Food and Agriculture
Organization of the
United Nations

Developing an indicator of price anomalies as an early warning tool: A compound growth approach

Felix G. Baquedano

GIEWS – Global Information and Early Warning System on Food and Agriculture

Developing an indicator of price anomalies as an early warning tool: A compound growth approach

Felix G. Baquedano

FOOD AND AGRICULTURE ORGANIZATION OF THE UNITED NATIONS
Rome, 2015

CONTENTS

Introduction

The food price surge in global markets in 2007/08 and then again in 2011, has spurred a lot of interest in creating an early warning indicator to detect abnormal growth in prices in consumer markets in the developing world, where advance warning of an impending food crisis can be critical. In these countries, on an early warning basis, sometimes market prices are the only source of information available to assess the severity of a local shock to either access or availability of food. Because prices summarize information held by a large number of economic agents, including their expectations regarding likely short-term developments in supply and demand, they are ideal to use as the basis of an early warning indicator (Dawe and Doroudian, 2012).

The objective of this paper is to present an indicator developed by the Global Information Early Warning System of the Food and Agriculture Organization (GIEWS) of the Food and Agriculture Organization of the United Nations (FAO) that can be used to identify abnormal price changes. The GIEWS indicator of price anomalies (IPA) relies on a weighted compound growth rate that accounts for both within year and across year price growth. Compound growth rates, which tend to smooth out volatile times series, are commonly used in the financial literature to quickly rank the returns of stocks or other financial assets based on their growth (Anson et al., 2011).

Other recent work in this area is that of Araujo et al., (2012). Araujo et al., rely on estimating a seasonally adjusted price trend and measuring the deviation of the current price from this trend. In this model, an alert is generated when these deviations from the trend price are greater than one standard deviation of the historical deviations[1]. While their approach works well for the Sahelian country markets they analyse, where only one production season exists and the commodity used is a non-tradable (millet), for other developing country regions, where seasonal trends can be dampened by imports or overlapping harvests from more than one season, their model can lead to confusing results due to possibly erroneous trend estimates. Currently, Araujo et al.'s work has been further developed by the World Food Programme (WFP) into the Alert of Price Spikes (ALPS) indicator (WFP, 2014). However, as we will latter discuss, it relies on assumptions that are rather weak, which can lead to the indicator not giving an alert when one is clearly required.

The IPA presented in this paper, nonetheless, borrows the definition of abnormal price growth from Araujo et al.'s, but applies it in the context of compound growth rates, as we will shortly discuss in the methodology section which follows. While other data sources are used in developing the IPA, particularly for CPI's, price data is exclusively obtained from GIEWS Food Price Monitoring and Analysis Tool (FPMA-T) which covers 93 countries and over 1 300 price series. Following our discussion of the methodology, we present two examples, the first applies the IPA to millet retail prices in Maradi (Niger). We compare our results to those of Araujo et al. who also used this market to test their approach. Our second market is for red beans retail prices in Managua (Nicaragua). In this market, prices were at considerably high levels due to a sharp reduction in production, but still well below their levels during the food price surge in 2007/08 and later in

[1] Besides a price alert Araujo et al. also consider a price watch. A price watch is a positive price difference from the trend that is less than one standard deviation.

2011. This series will also allow us to test the IPA on a series where there are several harvests in one season. In the final section we present our conclusions and discuss how the IPA is currently used.

Before proceeding with the discussion on the methodology used to calculate the IPA, it is appropriate to caution the reader that the indicator developed is just a rough guide of market dynamics. As such, one cannot rely on it as the sole element to consider when giving a food security alert. Instead, its results must be weighed with other available information on market fundamentals and possible short-term policy shocks that can explain these price movements. This is especially important when evaluating whether or not the observed shocks in prices will persist or are transitory.

Methodology and data

The basis for the IPA is a weighted sum of two Compound Growth Rates (CGR). A CGR is a geometric mean that assumes that a random variable grows at a steady rate, compounded over a specific period of time (Anson et al., 2011). Because it assumes a steady rate of growth, the CGR smooths the effect of volatility of periodic price movements. The CGR is the growth in any random variable from a time period t_0 to T_n, raised to the power of one over the length of the period of time being considered, as highlighted in Equation 1.

$$CGR_t = \left(\frac{P_{t_n}}{P_{t_0}}\right)^{\frac{1}{t_n - t_0}} - 1 \qquad (1)$$

Where P_{t_0} is the price at the beginning of the period, P_{t_n} is the price at the end of the period, and $t_n - t_0$ is the time in months between t_0 and t_n. This indicator is commonly used in the financial sector as a way to rank by annual rates of growth stocks or portfolio valuations. In this world, high growth implies high returns. In the markets of the developing world, periods of high growth in prices can have mixed results since the net effect will depend on whether or not a household (or country) is a net consumer or producer. For a household that is a net consumer, extended periods of high price growth can seriously reduce access to food. However, for a household that is a net producer, the same period can lead to a significant increase in income and better access to food.

We modify the CGR further to account for seasonality, which is an important characteristic of agricultural markets. As is well known, commodity prices are at their lowest at harvest and increase during the lean period which ends with the following season's harvest. In some countries, the increase in prices can be significant if say only one production season exists, capacity to store is minimal and there is no dependency on trade, as is the case for some Sahelian countries. In other countries, seasonal price increases are less strong, either because there are multiple overlapping production seasons, high capacity to store, or strong trade flows. Seasonal trends can also be influenced by intra-annual shocks, such as drops in supply due to either production or trade shortfalls. Capturing in a single index both sources of variation (within the year and across years), is challenging, since the index could adequately measure excessive growth in prices within the year but not across years (or vice-versa). Therefore, we define two CGR's to account for this problem, a Quarterly Compound Growth Rate (CQGR) and an Annual Compound Growth Rate (CAGR).

To account for seasonality, we do not define the production periods to calculate both CGRs, unlike Araujo et al.'s proposed approach that estimates a seasonally adjusted price trend. Instead, we try to account for seasonal factors by calculating the CQGR and CAGR as a moving average over the immediately preceding 3-month or 12-month period of month *t*, respectively. This approach is similar to Dawe and Doroudian's moving average price, but also differs as the model assesses the growth in prices instead of the actual price level.

Once both the CQGR and CAGR are calculated, we proceed to defining the threshold that identifies abnormal growth in prices. We define as abnormal price growth, an absolute positive change in the CGR, either annual or quarterly, that is at least one standard deviation of the mean CGR over a specific month. By restricting the calculations of the average and the standard deviation of the CGR to a particular month, we try to further account for seasonal effects, as can be seen in Equation 2.

$$\left(\frac{CXGR_{yt}-\overline{CXGR_t}}{\hat{\sigma}_{CXGR_t}}\right) = X\_IPA_t^Z \begin{cases} 0.5 \leq X\_IPA_t^Z < 1 & Price\ Watch(X\_IPA_t^W) \\ X\_IPA_t^Z \geq 1 & Price\ Alert(X\_IPA_t^A) \\ o.w. & Normal\ (X\_IPA_t^N) \end{cases} \quad (2)$$

Where $CXGR_{yt}$ is either the quarterly or annual compound growth rate in month *t* for year *y*, $\overline{CXGR_t}$ is the average of either the quarterly or annual compound growth rate for month *t* across years *y*, $\hat{\sigma}_{CXGR_t}$ is standard deviation of either the quarterly or annual compound growth rate for month *t* over years *y*, and $X\_IPA_t^Z$ is either the quarterly or annual indicator of price anomaly (watch/alert/normal). We use one standard deviation as the relevant threshold since we want to minimize the probability of missing a significant market event. Events that deviate by more than one standard deviation from historical CGR's are easy to identify and probably do not require any information from a model. As proposed by Araujo et al., when giving a warning using either the annual or quarterly $CXGR_{yt}$, we also consider deviations that are less than one standard deviation but greater than or equal to a half standard deviation. These events are of particular interest as they provide an early warning of possible severe market disruptions, especially when they are close to one standard deviation. Therefore, two types of warnings are defined: a price alert $(X\_IPA_t^A \geq 1)$ and a price watch $(0.5 \leq X\_IPA_t^W < 1)$.

