8 we also plot our estimate of the function $\upsilon(p_t) \equiv \dfrac{(1+r_0)p_0^*}{p_t} - 1$, evaluated using our estimates for r_0 and p_0^*. For any level of detrended price p_t, above or below p_0^*, this function $\upsilon(p_t)$ is an upper bound for the spread $\dfrac{E_t P_{t+1} - P_t}{P_t}$. The dashed lines show confidence regions of plus and minus two standard deviations for such function.

Figure 8 illustrates the nonlinearity of our predictions for price changes for cotton and maize. Indeed, each of the detrended base prices to the right of the vertical blue line, which is our estimate for the attractor in the latent stationary model, $(1 + r_0)p_0^*$, is in the second regime which has no stocks, and our estimate of the expected detrended price equals the

[11] To select the number of bins in the histogram we use the standard square-root choice.

location of the blue attractor on the horizontal axis. Every base price P_t changes in expectation by the lesser of rP_t, or the distance to the attractor, depending on the regime.

Figure 8 here

Figure 9 shows the time series for cotton and maize prices along with the corresponding conditional expected prices. The time series of prices plotted in Figure 9 are good examples of the justification of Deaton and Laroque (2003) for neglecting long run trends. Indeed the annual average price movement due to the trend for each illustrative sample appears generally negligible relative to annual price changes, and the series show intervals of sustained price rises often ending in price spikes, interspersed with brief intervals including large and heterogeneous price drops.

Figure 9 here

Paradoxically, our predicted prices - in green for the upward stochastic trend and red for the downward jumps in Figure 9 - show that a small negative trend is of no relevance for expected price changes in periods when prices are locally low. Such a trend is revealed in realized returns only indirectly and infrequently, accumulated in infrequent jumps down from boom prices. A notable feature of our predicted prices in both panels of Figure 9 is that precisely those large heterogeneous downward jumps are the price variations which are quite accurately predicted (as measured by the conditional coefficient of variation).

Our estimates, illustrated in Figure 6, Figure 7, Figure 8, and Figure 9, show that price persistence is very different depending on whether or not stocks are positive.[12] For current price above the threshold price $\lambda_0^t p_0^*$, the conditional expectation of price is, as noted above, independent of current price. The spread between expected price for next period and the current price decreases as the current price increases, dollar for dollar. Given any current price below the threshold $\lambda_0^t p_0^*$, storage arbitrage is active and the spread between the conditional expectation of price next period and current price is constant and equal to the interest rate r, independent of the trend.

The two highest maize prices in the lower panel of Figure 9 are from early years in the sample when such high prices are far above the conditional expectation of the following price, and are in fact followed by large downward jumps, consistent with the model presented here. Those are the two "outliers" in the lower panel of Figure 2 referenced in footnote 1. They are represented in the lower panels of Figure 7 and Figure 8 as the two lowest observations. In this context, they do not stand out as so anomalous.

However, Figure 9 shows how price autoregressions can scramble the regimes. Consider prices in the stockout regime (i.e. with red predicted prices in Figure 9). For many of those prices it is possible to find a very similar price that belongs to the alternate regime, with green predicted prices. Putting them together in a price autoregression like Figure 2 mixes both regimes, thus reducing the predictive power of observed price alone, ignoring the influence of calendar time via a small deterministic trend that scrambles the ordering of base prices relative to their ordering if detrended,

[12]The apparent positive trend in cotton and maize prices for most of the 1940s is consistent with the large levels of both private and public carryovers reported for cotton and maize in the same period. See for example Federal Reserve Bank Richmond, Va (1944, pp. 2-3) and Shepherd and Richards (1957, Figure 6, p. 975).

VIII. Concluding remarks

A large empirical literature on the existence, sign, size and direction of secular trends in the tradition of Prebisch (1950) and Singer (1950) is now beset by conflicting conclusions regarding the number and locations of breaks in the time series. The more numerous the economically exogenous breaks, the less useful are the trend estimates for those interested in price prediction. It is understandable that researchers focused on intermediate range forecasting based on current price have tended to neglect the possible influence of a small annual trend movement. Difficult to estimate and of uncertain persistence, any secular trend appears to add at most a minor constant to highly variable and otherwise unpredictable

A small latent secular trend can transform inferences about behavior of commodity prices. It linearizes highly nonlinear price autoregressions and generates sample returns that appear to be independent of current price. It renders observed price an unreliable indicator of forward returns. Indeed, the fact that price is observed to be equally persistent at all price levels, high and low, does not imply that forward returns are independent of the current state. The latter inference is an illusion generated by the very trends that seem so negligible.

How can small trends have such dramatic effects? The trend might have a negligible effect on short-run returns, but it renders current price a hazardous indicator of those returns. A latent trend in price series implies that an observed price might map to one of many different price expectations.

Consider a price observed early in a given sample (for example, maize price in 1940 in Figure 9), which lies below the corresponding value of the trending threshold. The distribution of the immediately subsequent proportional price change has a mean of r_0 times the 1940 price observation, consistent with active storage arbitrage. Consider, alternately, the case of a lower price observation that occurs at later time when it is above the down-trending threshold (for example, maize price in 1996 in Figure 9). In this case, with no stocks, the distribution of the subsequent price change has a mean that is a downward jump to the attractor indicated by the solid red square at 1997 in Figure 9.

Hence, price changes from similar prices observed sufficiently early and sufficiently late in the sample can have very distinct conditional distributions. For any given negative downward trend,no matter how small, if the sample size is large enough, trend-induced shuffling of observed prices with very different distributions of forward realizations can render the return realization correspondence approximately independent of current price, even though each price realization is a function of the current price, the harvest realization, and time.

In studies that recognize trends, it is a common practice to detrend the sample observations in a preliminary step, and then estimate the model relying on standard consistency proofs for models without trends as if they applied to the detrended model. We implement a new one step consistent econometric procedure in an empirical model that explicitly recognizes the influence of the trend on arbitrage.

In our samples, a Kolmogorov-Smirnov test indicates that the sample distribution of cotton prices detrended using our simultaneously estimated trend parameter is significantly different from the distribution implied by the usual approach using a preliminary estimate of the trend, while the difference is not significant for maize. Only after implementing our simultaneous approach can we know whether prior detrending is sufficient. Hence this approach is important even if a trend has been recognized in existing studies and estimated in a preliminary step. The restriction on interactions with other parameters implied by prior detrending might well affect inferences about the trend and the behavior of detrended endogenous variables, and the predictability of price changes, as observed in our sample of cotton prices.

We do not claim that our empirical estimates capture all the features of our illustrative samples. A more accurate identification of secular trends, breaks in trends, stochastic trends, or persistent ergodic shifters should only increase the force of our findings regarding the effects of ignoring such phenomena in empirical analysis.

Our empirical approach could be useful to meet challenges in estimation of models with trends that are incorporated into the structure of the model, without log-linearization of the Euler equations.[13] Such challenges are encountered in DSGE models with trends in endogenous variables, as discussed in the survey by Fernández-Villaverde, Rubio-Ramírez and Schorfheide (2016, p. 650).Our work shows that simultaneous estimation of trends with other parameters, as opposed to the standard practice of preliminary detrending using an estimated trend, can be crucial for correct inferences about the predictability of price changes and the nature of their distributions.

References

Andrews, D.W.K. 1987. "Consistency in Nonlinear Econometric Models: a Generic Uniform Law of Large Numbers." *Econometrica,* 55(6), 1465-1471.

Bobenrieth, J.R.A., Bobenrieth, E.S.A., and B.D. Wright. 2022. "Estimation of Nonlinear Models with Exponential Trends." Working paper, Universidad del Bío-Bío, Pontificia Universidad Católica de Chile, and University of California at Berkeley.

