

A Refreshing Perspective on Seasonality

Luc Christiaensen, Christopher L. Gilbert, and Jonathan Kaminski

Overview

Conventional wisdom: *We all know about seasonality in rural livelihoods, but it is very unclear precisely what it is we know, and it is considered less and less frequently by development economists and policy makers.*

Findings:

- *The commonly used methods to estimate the seasonal gap in crop prices (dummy variable or moving average deviation) can yield substantial upward bias. This bias can be partially circumvented by using more parsimonious methods (trigonometric or sawtooth).*
- *Seasonal price variations are substantial and widespread. Prices during the peak months are on average estimated to be 28 percent higher than those during the troughs in the seven African countries examined. Food price volatility is much higher still, with seasonal variation explaining only a fraction of overall food price volatility (17 percent on average).*
- *Among staple crops, the seasonal gap is highest for maize (33 percent) and lowest for rice (16.5 percent). The gap is on average two and a half to three times larger than on the corresponding international reference market (South African Futures Exchange for maize, and Bangkok international market for rice). This finding suggests that there is substantial excess seasonality in African staple markets.*
- *Country-specific circumstances do not appear to affect the extent of seasonality—the main exception being maize prices in Malawi.*
- *Finally, evidence from Tanzania shows that food price seasonality can translate into seasonal variation in caloric intake, with seasonal differences in caloric intake of 10 percent among poor urban households and rural net food sellers.*

Policy message: *Seasonality in African staple prices is widespread, well above what is observed in international reference markets, and shown to affect caloric intake among certain population groups. These findings confirm that it is premature to ignore seasonality in the African development debate. Entry points for reducing food price seasonality include better access to financial markets for households, more secure storage at the village level, reduction in transport costs, and increased intra-African food trade. The relative effectiveness of these policies requires further investigation.*

The Issue: Is Seasonality in Food Prices and Food Consumption Important?

Seasonality in food prices and consumption was much studied in the 1990s, and was shown to be associated with significant fluctuations in hunger and nutrition. Since then, the topic has largely disappeared from the policy debate, especially among development economists. The general perception of improved integration of local food markets may have partly motivated this neglect. Nevertheless, substantial seasonality in price movements is still possible, even when domestic food markets are better integrated. This can happen, for example, if the timing of production is highly correlated across markets and commodities, and if domestic food markets are poorly integrated with world markets (or those in neighboring countries).

A certain degree of seasonality in food prices is unavoidable. Agricultural production is cyclical, necessitating intertemporal arbitrage. Storage costs ensue, driven by postharvest loss and the opportunity cost of capital. This drives a wedge between prices before and after the harvest. This price gap can be compounded by market power along the marketing chain and in storage, high transaction costs due to poor infrastructure and fuel costs, transport monopolies, and credit constraints for producers and traders.

Thus, seasonality is widely acknowledged to be part of African (rural) livelihoods. But what exactly do we know? The most salient aspect of seasonality in Africa is food price seasonality, as well as its effects on food consumption and nutrition.¹ Despite wide recognition that food prices are seasonal, there has been little systematic analysis of the extent of seasonal variation across countries and markets, or even how this should be measured. A companion study further assesses (for one country, Tanzania) whether seasonality in food prices also leads to similar variations in food consumption, on which there is even less evidence (Kaminski, Christiaensen, and Gilbert 2016). The findings show that it is premature to ignore seasonality in the African development debate.

The Analysis: Challenges in Estimating Seasonality

The Data

The study examines the extent of seasonal patterns in food prices for 13 crops and food products across 193 market locations in seven countries (Burkina Faso, Ethiopia, Ghana, Malawi, Niger, Tanzania, and Uganda). The data, which cover 2000–12, come mainly from national statistical offices and (in the case of Uganda)

a private marketing agency. They cover the most important staple cereals (maize, millet, rice, sorghum, and teff), cassava, several important fruits and vegetables, and eggs. The data set yields a total of 1,053 location-food crop pairs.

The problem of short data series. An important statistical problem that arises in analyzing seasonality is to disentangle seasonal movements from the longer-term trend in prices or consumption on the one hand, and irregular movements on the other. This problem is acute when the number of data points is small. For example, with 10 years of data, the study will have only 10 observations on January prices. Seasonality estimates could therefore be unduly influenced by irregular price movements. A second statistical problem is that data series are often incomplete. Sample start and end dates differ across series, but the more serious problem is gaps within the series. Short data samples and missing observations in monthly series are frequent challenges in representing seasonality in developing country prices, and this study is no exception. These challenges only multiply when analyzing seasonality in consumption, with only five years of monthly consumption data being available.