Dealing with Type I and Type II errors

As with any indicator, we have to balance the two types of errors. The first can occur when the indicator signals a price alert/watch when the reality is that markets are behaving normally. This is known as a Type I error (false positive), which leads to a significant reduction in confidence of the indicator. A Type II error occurs when no price alert/watch is given when one should have been issued, or false negative. This type of error is of more concern for early warning purposes as one does not want to miss an impending market shock. However, both of these errors are inter-related since mitigating for a Type I error will increase the probability of a Type II error and vice-versa.

To deal with these errors, we introduce two modifications to Equations 1 and 2 above. The first is to deflate the quarterly and annual compound growth rates of prices, shown in Equation 1 by the volatility of prices during the same period. As discussed earlier, the compound growth rate smooths out the volatility of prices, so by not accounting for risk the estimated growth rate might still be too high. By reducing the slope of the compound growth rate, the deviations with respect to the average of the compound growth rate at month *t*, will be smaller, thus reducing the probability of a Type I error. We define volatility as the standard deviation of log differences. To adjust the compound growth rates, we estimate a quarterly and annual measure of volatility and then we deflate the $CQGR_{yt}$ and $CAGR_{yt}$ by its respective measure, as shown in Equation 3.

$$CXGR_t \times \left(1 - \sigma_{[P_{t_0} - P_{t_n}]}\right) = vCXGR_t \quad (3)$$

Where the term $\sigma_{[P_{t_0} - P_{t_n}]}$ represents the standard deviation of prices over the period $P_{t_0} - P_{t_n}$, which is equal to either three or twelve months depending on which compound growth rate is calculated. And $vCGR_t$ is the volatility adjusted compound growth rate for either the annual or quarterly series. All other terms in Equation 3 are identical to those in Equation 1.

Our second modification seeks to reduce the probability of a Type II error by modifying the calculations of the mean and standard deviation in Equation 2. As stated in Equation 2, the historical standard deviation and mean is calculated giving equal weights to all-time periods in a price series. So, a period of high and volatile prices at the beginning of the period, will have the same weight as a more recent period of low and less volatile prices. The result of this is that the threshold for a CGR to be considered abnormal may be higher than it needs to be, thus resulting in a Type II error. Instead of using the simple mean and standard deviation, we substitute them for a weighted mean and standard deviation. The weights are increasing time weights, so the more recent past has a higher weight in the calculation of the mean and standard deviation than the beginning of the price series. The weighted mean is defined as follows:

$$\overline{vCXGR}_{Wt} = \frac{1}{\sum_{y=1}^{Y} w_y} \sum_{y=1}^{Y} w_y vCXGR_{yt} \quad (4)$$

Where $\overline{vCXGR}_{Wt}$ is the weighted average for month *t* of the *x* (quarterly or annual) volatility adjusted compound growth rate, w_y is the weight for year y, $vCXGR_{yt}$ is the unweighted volatility adjusted compound growth rate in year *y* in month *t*, and $\sum_{y=1}^{Y}$ is the summation operator over years *Y*. The weighted standard deviation is then estimated as follows:

$$\hat{\sigma}_{vCXGR\_Wt} = \sqrt{\frac{\sum_{y=1}^{Y} w_y \left(vCXGR_{yt} - \overline{vCXGR}_{Wt}\right)^2}{\sum_{y=1}^{Y} w_y (\hat{Y} - 1)/\hat{Y}}} \quad (5)$$

Where $\hat{\sigma}_{CXGR\_Wt}$ is the weighted standard deviation in month *t* of the $vCXGR_{yt}$, $\hat{Y}$ is the total number of weights and all other terms are identical to those defined in Equations 3 and 4. The calculations of the $X\_IPA_t^Z$, as stated in Equation 2, is then modified by substituting for the weighted mean and standard deviations as shown in Equation 6.

$$\left(\frac{vCXGR_t - \overline{vCXGR}_{Wt}}{\hat{\sigma}_{vCXGR\_Wt}}\right) = X\_IPA_t^Z \begin{cases} 0.5 \leq X\_IPA_t^Z < 1 & Price\ Watch(XP_t^W) \\ X\_IPA_t^Z \geq 1 & Price\ Alert(XP_t^A) \\ o.w. & Normal\ (XP_t^N) \end{cases} \quad (6)$$

Where all terms in Equation 6 have been previously defined in Equations 3 to 5. The indicator of price anomalies (IPA_t) for month *t* is then obtained by the following weighted sum:

$$IPA_t = \gamma Q\_IPA_t^Z + (1-\gamma)A\_IPA_t^Z \quad (7)$$

A critical part component of the IPA_t is the value of γ. In other words which growth rate is more important for the final indicator: quarterly or annual price growth? To answer this question we rely on elements of Principal Component Analysis (PCA), which is beyond the scope of this paper, but the work of Shelns (2013) and Smith (2002) provide a good introduction to the subject. In essence, what we are trying to determine is the individual weight of $vCQGR_{yt}$ and $vCAGR_{yt}$ in the sum of their variances. The PCA allows us to calculate the eigenvalues[2] for both of these compound growth rates. The sum of the eigenvalues is exactly equal to the sum of the variances of $vCQGR_{yt}$ and $vCAGR_{yt}$ in the variance covariance matrix. Thus, the ratio of each eigenvalue to the sum of the variances gives us the value for γ. In Annex 1, we provide a detailed example of these calculations.

Data sources and estimation

The sole source of price information used in constructing the IPA_t, is the GIEWS' FPMA-T which contains more than 1 300 prices series for 93 countries covering the main food grains (maize, rice, and wheat), as well as other regionally-important staples. These series are obtained from multiple national and international sources which are cited in the Tool. The FPMA-T also contains benchmarks prices for 22 major exporters of maize, rice and wheat. These series are available in nominal and real terms. The series in real terms, which are the ones we consider for the IPA_t, have been deflated using CPI (2005=100) information from the International Monetary Fund (IMF) or national data sources. The IPA_t is currently estimated on a monthly basis for all series and countries in the FPMA-T.

Applying the indicator to real world scenarios: Millet in Niger and Red Beans in Nicaragua

Millet, Retail, Maradi (Niger)

In this section, we begin by presenting the results of the IPA for the millet retail price series in Maradi (Niger), West Africa. The market of Maradi is considered a national and regional reference market in the West African Sahel, with shocks in this market being transmitted across several markets in the region (Araujo et al., 2012). We also choose this market/commodity pair because it is the main series used by Araujo et al.'s to highlight their method, to which we will briefly compare and contrast. Immediately following this discussion,

[2] An eigenvalue λ is the corresponding scale factor to an eigenvector. When there is some real number λ, such that $Av = \lambda v$, we say that v is an eigenvector of A and λ (the scale factor) its' eigenvalue.

we proceed to discuss the results for the wholesale prices red beans in Managua (Nicaragua), Central America.