Bobenrieth, E.S.A., Bobenrieth, J.R.A., Guerra, E.A., Wright, B.D., and D. Zeng. 2021. "Putting the Empirical Commodity Storage Model Back on Track: Crucial Implications of a 'Negligible' Trend." *American Journal of Agricultural Economics,* 103(3), 1034-1057.

[13] Log-linearization of Euler equations is known to raise issues of accuracy, see for example Carroll (2001) and Ludvigson and Paxson (2001).

Cafiero, C., Bobenrieth, E.S.A., Bobenrieth, J.R.A., and B.D. Wright. 2011. "Empirical Relevance of the Competitive Storage Model." *Journal of Econometrics,* 162(1), 44-54.

Carroll, C.D. 2001. "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)." *The B.E. Journal of Macroeconomics,* 1(1), 1-38.

Cashin, P., and C.J. McDermott. 2002. "The Long-Run Behavior of Commodity Prices: Small Trends and Big Variability." *IMF Staff Papers*, Palgrave Macmillan, 49(2), 175-199.

Deaton, A. 2010. "Understanding the Mechanisms of Economic Development." *Journal of Economic Perspectives*, 24(3), 3-16.

Deaton, A., and G. Laroque. 1992. "On the Behaviour of Commodity Prices." *Review of Economic Studies*, 59(1), 1-23.

Deaton, A., and G. Laroque. 1996. "Competitive Storage and Commodity Price Dynamics." *Journal of Political Economy*, 104(5), 896-923.

Deaton, A., and G. Laroque. 2003. "A Model of Commodity Prices After Sir Arthur Lewis." *Journal of Development Economics*, 71(2), 289-310.

Federal Reserve Bank, Richmond, Va. 1944. "Monthly Review of Financial and Business Conditions." Fifth Federal Reserve District.

Fernández-Villaverde, J., Rubio-Ramírez, J.F., and F. Schorfheide. 2016. "Solution and Estimation Methods for DSGE Models." *Handbook of Macroeconomics*, 2A, 527-724.

Gouel, C., and N. Legrand. 2017. "Estimating the Competitive Storage Model with Trending Commodity Prices." *Journal of Applied Econometrics*, 32(4), 744-763.

Guerra, E.A., Bobenrieth, E.S.A., Bobenrieth, J.R.A., and C. Cafiero. 2015. "Empirical Commodity Storage Model: the Challenge of Matching Data and Theory." *European Review of Agricultural Economics*, 42(4), 607-623.

Gustafson, R.L. 1958. "Carryover Levels for Grains." USDA, Technical Bulletin 1178.

United States Department of Agriculture, Washington, DC.

Hansen, B.E. 2017. "Regression Kink With an Unknown Threshold." *Journal of Business & Economic Statistics,* 35(2), 228-240.

Harvey, D.I., Kellard, N.M., Madsen, J.B., and M.E. Wohard. 2010. "The Prebisch-Singer Hypothesis: Four Centuries of evidence." *The Review of Economics and Statistics,* 92(2), 367-377.

Harvey, D.I., and S.J. Leybourne. 2015. "Confidence Sets for the Date of a Break in Level and Trend When the Order of Integration is Unknown." *Journal of Econometrics,* 184(2), 262-279.

Jansen, J.P., Smith, A., and C.A. Carter. 2018. "Commodity Price Comovement and Financial Speculation: The Case of Cotton." *American Journal of Agricultural Economics,* 100(1), 264-285.

Jennrich, R.I. 1969. "Asymptotic Properties of Non-Linear Least Squares Estimators." *The Annals of Mathematical Statistics,* 40, 633-643.

Kim, D., and P. Perron. 2009. "Unit Root Tests Allowing for a Break in the Trend Function at an Unknown Time Under Both the Null and Alternative Hypotheses." *Journal of Econometrics*, 148(1), 1-13.

Lai, T. L. 1994. "Asymptotic Properties of Nonlinear Least Squares Estimates in Stochastic Regression Models." *The Annals of Statistics,* 22(4), 1917-1930.

Ludvigson, S., and C. Paxson. 2001. "Approximation Bias in Linearized Euler Equations." *The Review of Economics and Statistics,*, 83(2), 242-256.

Osborne, T. 2004. "Market News in commodity Price Theory: Application to the Ethiopian Grain Market." *Review of Economic Studies,* 71(1), 133-164.

Perron, P. 1989. "The Great Crash, the Oil Price Shock and the Unit Root Hypothesis." *Econometrica*, 57(6), 1361-1401.

Perron, P., and X. Zhu. 2005. "Structural Breaks with Deterministic and Stochastic Trends." *Journal of Econometrics*, 129(1-2), 65-119.

Pfaffenzeller, S. 2013. "Updating the Grilli and Yang Commodity Price Index - Beyond

2003." Note. http://www.stephan-pfaffenzeller.com/cpi.html

Pfaffenzeller, S., Newbold, P., and A. Rayner. 2007. "A Short Note on Updating the Grilli and Yang Commodity Price Index." *World Bank Economic Review,* 21(1), 151-163.

Pollard, D., and P. Radchenko. 2006. "Nonlinear Least-Squares Estimation." *Journal of Multivariate Analysis,* 97(2), 548-562.

Pötscher, B.M., and I.R. Prucha. 1989. "A Uniform Law of Large Numbers for Dependent and Heterogeneous Data Processes." *Econometrica,* 57(3), 675 - 683.

Pötscher, B.M., and I.R. Prucha. 1994. "On the Formulation of Uniform Laws of Large Numbers: a Truncation Approach." *Statistics,* 25(4), 343-360.

Prebisch, R. 1950. "The Economic Development of Latin America and its Principal Problems." *Economic Bulletin for Latin America,* 7, 1-12.

Roberts, M.J., and W. Schlenker. 2013. "Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate." *American Economic Review*, 103(6), 2265-2295.

Samuelson, P.A. 1971. "Stochastic Speculative Price." *Proceedings of the National Academy of Science*, 68, 335-337.

Scheinkman, J.A., and J. Schechtman. 1983. "A Simple Competitive Model with Production and Storage." *Review of Economic Studies*, 50, 427-441.

Shannon, C.E. 1948. "A note on the concept of entropy." *Bell Syst. Technical Journal*, 27, 379-423.

Shepherd, G., and A. Richards. 1957. "Effects of the USDA Corn Storage Program on Corn Carryover Stocks and Corn Utilization." Research Bulletin 446. Iowa Agricultural Experiment Station Iowa State College. Ames. Iowa.

Singer, H. 1950. "The Distribution of Gains Between Investing and Borrowing Countries." *American Economic Review, Papers and Proceedings*, 40(2), 473-485.

Skouras, K. 2000. "Strong Consistency in Nonlinear Stochastic Regression Models." *The Annals of Statistics,* 28(3), 871-879.

van de Geer, S. 1990. "Estimating a Regression Function." *The Annals of Statistics,* 18(2), 907-924.

Wald, A. 1949. "Note on the Consistency of the Maximum Likelihood Estimate." *The Annals of Mathematical Statistics,* 20(4), 595-601.

Williams, J.C., and B.D. Wright. 1991. *Storage and Commodity Markets.* Cambridge University Press.

Wright, B.D., and J.C. Williams. 1982. "The Economic Role of Commodity Storage." *The Economic Journal*, 92 (367), 596–614.

Wu, C.F. 1981. "Asymptotic Theory of Nonlinear Least Squares Estimation." *The Annals of Statistics,* 9(3), 501-513.