Which Empirical Approach to Take?

A measure that is commonly used in the development literature to characterize seasonality is the seasonal gap, that is, the ratio of the highest over the lowest monthly price (or consumption), or the ratio of the highest monthly deviation from the trend over the lowest monthly deviation. So, measuring seasonality requires estimating a trend and estimating the monthly deviation from it. A traditional approach to estimate the trend has been to use a 12-month centered moving average. This average has the advantage of enabling the annual increment to vary across time, but the approach is weak when the sample is short and there are missing values. For example, using moving averages sacrifices the initial and final six months of data, which is a major loss when time series are short. To calculate the moving average, data gaps must be filled (with little guidance on how to do so). Both disadvantages can be overcome by using a monthly dummy variable regression with a trend instead. The estimated monthly dummies then represent the deviations from the trend, that is, the seasonal factors. (When specified in first differences, the monthly dummy variable regression typically enables a stochastic trend as well.)

An important attraction of these unrestricted approaches to seasonality measurement is that no a priori structure is imposed on the form of seasonality (each month's deviation from the trend is calculated separately). The approach can then easily accommodate crops for which there are two annual harvests. The disadvantage is that a long time series is necessary to obtain accurate estimates, since only a single observation per year is used to estimate each seasonal factor/month. This is especially problematic when samples are short and the peak and trough months that are necessary to calculate the seasonal gap are not known a priori by the analyst, as is the case in many developing countries. Intuitively, although the empirical estimates of the seasonal factors (or monthly dummies) are each unbiased, each empirical estimate of a seasonal factor represents a draw

from a distribution, which usually deviates slightly from its true point value. As a result, by taking each time the maximum and minimum values of all the seasonal factors, the gap will be overestimated. The upward bias is larger the shorter is the sample and the less well defined is the seasonal pattern.

The extent of this problem is shown by the Monte Carlo simulations reported by Gilbert, Christiaensen, and Kaminski (2017). The dummy variable procedures perform poorly for samples of the length typically found in developing countries (about a decade or so). Taking 10 years of data with no seasonality genuinely present, apparent but spurious seasonal factors purport to imply a seasonal gap of 15 percentage points. When seasonality is genuinely present, biases are still likely with short data series, although much smaller (4 percentage points when using 10 years of data). The biases double again (to an estimated 8 percentage point gap) when seasonality is poorly defined. As predicted, the bias in the seasonal gap declines as sample size increases. Each seasonal factor is then estimated more precisely, such that the maximum and minimum identified by the data more likely represent the true peak and trough price months.

To mitigate such estimation bias in short samples, the study proposes the use of two more parsimonious approaches. By imposing a harvest-based pattern on the monthly seasonality factors, parsimonious seasonality models reduce the influence of any single monthly price. Consequently, there is a much lower probability of incorrect peak and trough identification (for example, through an error of a single month in either direction). Two alternative specifications are considered (box 16.1 provides further detail):

- *Trigonometric structure.* Here the analysis assumes that price variations follow a pure sine wave over time (defined by two cosine parameters). Although the trigonometric specification is parsimonious, it is restrictive in that the postharvest price decline is symmetric with respect to the preharvest price rise. In practice, for many crops, prices drop more rapidly postharvest than they rise in the remainder of the crop year.

Box 16.1 Metrics and Method

This box explains more rigorously the more parsimonious approaches to measuring seasonality in prices.

Trigonometric Seasonality

In this case, the seasonal pattern is defined by a pure sine wave. The simplest two-parameter sinusoidal trigonometric seasonality representation for month m is

$$s_m = \alpha \cos\left(\frac{m\pi}{6}\right) + \beta \sin\left(\frac{m\pi}{6}\right) \quad (\text{B16.1.1})$$

box continues next page

Box 16.1 Metrics and Method (continued)

With trending data, the estimating equation for the price in month m of year y is

$$\Delta p_{ym} = \gamma + \Delta s_m + u_{ym} = \gamma + \alpha \Delta \cos\left(\frac{m\pi}{6}\right) + \beta \Delta \sin\left(\frac{m\pi}{6}\right) + u_{ym} \tag{B16.1.2}$$

where u is an error term.

Equation B16.1.2 is estimable by least squares. The seasonal factor, s_{m^r} may be re-expressed as a pure cosine function:

$$s_m = \lambda \cos\left(\frac{m\pi}{6} - \omega\right) \tag{B16.1.3}$$

where $\lambda = \sqrt{\alpha^2 + \beta^2}$ and $\omega = \tan^{-1}(\alpha/\beta)$. The parameter λ measures the amplitude of the seasonal cycle and implies a seasonal gap of 2λ . If the specification is valid, least squares estimation of equation B16.1.1 yields unbiased and consistent estimates of the α and β coefficients in equation B16.1.2. However, the implied seasonal gap, 2λ , is a nonlinear, nonnegative function of these estimates and will therefore also be biased upward. The trigonometric approach is illustrated using tomato price data from Morogoro, Tanzania (figure B16.1.1, panel a).