The price series we use to highlight our method is obtained from the same source as Araujo et al.'s, the market information system of Niger through the FPMA-T. However, our series differs from theirs in two ways. First, our series is at the retail level and Araujo et al.'s is from a wholesale market. Second, the observation period is different, Araujo et al.'s work covers the period from January 1990 to October 2008 (226 observations), this work covers the period April 1995 to April 2012 (206 observations). Nonetheless, as we will highlight in our discussion, these differences play no important role towards the conclusions reached.

We start by evaluating the effect on the IPA_t of adjusting the Compound Growth Rates (CGRs) by the volatility of the price series and using a weighted mean and standard deviation. As discussed previously, the effect of adjusting the CGRs by the volatility of the price series was to further reduce its slope, under the assumption that the price growth may still be too high given the observed level of volatility. Thus, by reducing the slope, we expect average growth rates to be smaller reducing the probability of generating an alert. The weighing, by a time weight of both the standard deviation and mean should contribute to reduce the probability of a Type II error ($Pr(II)$), or failing to give a price warning when one should have been given, as it becomes a bit easier to generate an alert in a case where the recent past has been less volatile relative to a period further out in time, as the latter period would have more influence on the value of σ. Looking at the results from Table 1, for both the quarterly and annual CGRs, the means and standard deviations are significantly lower for the weighted and volatility adjusted series. But we find no significant difference with respect to values of the IPA's, whether quarterly, annual or the final weighted measure. Another important finding from Table 1, is that within year price changes explain most of the IPA, as implied by the value of γ of 89 percent.

Table 1: Comparison of the weighted and unweighted mean and standard deviation in the Indicator of Price Anomalies (IPA) for retail millet prices in Maradi (Niger)

	Mean	St. Dev.	Indicator
Compound Quarterly Growth Rate (CQGR)			$Q\_IPA_t^z$
Un-weighted no volatility adjustment	0.016	0.052	-0.55
Weighted and volatility adjusted	0.010	0.036	-0.30
t-stat (two tail)	-8.16 1/	-18.03 1/	1.48
Compound Annual Growth Rate (CAGR)			$A\_IPA_t^z$
Un-weighted no volatility adjustment	0.013	0.033	-3.20
Weighted and volatility adjusted	0.007	0.022	-1.18
t-stat (two tail)	-27.10 1/	-43.33 1/	0.90
Indicator of Price Anomaly (IPA)			IPA_t
Un-weighted no volatility adjustment			-0.83
Weighted and volatility adjusted			-0.39
t-stat (two tail)			1.51
Gamma 2/			0.894

Source: Author's calculations.
1/ (Pr<0.01).
2/ Gamma is estimated with the aid of principal component analysis.

However, these results tell us nothing of what the effect on the $\mathrm{Pr}(II)$ is when adjusting the CGR's for price volatility and using weighted means and standard deviations. To evaluate what happens to the $Pr(II)$, we determine this probability given 1 percent and 5 percent probabilities of committing a Type I error, or giving an alert when one should not have been given. We focus on evaluating $Pr(II)$ for the case when the IPA generates a price alert, or a value of the $IPA \geq 1$, and not a price watch, as the consequences of not giving a signal (or an incorrect one) is more significant for this case. The null hypothesis of our test is that the true mean of the IPA is represented by the average of the generated price alerts, under the given probability of a Type I error, or α level. And the alternate hypothesis is that the true mean is different from the observed mean. To carry out this test we rely on the standard normal Z-statistic. Given that we are evaluating the proportion of prices alerts generated by the IPA in our sample the Z-statistic is defined as:

$$Z = \frac{X-np}{\sqrt{npq}} = \frac{X-\mu}{\sigma} \quad (8)$$

Where Z is the Z-statistic; X is the critical value for rejecting the null under a given level of α; n is the number of observations; p is the proportion of the sample that is a price alert; q is $1-p$; μ is the mean of the number of alerts generated by the IPA, which in this case is equal to the total number of alerts; and σ is the standard deviation of the observed sample.

The total number of n observations in our sample is 169, with 37 observations lost in the estimation of the CGR's and calculations of the weighted mean and standard deviation. The unadjusted IPA has a mean of 13 price alerts, for a proportion of p alerts in the sample of 8 percent (Table 2). By contrast, the volatility adjusted and weighted IPA, or $vIPA$, has a mean of 25 alerts with a proportion of p alerts of 15 percent. For both the IPA and the PA , the critical z-value (the z-value at which the null hypotheses is rejected at a given level of α) generated by the observed μ in Equation 8, is ±1.96, when the $Pr(\alpha = 0.05)$ and ±2.58 when $Pr(\alpha = 0.01)$. And corresponding lower and upper Confidence Intervals (CIs) (or the interval of X values in Equation 8 at which the null hypotheses is not rejected) of 6.2/19.8 for the IPA when $Pr(\alpha = 0.05)$ and 4.1/21.9 for $Pr(\alpha = 0.01)$. Similarly, for the $vIPA$, these same CIs are shown in Table 2.

To test the alternative hypotheses that $\mu \neq 13(25)$ or $\mu = \bar{\mu}$, we must first generate a value for $\bar{\mu}$. The value of $\bar{\mu}$ is simulated from a distribution with a mean μ and standard deviation σ. Then 1 million random samples of $\bar{\mu}$ are obtained to calculate an equal number of lower and upper z-statistics, given the corresponding Confidence Intervals under the null hypotheses for the appropriate $Pr(\alpha)$, as shown in Table 2. Finally, each lower and upper z-statistics is then used to determine the area under the normal probability curve, with this area being equal to the $Pr(II)$. The results of these calculations are presented in Table 2 as the $Pr(\beta|\alpha) = Pr(II)$ and represent the average value of the total number of random samples. At first glance, we find that the average of the simulated $\bar{\mu}$ is fairly close to that of the mean for both the IPA and the $vIPA$ (Table 2). Moreover, irrespective of the $Pr(\alpha)$, the $Pr(\beta|\alpha)$, is also very similar for both indicators. However, $Pr(\beta|\alpha)$ is smaller for the $vIPA$ at the two evaluated $Pr(\alpha)$. And consistent with the theory as the $Pr(\alpha)$ increase the $Pr(\beta|\alpha)$ is also greater, but this probability continues to be smaller for the $vIPA$. The fact that the $vIPA$ has a lower $Pr(\beta|\alpha)$ relative to the IPA as a result of weighing the mean and standard deviation, leads us to conclude that the former indicator would be more preferable as an early warning indicator.

Table 2: Probabilities of giving (not giving) a price alert when one should not (should) have been given for three versions of a price warning indicator for millet in Maradi (Niger)

	IPA[1/]	vIPA[2/]	APAI[3/]
Alerts	13	25	35
Total Observations (N)	169	169	169
Proportion of Sample with an Alert (p)	0.08	0.15	0.21
Proportion of Sample not in Alert (q)	0.92	0.85	0.79
Standard Deviation	3.464	4.615	5.268
Mean	13	25	35
Simulated Mean	12.9970	25.0022	34.9870
Lower 95% Confidence Interval	6.21	15.95	24.67
Upper 95% Confidence Interval	19.79	34.05	45.33
Lower 99% Confidence Interval	4.08	13.11	21.43
Upper 99% Confidence Interval	21.92	36.89	48.57
Prob(α)[4/]	0.05	0.05	0.05
Prob(β|α)[5/]	0.8341	0.8339	0.8343
Power(1-β)	0.16590	0.16606	0.16569
Prob(α)[4/]	0.01	0.01	0.01
Prob(β|α)[5/]	0.9314	0.9313	0.9315
Power(1-β)	0.06863	0.06871	0.06848

Source: Author's calculations.
1/ IPA: Indicator of Price Anomalies.
2/ vIPA: Volatility adjusted compound growth rates and weighted mean and standard deviation.
3/ APAI: Araujo et al.'s price anomaly indicator.
4/ Prob(α): Probability of giving a price alert when one should not have been given.
5/ Prob(β|α): Probability of not giving a price alert when one should have been given.