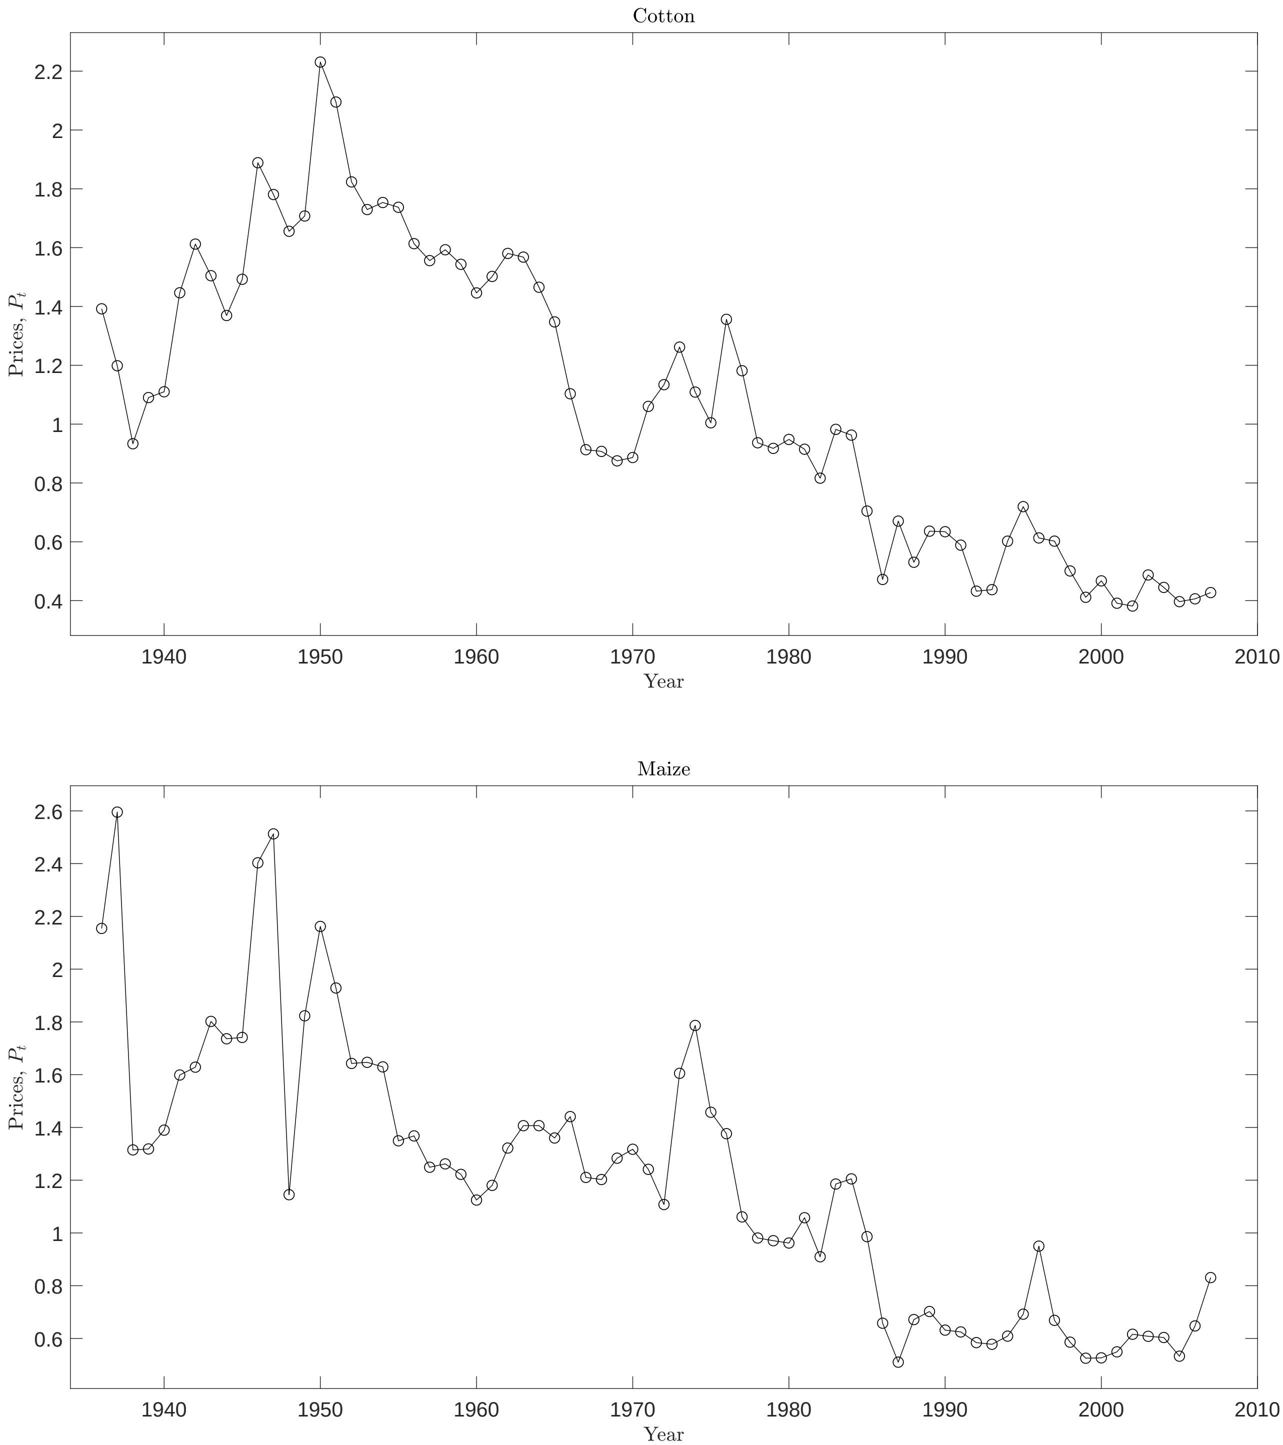

Figure 1: Annual prices for cotton and maize. 1936-2007

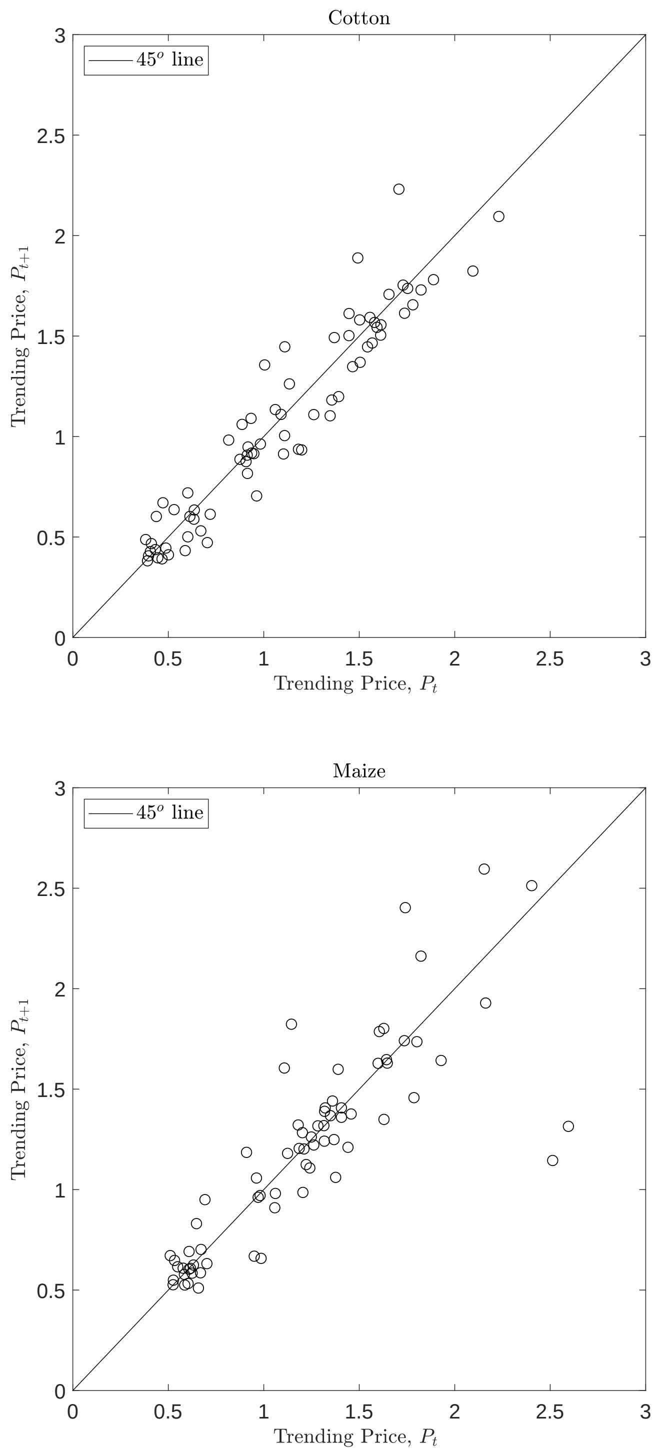

Figure 2: Annual price realizations at t and t+1. 1936-2007

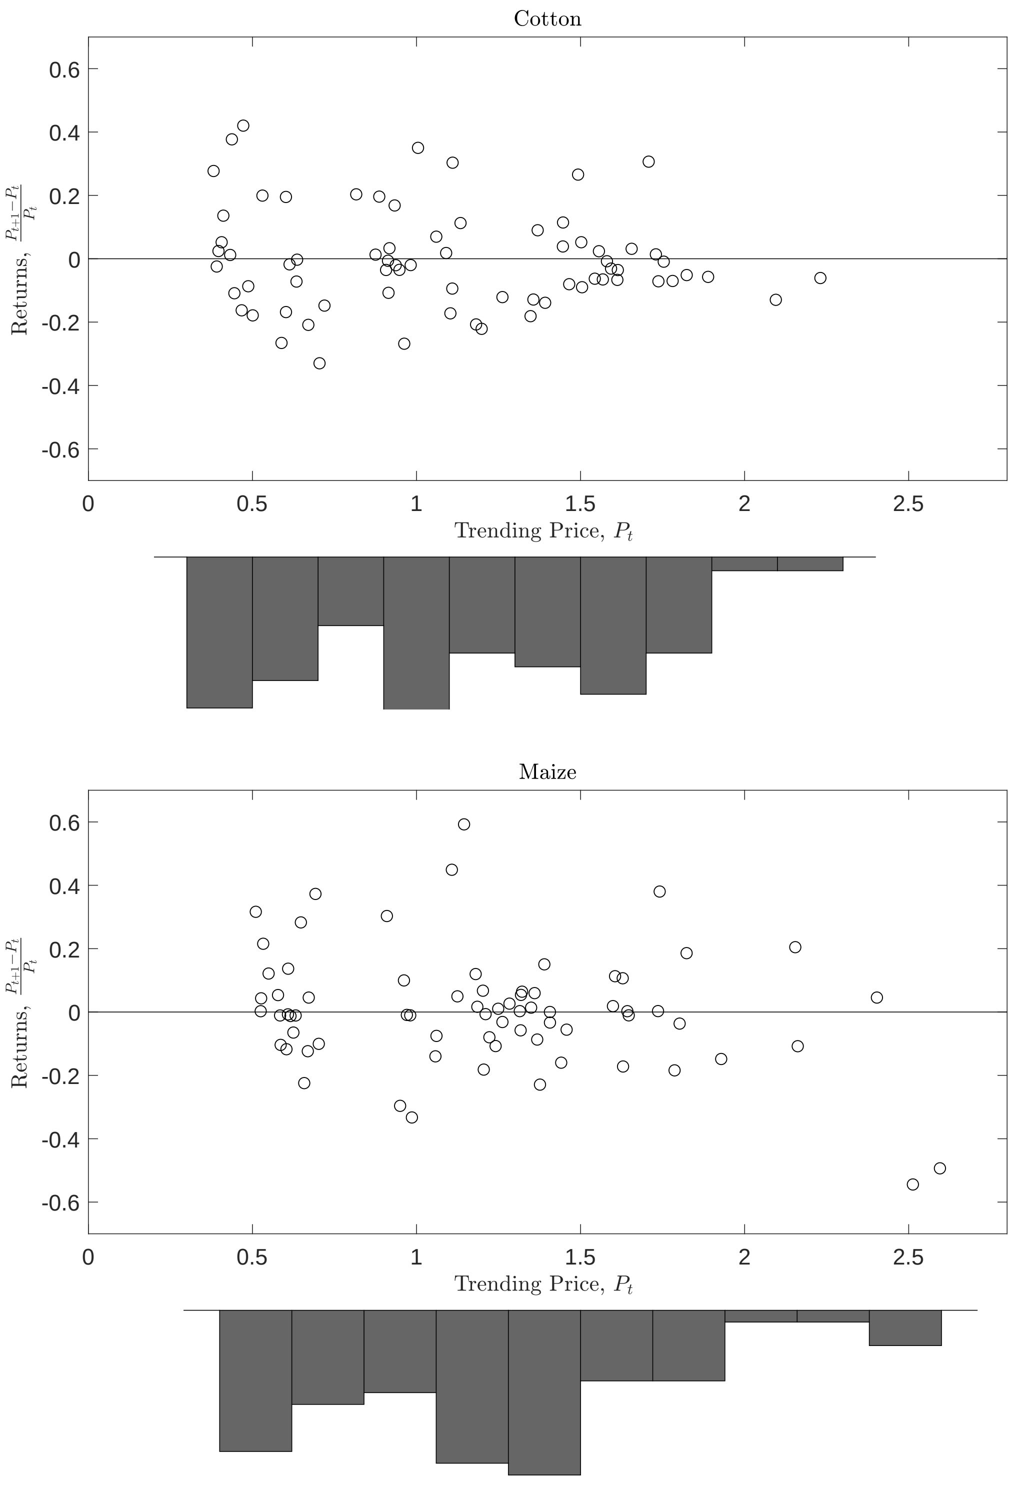

Figure 3: Annual proportional price changes between t and t+1 mapped from price at t. 1936-2007