We now compare our results to Araujo et al.'s method. We begin by briefly discussing their identifying process of an abnormal price series. Since their work is published, and for the sake of brevity, we will not enter into too much detail into their method, but only highlight those components that contrast with the IPA_t. Nonetheless, we highly encourage the reader to peruse this very important work, which contributes[3] greatly to the subject. The first major difference between both methods is how a price deviation is measured. Araujo et al. measure the deviation of a price P_t in month *t* to its trend $\overline{P}_t$ in month *t*. As previously highlighted, we instead measure deviations as the difference between the quarterly (annual) compounded growth rate $vCXGR_{yt}$ and its weigthed average $\overline{vCXGR}_{Wt}$ in month *t*. To estimate $\overline{P}_t$ Equation 9 is fitted.

$$P_t = \alpha T_t + \sum_{s=1}^{12} \beta_s M_{st} + \varepsilon_t \qquad (9)$$

Where P_t is the log normal real price in month *t* for a particular commodity/market pairing, T_t is a time trend meant to capture long-term movements, and M_{st} is a seasonal dummy equal to zero during the lean season and one during the production cycle. This definition of the seasonal dummy relies on a very important assumption for early warning purposes. Araujo et al. assume that, if prices start to increase during the production cycle and especially close to the harvest, this signals a potential bad year, particularly in years

[3] A major contribution of this work is the process Araujo et al. use to determine leading markets, both within a country and across countries.

when a production shortfall is anticipated. So, by measuring the deviations in these months from the trend, one can anticipate a potential crisis. In Niger, the production season roughly begins with early planting in May and concludes in November with the harvest. While this assumption is sound, especially for the Sahelian countries they analyse, it cannot be replicated across different country/commodity settings where potentially a production shortfall in one season can be made-up in a subsequent period or by increasing imports, as tends to happen in other regions of the world that are better integrated to world markets. The IPA_t handles seasonal effects differently by trying to compare the growth of a price P_t, both quarterly and annually, to a historical growth rate for that same month, thus avoiding an ill-defined seasonal dummy.

A more concerning problem with this fitted trend is the assumption implied by Equation 8, and reaffirmed by Araujo et al., and that is that prices are trend stationary. This is rarely the case as most time series of market prices tend to be non-stationary. A non-stationary series implies that a shock in month t on P_t will be permanent, instead of dying out in subsequent periods. Moreover, the stationary properties of a series can change depending on the length of the sub-sample chosen (Banerjee et al., 1998). Also different conclusions can be reached depending on which unit root test is chosen as they tend to be of very low power (Banerjee et al., 1998 and De Boef and Granato, 1999).

Araujo et al. rely on two types of unit root test to evaluate the hypotheses of non-stationarity in their time series, the Augmented Dikey Fuller (ADF) (Fuller, 1976) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) (Kwiatkowski et al., 1992). The latter test has been shown not only to have a higher power (or a lower probability of a Type II error) but it also tests a different hypotheses. The null hypotheses of the KPSS is that the series is stationary, the opposite of the ADF. We supplement these tests with the Dikey Fuller-Generalized Least Squares (DF-GLS) (Elliot et al., 1996) and the Phillips-Perron test (PP) (Phillips and Perron, 1998), both with higher power than the ADF. We organize our results from the test with the least power (ADF) to the one with the highest power (KPSS). To be consistent with Araujo et al., we test the hypotheses that the Maradi (Niger) millet price series is trend stationary. Our results contrast with Araujo et al.'s[4], where both the ADF and KPSS conclude that prices are trend stationary. In our case, the ADF concurs with their previous work but the KPSS rejects the null (Table 3). The DF-GLS supports the findings of the KPSS, leading us to conclude that the millet retail price series for Maradi (Niger) is non-stationary. The different conclusion is definitely a result of using a different time period and different series (retail instead of wholesale).

[4] We cannot fully compare our results to Araujo et al., since they only provide the p-value and LM-statistic, instead of providing both the test statistic and critical value for the rejection of H0 like we do.

Table 3: Unit root tests

	ADF	Phillips-Perron	DF-GLS	KPSS
Millet, Retail, Maradi (Niger)				
Test Statistic	-4.17 1/	-3.67 2/	-2.5	0.47 1/
Test Result-Araujo et al.	Rejected	-	-	Not Rejected
Critical Value for rejection of HO	-4.01	-3.44	-2.65	0.22
Null Hypothesis (HO)	Non-stationary	Non-stationary	Non-stationary	Stationary
Red beans, Retail, Managua (Nicaragua)				
Test Statistic	-3.41 3/	-3.17 3/	-3.27 2/	0.37 1/
Critical Value for rejection of HO	-3.14	-3.14	-2.95	0.21
Null Hypothesis (HO)	Non-stationary	Non-stationary	Non-stationary	Stationary

Source: Author’s calculations.
1/ (Pr<0.01).
2/ (Pr<0.05).
3/ (Pr<0.10).

The practical implications of our unit root results is that, in order for Equation 8 to represent a stable system, we would have to estimate it using the first difference of the log price (ΔP_t). Implying that Equation 9 would be estimated only with the seasonal dummy as the differencing would zero out the time trend. However, again for consistency with Araujo et al., we will assume stationarity, based on the results of the ADF and PP, and estimate Equation 9 as specified. In practice, however, basing our tests on this very strong assumption could potentially lead to erroneous signals. Explicitly Araujo et al.’s price anomaly indicator $(APAI_t)$ is defined in Equation 10.

$$\left(\frac{I_t}{\hat{\sigma}_I}\right) = \left(\frac{P_t - \overline{P}_t}{\hat{\sigma}_I}\right) = APAI_t \begin{cases} 0.5 \le APAI_t < 1 & Price\ Watch \\ APAI_t \ge 1 & Price\ Alert) \\ o.w. & Normal \end{cases} \tag{10}$$

Where I_t is the difference of P_t in month *t* from its trend, $\overline{P}_t$, $\hat{\sigma}_I$ is the standard deviation of I_t over the whole series and $APAI_t$ represents the value of the ratio of I_t to $\hat{\sigma}_I$. A value greater than one implies that the deviation of P_t from its trend is higher than one standard deviation.

Figure 1, shows the results for the $vIPA$ and the $APAI$. As previously discussed, we fitted[5] Araujo et al.’s model under the assumption that the millet price series for Maradi (Niger) was trend stationary. As shown in Figure 1, this is a very weak assumption, since it becomes increasingly difficult for the $APAI$ to generate a warning, given that the warning threshold becomes higher as the price trend continues to increase. For example, during the period of October 2011 to April 2012, where prices continually increased as a result of a sharp decline in production (FAO, 2011), no warning of any kind was given by the $APAI$, contrary to the signals sent by the $vIPA$ for the same period. When we compare the power of the $APAI$, relative to the $vIPA$, we also find that the former has a smaller power, implying a higher $Pr(\beta|\alpha)$ (Table 2).

[5] The fitted values of Equation 8 are available in the Annex.