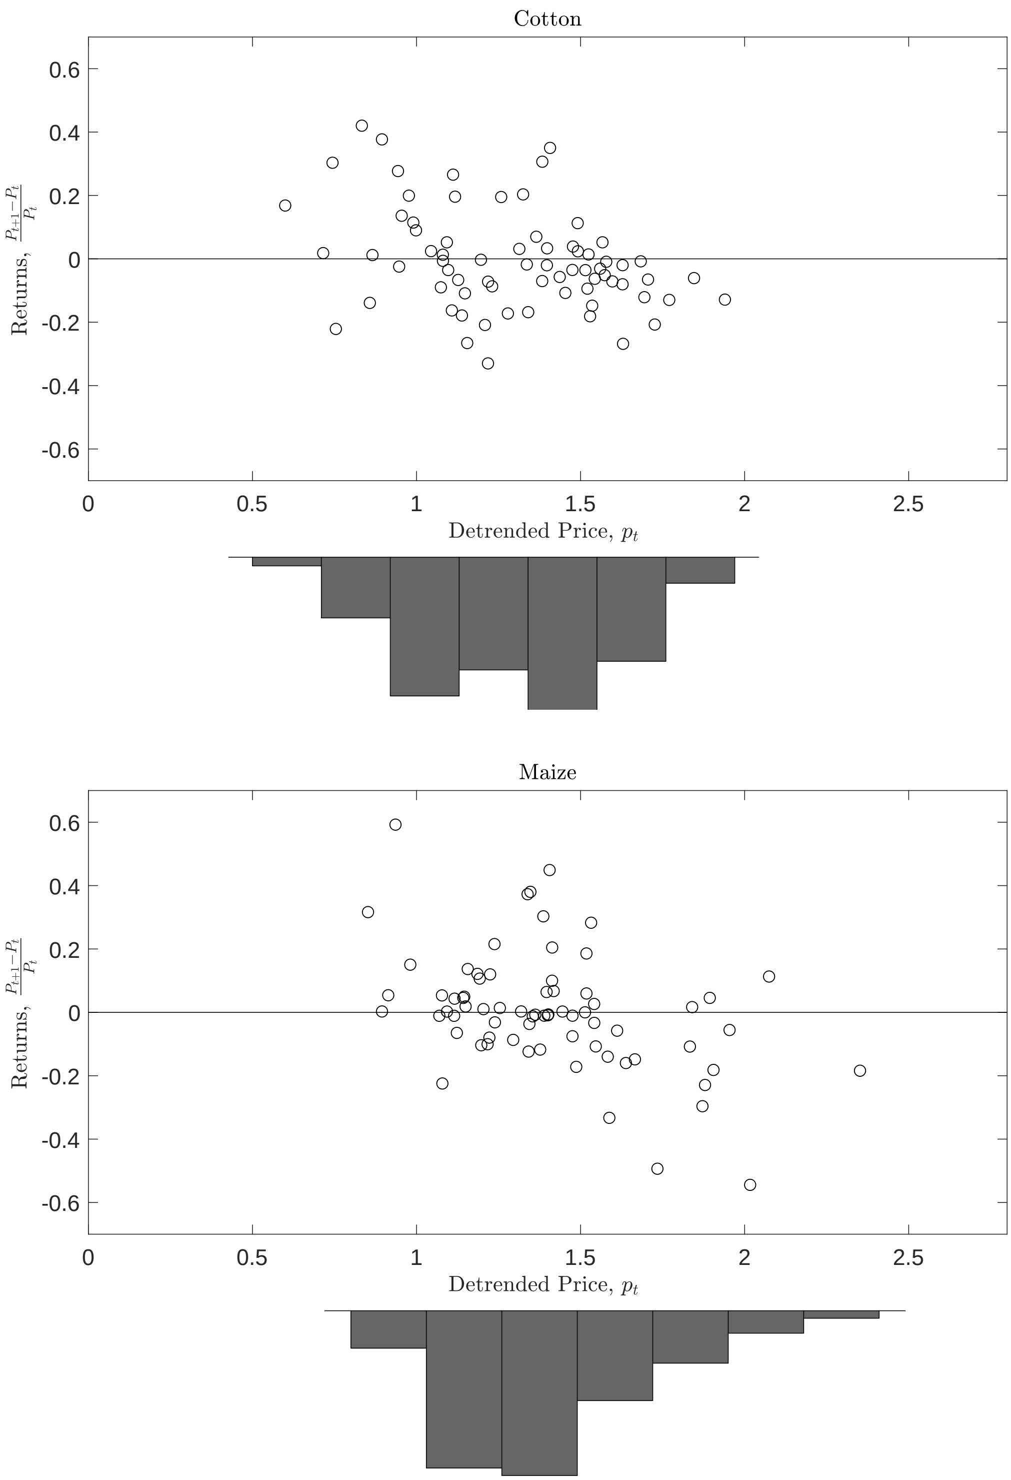

Figure 4: Preliminary Detrending. Annual proportional price changes between t and t+1 mapped from detrended price. 1936-2007