Figure 1: Indicator of Price Anomalies (IPA) and Araujo et al. price anomaly indicator for retail millet prices in Maradi (Niger)

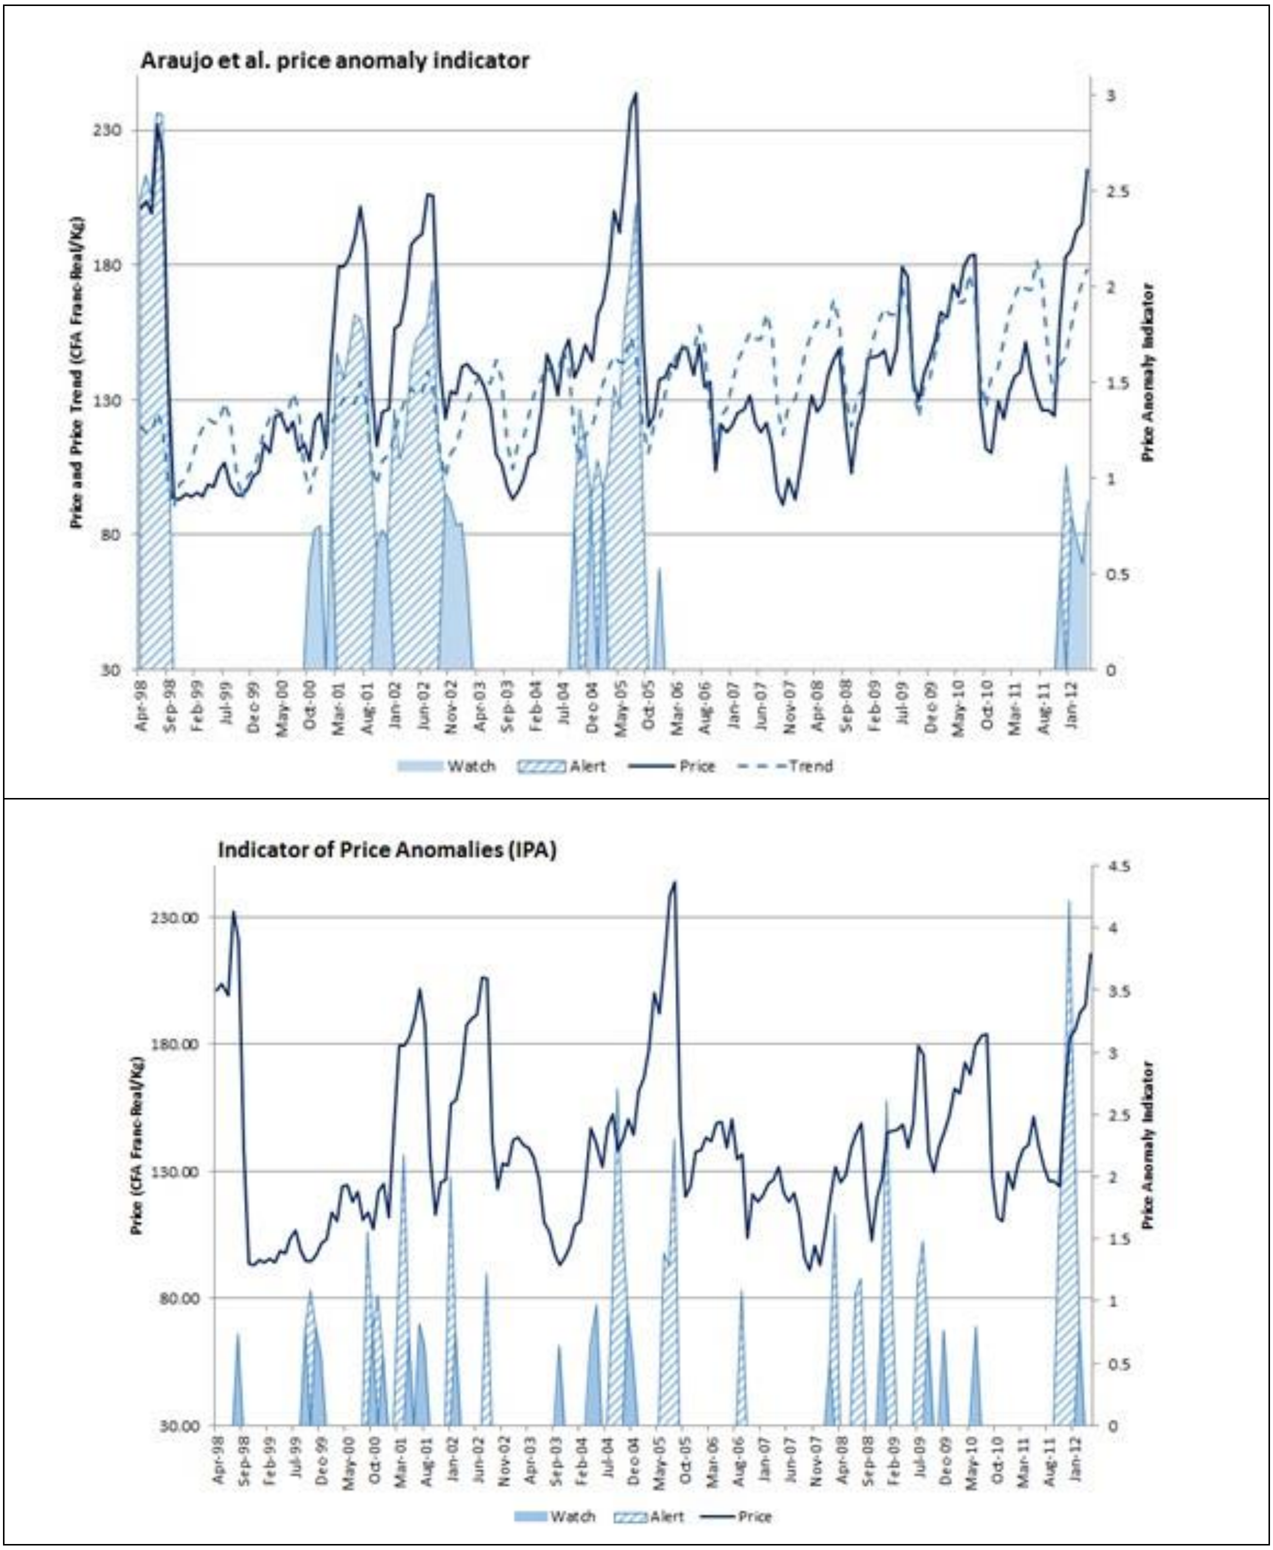

Source: Author's calculations. National CPI (2005=100) from BCEAO used to deflate prices.

Red Beans, Retail, Managua (Nicaragua)

The red beans retail market in Managua (Nicaragua) is of interest to us for three reasons. First, red beans in Central America are cultivated in three consecutive seasons. This allows us to compare the IPA to Araujo et al.'s method in a market with a more complex seasonal structure. Second, Nicaragua is the main producer and intra-regional exporter of red beans to the region during the lean season, which roughly goes from April to August. Since red beans markets in this region are well integrated, this allows us to test the efficiency of estimating a regional commodity α for the IPA instead of one for each commodity/pairing in the FPMA-T. Finally, Central American red bean markets have been severely affected by a drop in production in Nicaragua and Honduras, the second largest producer and exporter. This allows us to highlight the indicator during a high price period.

In Nicaragua, the harvest of the first bean season occurs from August to October, the second season harvest goes from December to January and finally the third season harvest starts in mid-February ending in mid-March. The red beans retail price series in the Managua market (the Capital) is obtained from the Ministry of Agriculture of Nicaragua through the FPMA-T. The period used for our analysis covers January 2000 to March 2014 (171 observations). Prices have been deflated using national CPI (2005=100) as reported by the IMF.

The first result that we discuss for this series, is the effect of using a regional estimate for γ as opposed to a country/market specific weight on the warnings given by theIPA_t. To estimate the regional γ, we use the 11 price series for red beans (both retail and wholesale) for El Salvador, Honduras and Nicaragua available in the GIEWS' FPMA-T for the period February 2007 to March 2014. We use this shorter period because it is the period where all series coincide with each other. For the Nicaragua red beans retail price in the Managua market, we use the full period of the data. Both alphas were estimated with the aid of principal component analysis (see Annex 1 for a more detailed discussion). The estimate of γ for the red beans/Managua is 0.899, the regional estimate for γ is somewhat lower at 0.846 (Table 4). In our discussion of the effects of using a regional γ, we use the estimates from the weighted and volatility adjusted compound growth rates, or $vIPA$, which we compare to the unweighted series immediately after this discussion.

Table 4: Changes in warnings in the indicator of price anomalies from a country specific alpha to a regional alpha

From (Country Alpha 0.899)	To (Regional Alpha 0.846)		
	Normal	Watch	Alert
Normal	-	4[1/]	0
Watch	0	-	1
Alert	0	0	-

Source: Author's calculations.
Note: Results from a z-test (two-tail). H0: Mean difference is equal to zero.
1/ (Pr<0.01).

Using a regional γ as opposed to a commodity/market specific one, results in four price points changing from being considered normal to a watch (Table 4). With the exception of the $vIPA$ for January 2005 that went from -0.22 with the commodity/market specific α to 0.89 with the regional one, the other three points where close to being considered a watch already, June 2005 went from an $vIPA$ of 0.49 to 0.65, May 2008 from 0.44 to 0.62 and July 2008 from 0.41 to 0.52. We also find that the power of the $vIPA$ calculated with a regional weight is higher than the one with country specific γ, thus implying a lower $Pr(\beta|\alpha)$.Given that all these adjustments happened in periods of high prices (see Figure 2) it would suggest that we would've missed highlighting these events, suggesting a potential improvement of the indicator. Thus in our further discussions in this section we will only concentrate on the results obtained from using a regional γ.

When using a weighted mean and standard deviation for the calculation of the IPA, we reach the same conclusion as before, there is no significant difference from weighing in the final indicator but mean and standard deviations for the quarterly and annual indicator are significantly different (Table 5). In both the quarterly and annual compound growth rates the means are significantly higher for the weighted estimate versus the unweighted one, a result opposite to the Niger series where they were lower. For the standard deviation the results are mixed, for the quarterly growth rate they are significantly higher for the weighted estimate, while for the annual growth rate they are significantly lower (Table 5). These results imply that in the weighted version the threshold to generate a waning with the $Q\_IPA_t^z$ is much higher than with the $A\_IPA_t^z$, which has a significantly lower standard deviation.

Table 5: Comparison of weighted and un-weighted mean and standard deviation in Indicator of Price Anomalies (IPA) or retail prices of red beans in Managua (Nicaragua)

	Mean	St. Dev.	Indicator
Compound Quarterly Growth Rate (CQGR)			$Q\_IPA_t^z$
Un-weighted no volatility adjustment	-0.0013	0.036	0.61
Weighted and volatility adjusted	0.0002	0.199	0.25
t-stat (two tail)	5.30 1/	15.00 1/	-0.99
Compound Annual Growth Rate (CAGR)			$A\_IPA_t^z$
Un-weighted no volatility adjustment	0.010	0.257	0.74
Weighted and volatility adjusted	0.024	0.224	0.93
t-stat (two tail)	6.57 1/	-22.30 1/	1.55
Indicator of Price Anomaly (IPA)			IPA_t
Un-weighted no volatility adjustment			0.63
Weighted and volatility adjusted			0.36
t-stat (two tail)			-0.90
Alpha			0.846

Source: Author's calculations.
1/ (Pr<0.01).

Figure 2: Indicator of Price Anomalies (IPA) and Araujo et al. price anomaly indicator for retail prices of red beans in Managua (Nicaragua)