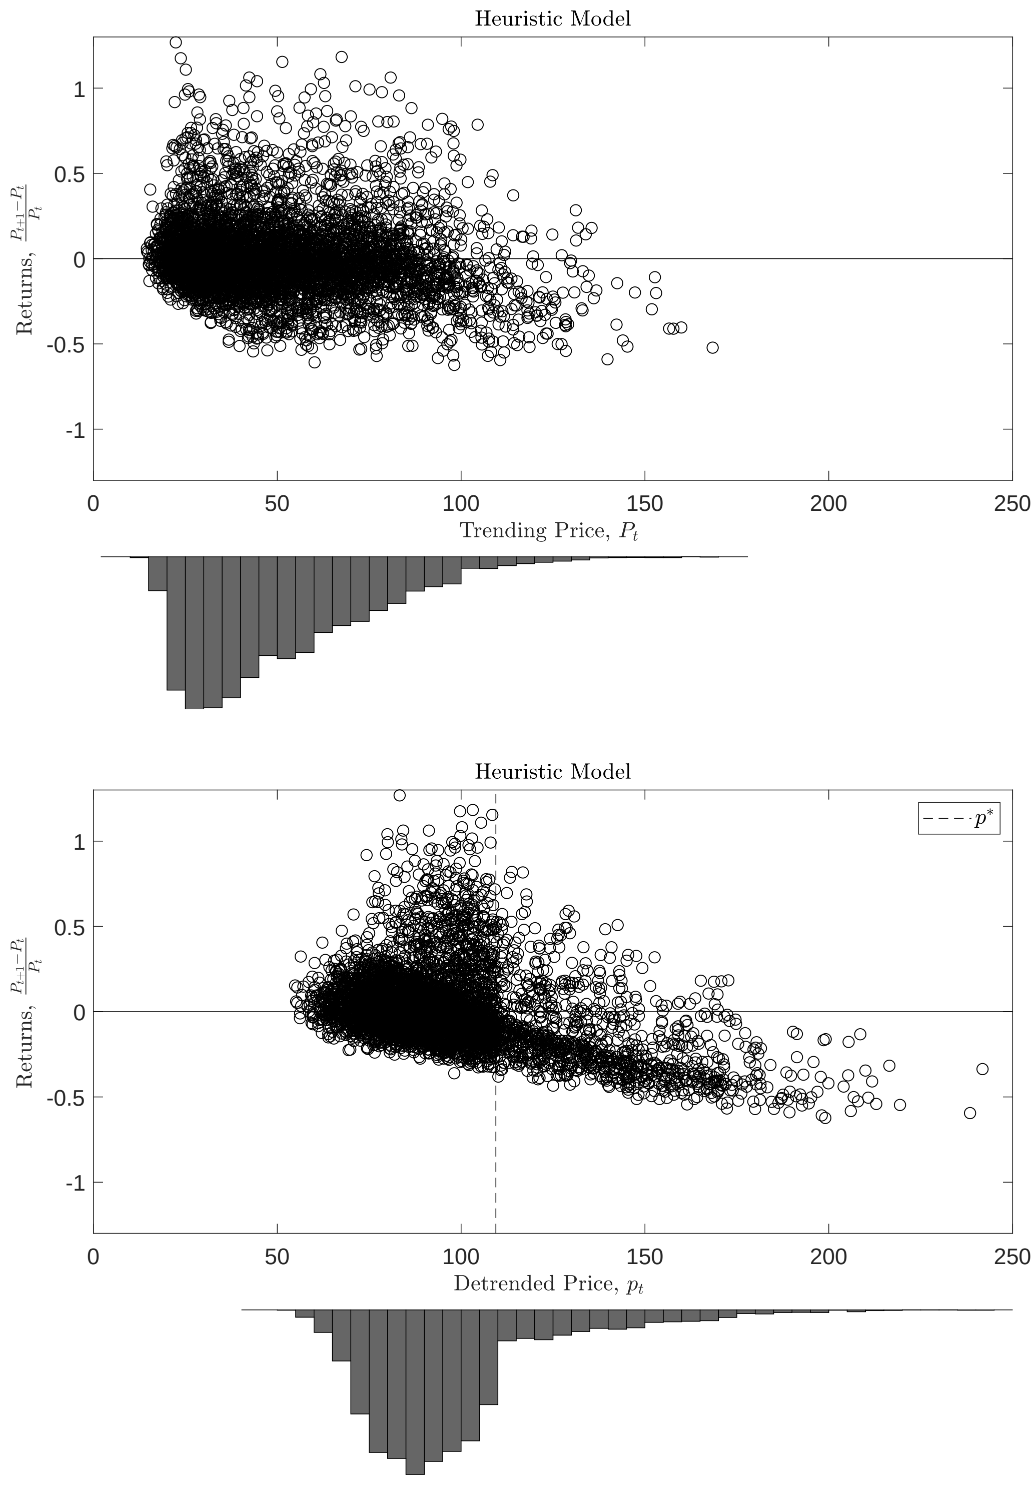

Figure 5: MonteCarlo: proportional price changes

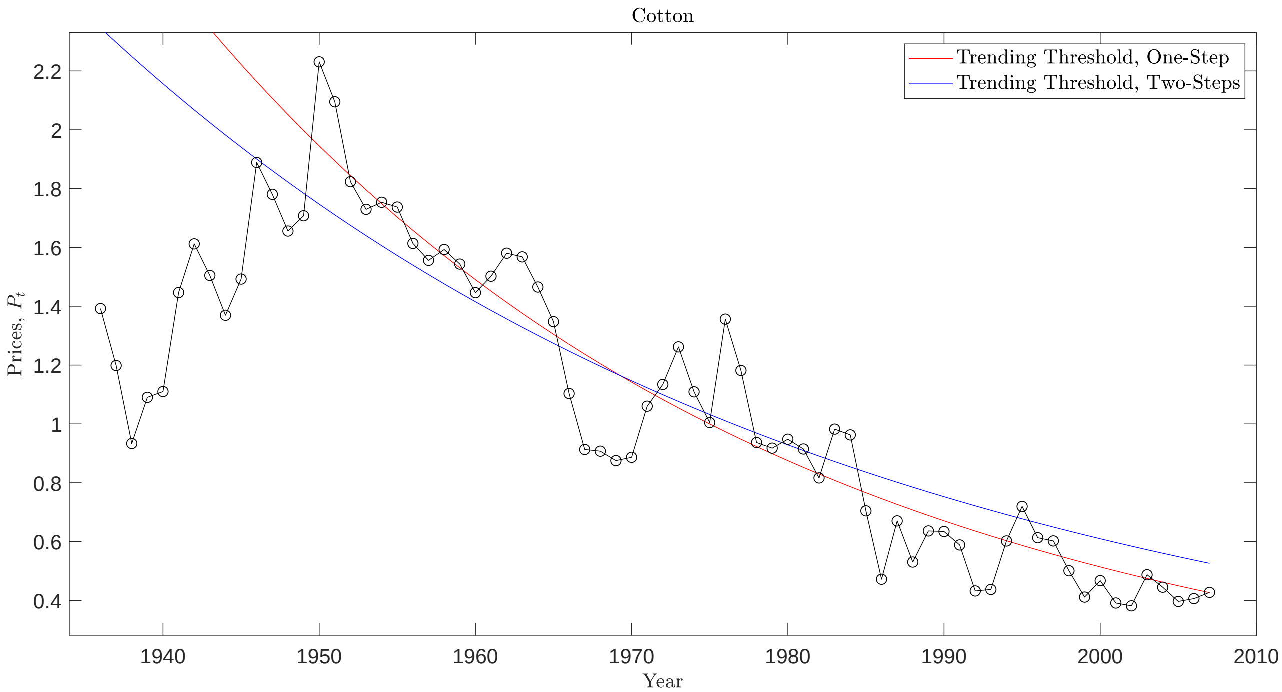

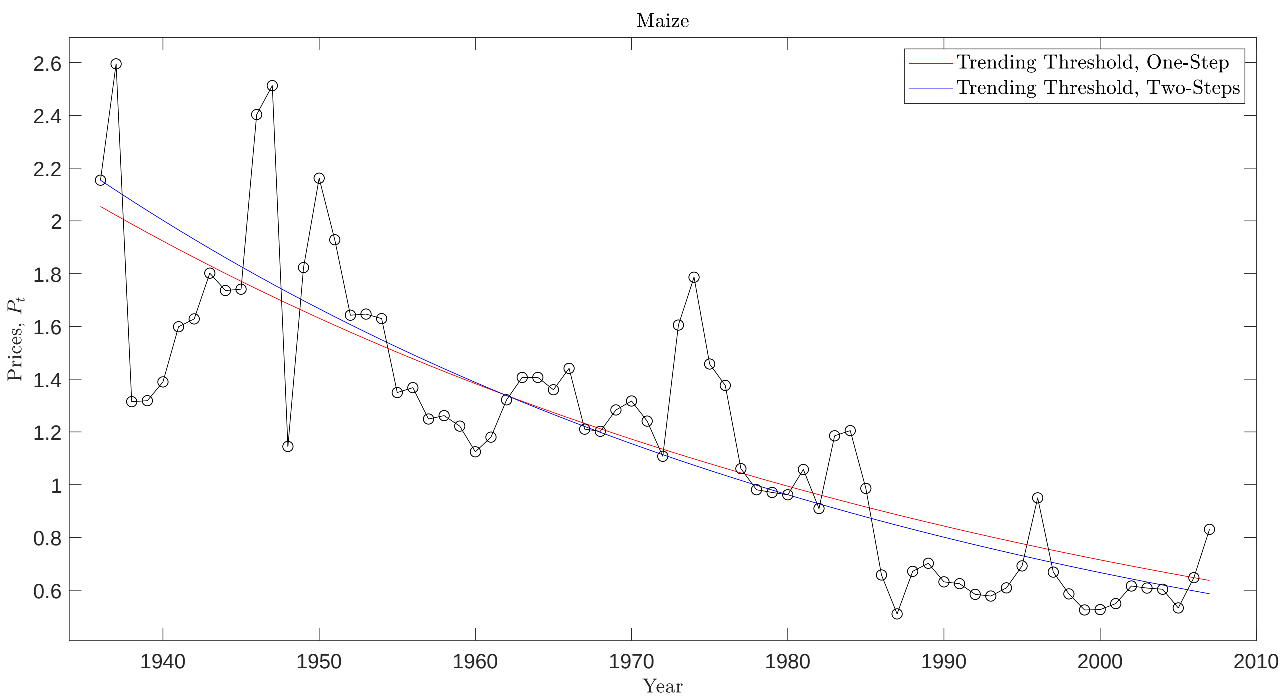

Figure 6: Annual prices and comparison of one-step and two-step estimation of the trending price threshold. 1936-2007

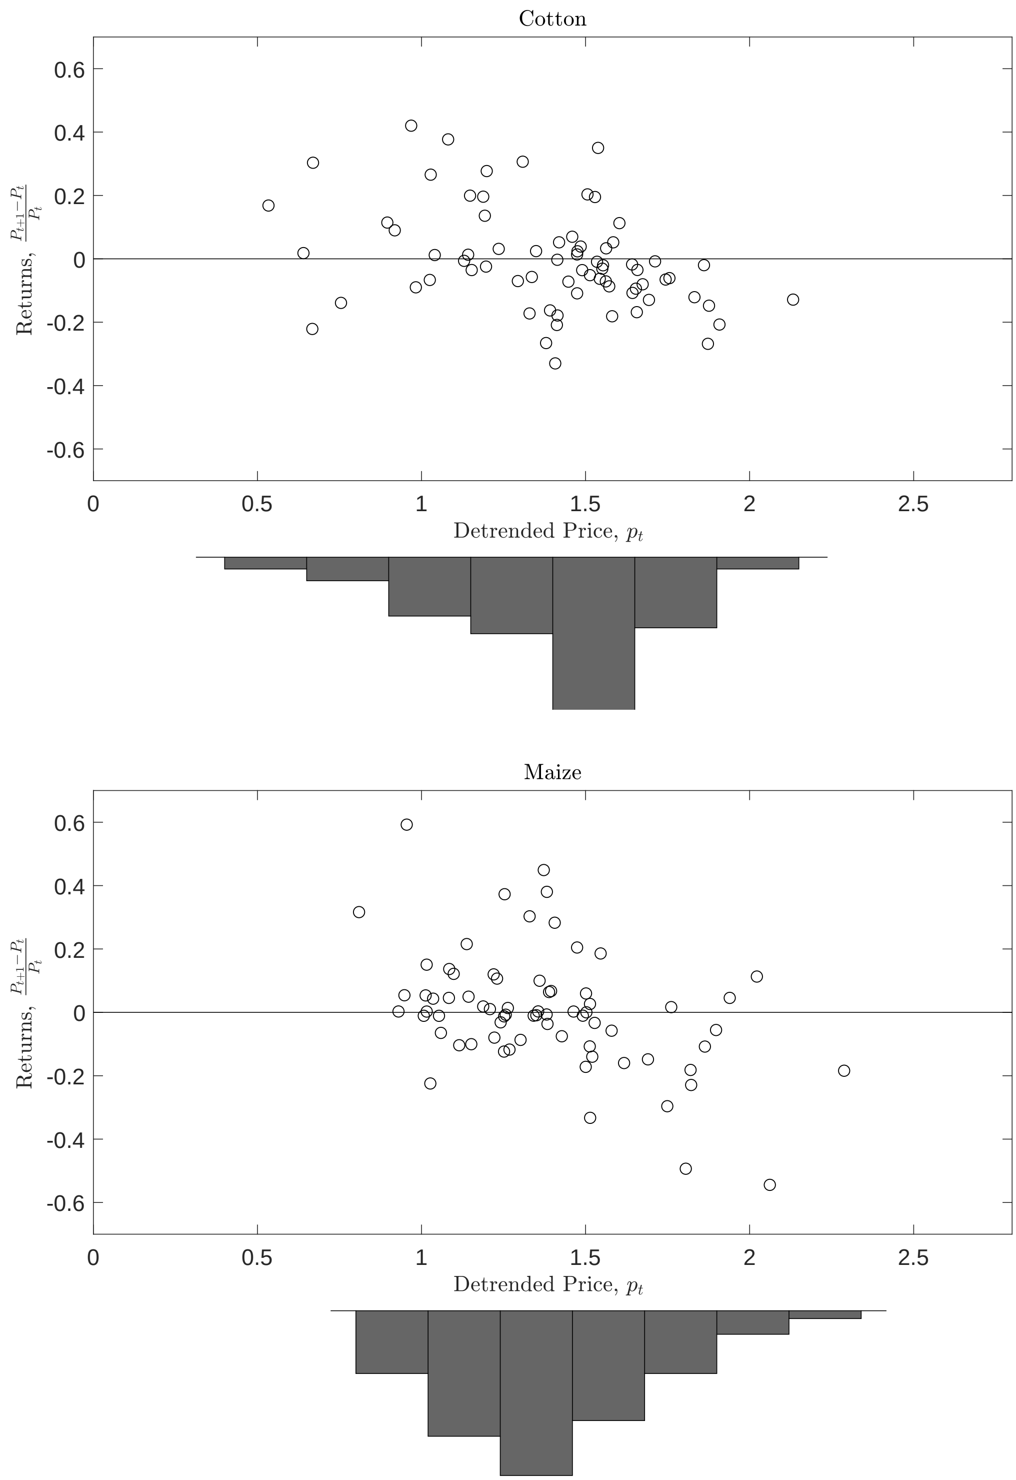

Figure 7: One step Estimates. Annual proportional price changes between t and $t+1$ mapped from detrended price. 1936-2007