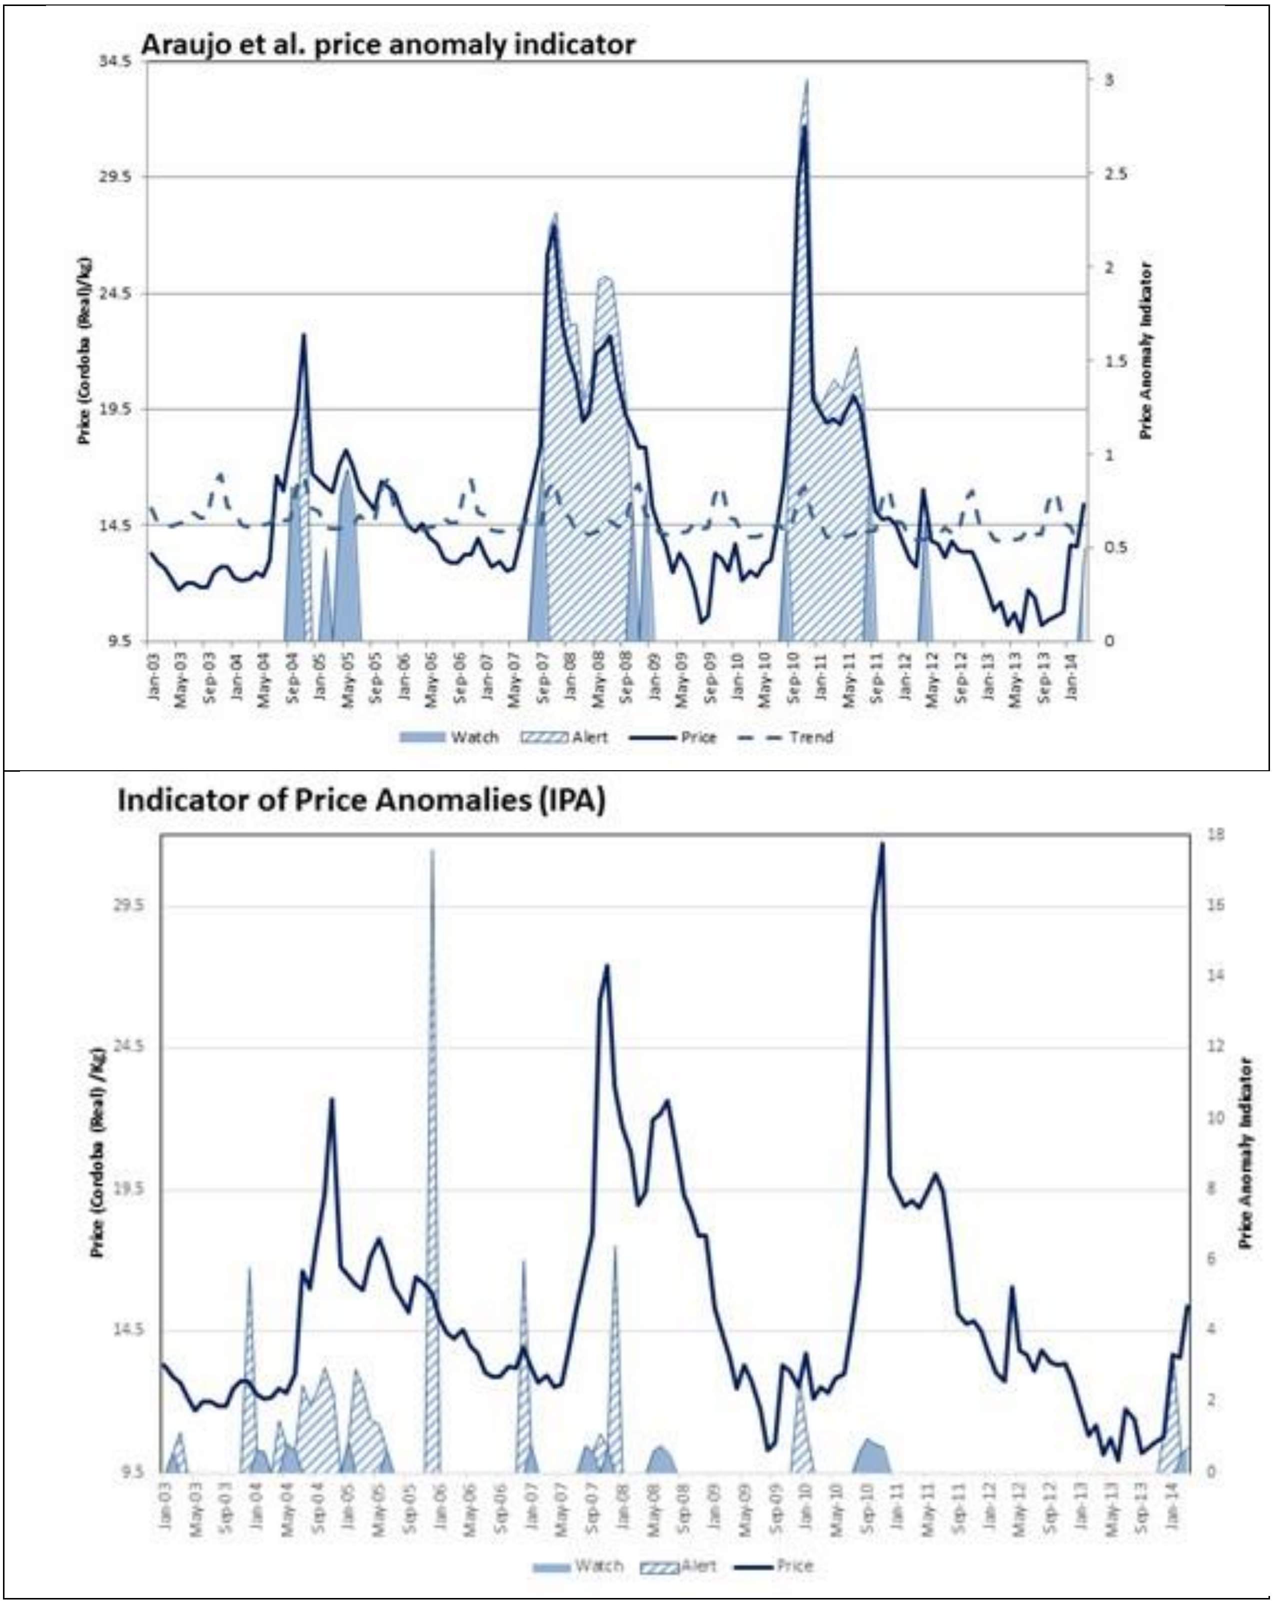

Source: Author's calculations. National CPI (2005=100) from the IMF.

To compare the warnings from the $APAI$ to the $vIPA$, we again rely on the assumption that prices are trend stationary. When evaluating our results from the $vIPA$ to the $APAI$, we find that a significant number of series go from a state of warning (watch/alert) to normal in the $vIPA$, implying that no alert should have been given (Table 6). But more importantly we also find 16 series that would have been missed, since they were considered normal when a warning should have not probably been issued. As previously discussed for the millet Niger case, assuming trend stationarity makes it significantly difficult to issue an alert and this is again highlighted in Figure 2 and in the number of the series that changed from not being a warning to a watch or alert, as just discussed. A clear example of this is the recently-concluded 2013/14 bean production season. Following a bumper red bean crop in 2012/13, farmers in Nicaragua significantly reduced area planted for 2013/14 due to the low prices resulting in a sharp decline in production. The prospects of sharp increases in consumer prices, in response to the shortfall in supplies, were very apparent by the end of the second season harvest in December 2013. However, the $APAI$, did not provide any warning until March 2014, the end of third bean season. By contrast, the $vIPA$ has been signaling some sort of warning (watch/alert) since December 2013 (FAO, 2014). This example further highlights the implications of using a non-stationary series to define the warning threshold and the effects of possibly ill-defined seasonal dummies. Moreover, the higher power of the $vIPA$, relative to the $APAI$, provides support to the higher risk one runs with the latter indicator when applying to a market with multiple seasons and more dynamic market structures (Table 7).

Table 6: Changes in warnings between the indicator of price anomalies and the Araujo et al.'s price anomaly indicator for the red bean retail market in Managua (Nicaragua)

From $APAI$[2/]	To $vIPA$[1/]		
	Normal	**Watch**	**Alert**
Normal	-	9	7
Watch	9	-	5
Alert	16	7	-

Source: Author's calculations.
1/ vIPA: Volatility adjusted compound growth rates and weighted mean and standard deviation.
2/ APAI: Araujo et al.'s price anomaly indicator.

Table 7: Probabilities of giving/not giving a price alert when one should not/should have been given for four versions of a price warning indicator for red beans in Managua (Nicaragua)

	vIPA-CG[1/]	vIPA-RG[1/]	IPA-RG[2/]	APAI[3/]
Alerts	19	20	34	31
Total Observations (N)	135	135	135	135
Proportion of Sample with an Alert (p)	0.14	0.15	0.25	0.23
Proportion of Sample not in Alert (q)	0.86	0.85	0.75	0.77
St.Dev	4.041	4.128	34	4.887
Mean	19	20	34	31
Simulated Mean	18.994	19.996	33.998	30.993
Lower 95% Confidence Interval	11.08	11.91	24.11	21.42
Upper 95% Confidence Interval	26.92	28.09	43.89	40.58
Lower 99% Confidence Interval	8.59	9.37	21.01	18.41
Upper 99% Confidence Interval	29.41	30.63	46.99	43.59
Prob(α)[4/]	0.05	0.05	0.05	0.05
Prob(β|α)[5/]	0.8342	0.8341	0.8344	0.8342
Power(1-β)	0.16584	0.16589	0.165578	0.16578
Prob(α)[4/]	0.01	0.01	0.01	0.01
Prob(β|α)[5/]	0.9314	0.9314	0.9316	0.9314
Power(1-β)	0.06857	0.06861	0.0684	0.06856

Source: Author's calculations.
1/ vIPA: Volatility adjusted compound growth rates and weighted mean and standard deviation.
2/ IPA: Indicator of Price Anomalies.
3/ APAI: Araujo et al.'s price anomaly indicator.
4/ Prob(α): Probability of giving a price alert when one should not have been given.
5/ Prob(β|α): Probability of not giving a price alert when one should have been given.

Conclusions

Currently, the indicator of price anomalies mainly informs the Food Price Monitoring and Analysis website and bulletin, published on monthly intervals by the FAO/GIEWS. Validation of the signals generated by the indicator of price anomalies is undertaken by weighing market fundamentals (production, trade, stocks) and local market/policy conditions. At times, even if a price warning is given by the indicator, it is not always considered, for example an alert is generated but in the month where the harvest is undergoing and the expectations for production are very good.

The main advantage of the indicator of price anomalies proposed in this paper is its simplicity. It can be used in different markets without concerns as to whether or not the market year has been well defined. The indicator is also more robust, since you do not have to deal with potentially non-stationary series, as you are directly evaluating price growth which is stationary. The price anomaly indicator by being calculated over a particular month over various years, also allows one to answer the question of whether or not a small change in price is normal for that month. A next step in this work is to adapt the IPA_t for prices series where weekly information is available. The IPA_t, cannot yet be considered a true early warning indicator, since it is summarizing price movements that are monthly.