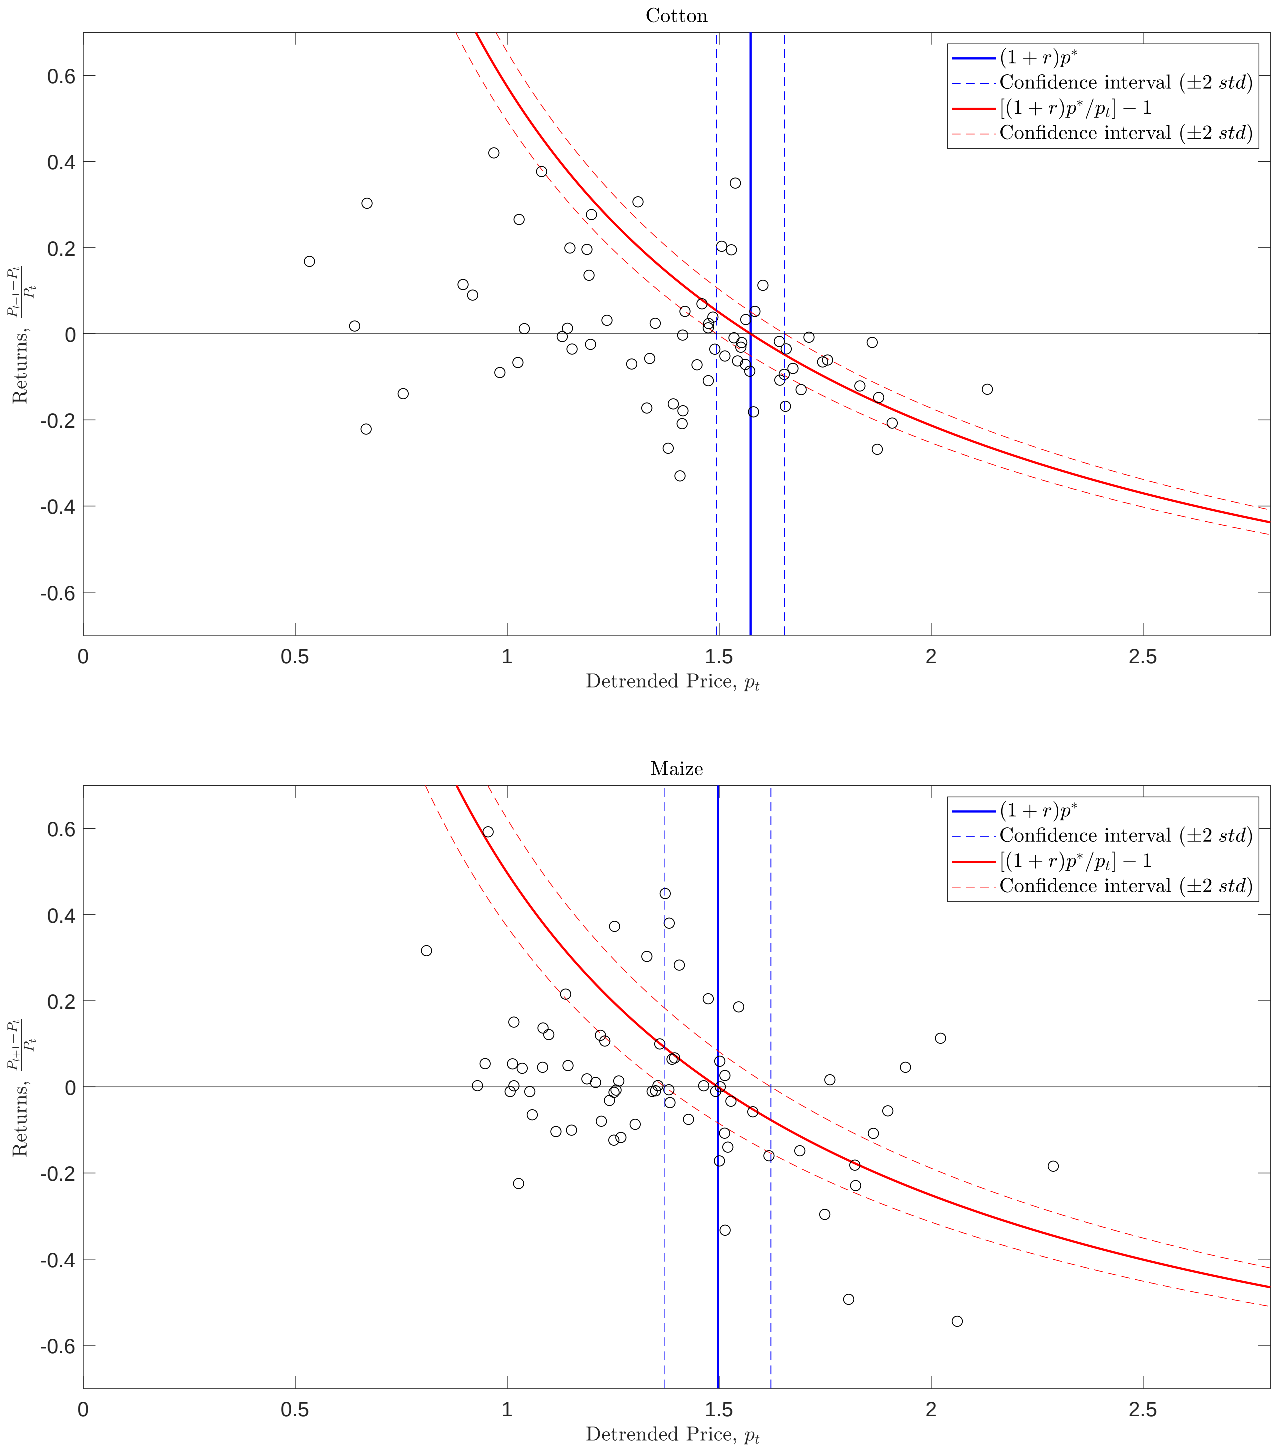

Figure 8: Annual proportional price changes and model predictions

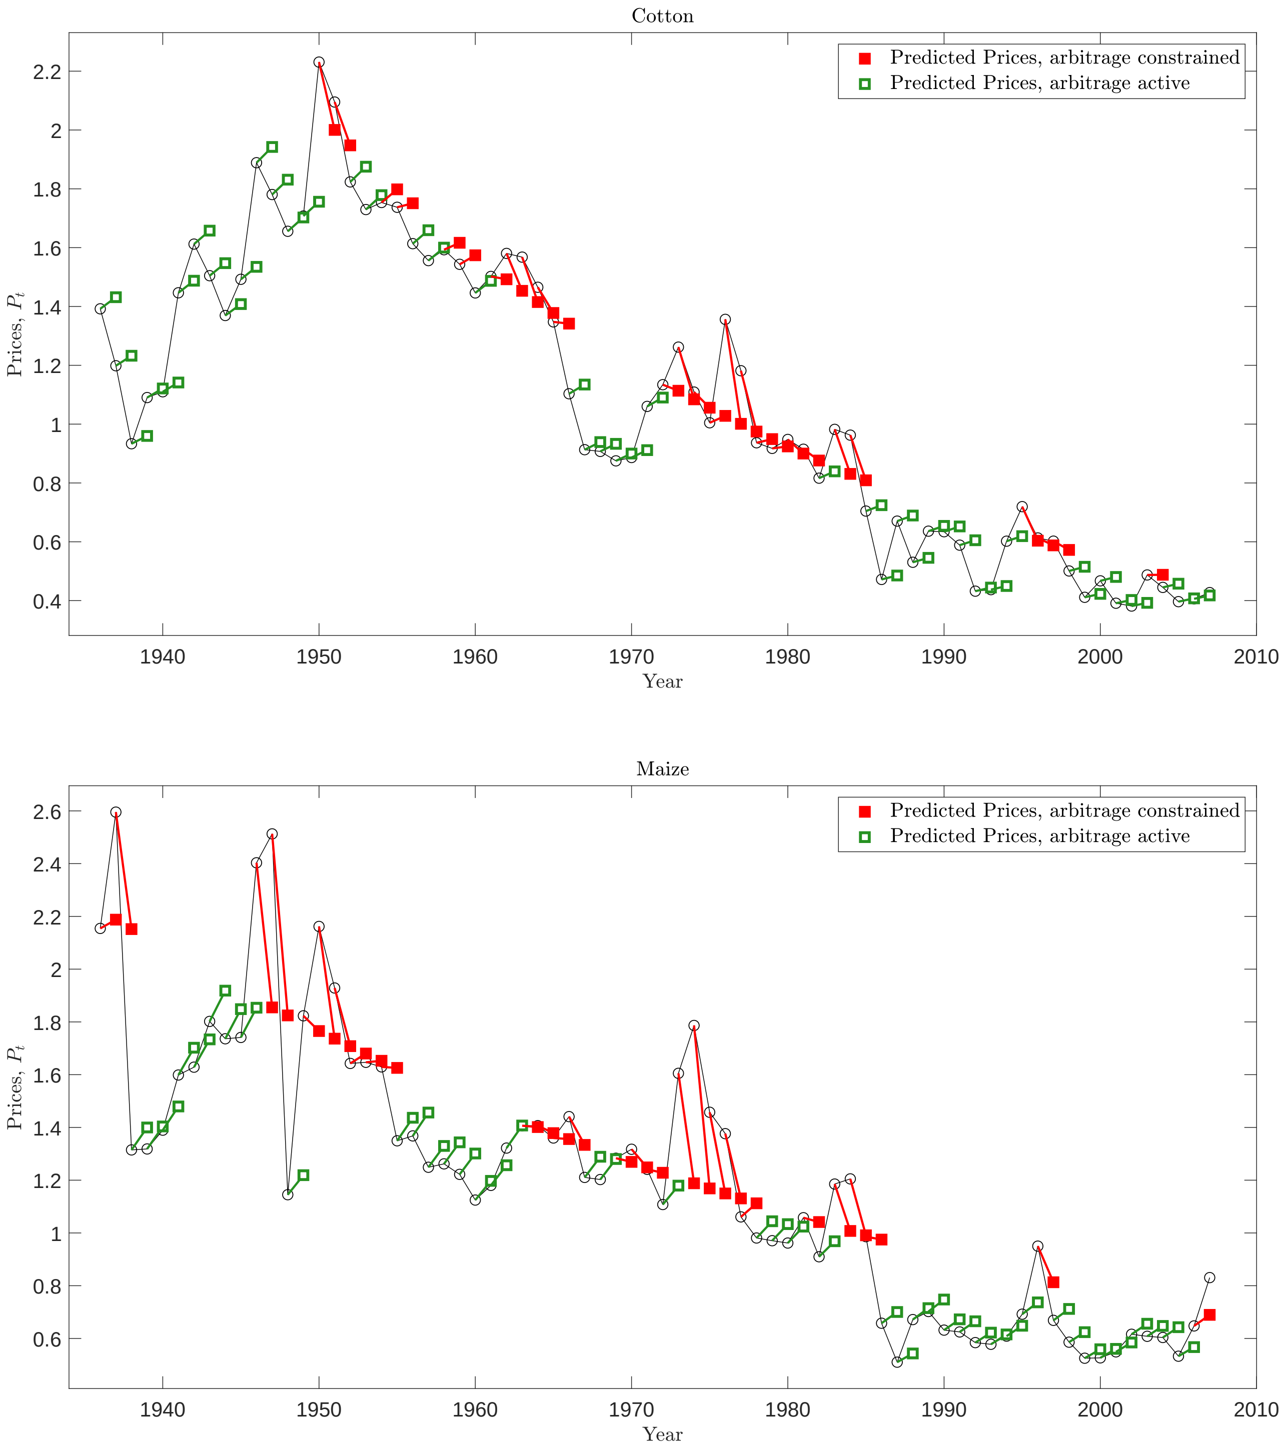

Figure 9: Annual prices and predicted prices. 1936-2007

Table 1
Estimates of $(1+r_0)$, p_0^*, $\lambda_0, (1+r_0)p_0^*$, and percentage of stockouts.†

	$1+r$	p^*	λ	$(1+r)p^*$	% stockouts
Cotton	1.0284 (0.2041)	1.5307 (0.3063)	0.9737 (0.0006)	1.5741 (0.0402)	38%
Maize	1.0647 (0.1217)	1.4060 (0.1711)	0.9836 (0.0010)	1.4969 (0.0626)	39%

† Estimated standard errors in parenthesis.