References

Anson, M.J.P., Fabozzi, F.J., Jones, F.J. 2011. The handbook of traditional and alternative investment vehicles. John Wiley and Sons, Inc., Hoboken, New Jersey

Araujo, C., Araujo-Bonjean, C., Brunelin, S. 2012. Alert at Maradi: Preventing food crisis by using price signals. *World Development*, Vol 40(9), pp. 1882-194

Dawe, D. and Doroudian, A. 2012. A simple price monitoring tool to assess monthly changes in food prices. FAO, Rome

Elliott, G. R., T. J. Rothenberg, and J. H. Stock. 1996. Efficient tests for an autoregressive unit root. *Econometrica*, Vol. 64, pp. 813–836

Food and Agricultural Organization (FAO), 2011. Global Food Price Monitor. 8 November 2011. Rome, Italy. www.fao.org/giews

Fuller, W. A. (1976). Introduction to Statistical Time Series. New York: John Wiley and Sons.

Jackson, J.E., (1991). A User's Guide to Principal Components. New York: John Wiley and Sons.

Jolliffe, I.T., (2002). Principal Component Analysis, second edition (Springer).

Kwiatkowski, D.; Phillips, P. C. B.; Schmidt, P.; Shin, Y. (1992). "Testing the null hypothesis of stationarity against the alternative of a unit root". *Journal of Econometrics*, Vol. 54 (1–3).

Phillips, P. C. B., and P. Perron. 1988. Testing for a unit root in time series regression. *Biometrika* 75: 335-346.

Shelns, J., 2003. A tutorial of principal component analysis: Derivation, discussion and singular value decomposition. Version 1, University of California, San Diego.

Annex 1

Principal Component Analysis

The Principal Component Analysis (PCA) is a statistical procedure that uses an orthogonal transformation to convert a set of observations of possibly correlated variables into a set of values of linearly uncorrelated variables called principal components (Jackson, 1991 and Jolliffe, 2002). For example, if we have the matrices $\boldsymbol{Y}$ and $\boldsymbol{X}$, each of dimensions $m \times n$, and they are related by an orthogonal matrix $\boldsymbol{\Lambda}$ (of rank $r \leq m$), then $\boldsymbol{\Lambda}$ contains up to *m* eigenvalues λ (Shelns, 2003). An eigenvalue λ is the corresponding scale factor to an eigenvector. When there is some real number λ, such that $Av = \lambda v$, we say that v is an eigenvector of A and λ (the scale factor) its eigenvalue. The eigenvalues also poses a useful feature for our purposes, and that is that the sum of the eigenvalues is equal to the sum of the variances in the variance covariance matrix of *m* random variables. The transformation by the orthogonal matrix $\boldsymbol{\Lambda}$ results in such a way that the first principal component accounts for as much of the variability in the data as possible (Jackson, 1991 and Jolliffe, 2002). This allows us to use the λ's as a proxy to determine which compound growth rate (the $CQGR_t$ or $CAGR_t$) should have more weight in the IPA_t.

In essence the relationship we are trying to estimate is the following:

$$\sigma^2_{\sum Q,A} = \sigma^2_Q + \sigma^2_A = \lambda_Q + \lambda_A \quad \text{(A.1)}$$

Where Q refers to the $CQGR_t$ and A to $CAGR_t$, $\sigma^2_{\sum Q,A}$ is the sum of the individual variances of Q and A, σ^2_Q and σ^2_A are the individual variances for Q and A, and λ_Q and λ_A are the corresponding eigenvalues for Q and A, respectively. The weight α used to calculate the IPA_t is then calculated in Equation A.2.

$$\alpha = \frac{\lambda_Q}{\sigma^2_{\sum Q,A}} \quad \text{(A.2)}$$

The first step in calculating the eigenvalues is to calculate the variance covariance matrix, represented by the matrix Σ_{QA} in Equation A.3, where the diagonal elements are the variances of Q and A and the off diagonals represent their covariances. Before calculating matrix Σ_{QA}, it is important to centre Q and A around their means so that the series with the highest variance does not inadvertently dominate and be identified as the principal component erroneously (Shelns, 2003).

$$\Sigma_{QA} = \begin{bmatrix} \sigma^2_Q & \sigma^2_{Q,A} \\ \sigma^2_{Q,A} & \sigma^2_A \end{bmatrix} \quad \text{(A.3)}$$

The second and final step is to solve for the second order determinant defined in Equation A.4, where Σ_{QA} is the variance covariance matrix, the matrix λ contains the eigenvalues that we are searching for and the matrix I represents the identity matrix.

$$\left[\Sigma_{QA} - \lambda \mathrm{I}\right] = \begin{bmatrix} \sigma_Q^2 - \lambda_i & \sigma_{Q,A}^2 \\ \sigma_{Q,A}^2 & \sigma_A^2 - \lambda_i \end{bmatrix} = 0 \qquad \text{(A.4)}$$

The solution for the determinant of Equation A.4 can then be found by solving for Equation A.5, which is nothing more than a quadratic equation as shown in Equation A.6. This reduces the problem to finding the solution to a quadratic equation using the formula stated in Equation A.7 where the first result corresponds to λ_Q and the second to λ_A.

$$\left(\sigma_Q^2 - \lambda_i\right)(\sigma_A^2 - \lambda_i) - \left[\sigma_{Q,A}^2\right]^2 = 0 \qquad \text{(A.5)}$$

$$\lambda_i^2 - \left(\sigma_Q^2 + \sigma_A^2\right)\lambda_i + \left(\sigma_Q^2 \times \sigma_A^2 - \left[\sigma_{Q,A}^2\right]^2\right) = 0 \qquad \text{(A.6)}$$

$$\lambda_i = \frac{-(\sigma_Q^2+\sigma_A^2) \pm \sqrt{\left(\sigma_Q^2+\sigma_A^2\right)^2 - 4\times 1\times\left(\sigma_Q^2\times\sigma_A^2-\left[\sigma_{Q,A}^2\right]^2\right)}}{2\times 1} \qquad \text{(A.7)}$$

Annex 2

Coefficients of fitted trend for Araujo et al. Indicator of Price Anomalies (IPA)

	Coefficient	t-Stat
Millet, Retail, Maradi (Niger)		
Trend	0.0024	5.82[1/]
Seasonal Dummies		
D1	4.571	60.52[1/]
D2	4.633	62.76[1/]
D3	4.669	59.83[1/]
D4	4.697	49.41[1/]
D5	4.680	47.78[1/]
D6	4.677	47.64[1/]
D7	4.737	47.20[1/]
D8	4.686	42.87[1/]
D9	4.492	55.91[1/]
D10	4.401	60.12[1/]
D11	4.486	64.34[1/]
D12	4.504	61.26[1/]
n		205
F-stat		8837[1/]
R-Square		0.997
Unit Root Test of Residuals		
DF-GLS		-1.55
KPSS		0.54[1/]
Red beans, Retail, Managua (Nicaragua)		
Trend	-0.000402	-1.36
Seasonal Dummies		
D1	2.739	56.71[1/]
D2	2.694	54.78[1/]
D3	2.690	58.34[1/]
D4	2.689	52.22[1/]
D5	2.697	47.91[1/]
D6	2.702	48.43[1/]
D7	2.729	50.68[1/]
D8	2.712	50.22[1/]
D9	2.715	47.05[1/]
D10	2.809	37.93[1/]
D11	2.835	35.08[1/]
D12	2.746	47.47[1/]
n		171
F-stat		3239[1/]
R-Square		0.994
Unit Root Test of Residuals		
DF-GLS		-2.95[2/]
KPSS		0.39[1/]

Source: Author's calculations.
1/ (Pr<0.01).
2/ (Pr<0.05).