Sawtooth Seasonality

The pre- and postharvest price hike symmetry of the trigonometric specification limits its relevance to seasonality in Africa, where prices drop more rapidly than they rise. An alternative

Figure B16.1.1 Illustrating Trigonometric and Sawtooth Seasonality

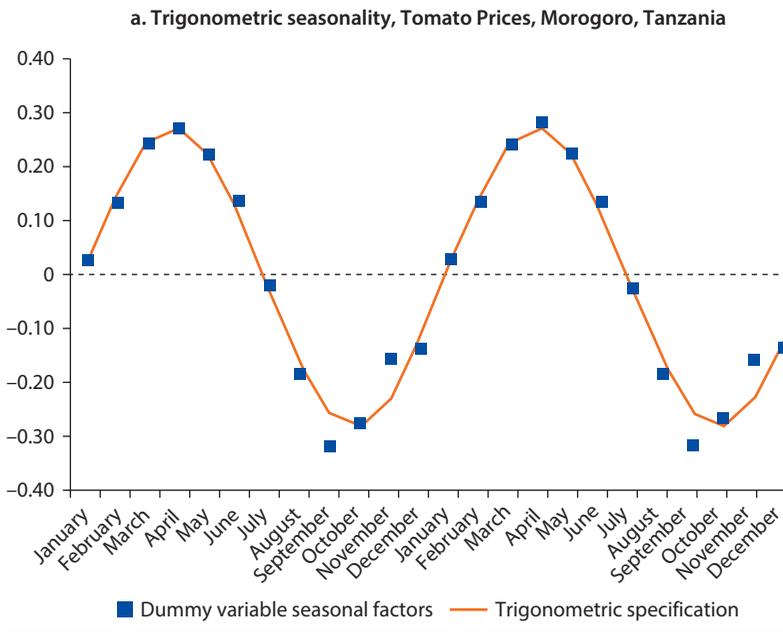
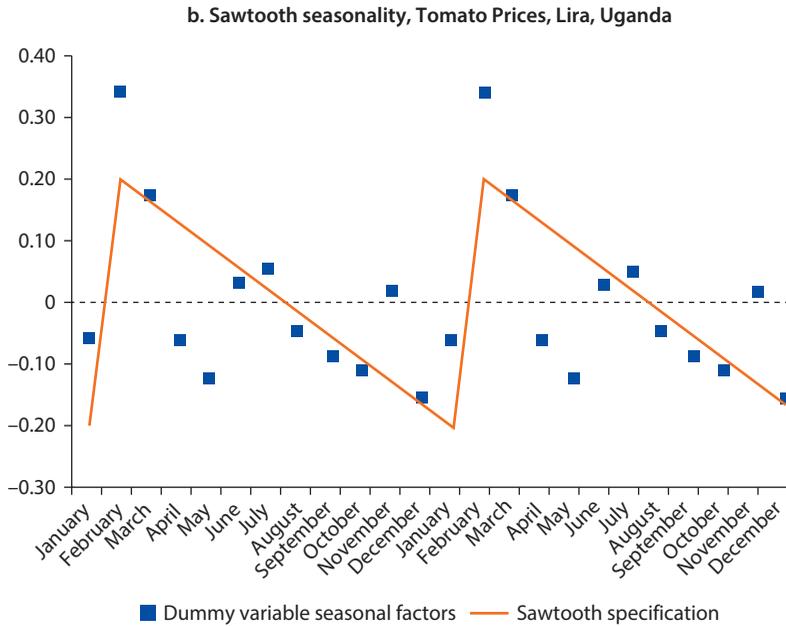


figure continues next page

box continues next page

Box 16.1 Metrics and Method (continued)

Figure B16.1.1 Illustrating Trigonometric and Sawtooth Seasonality (continued)



parametric specification is a sawtooth function in which prices fall sharply postharvest and then rise at a steady rate through the remainder of the crop year. Suppose the peak seasonal factor of λ occurs in month m^* and that price falls by the seasonal gap of 2λ to $-\lambda$ in the harvest month m^*+2 . The seasonal factor then rises steadily by an amount $\frac{\lambda}{5}$ over the remainder of the year. Conditional on knowing the peak price month, m^* , the amplitude parameter λ may be estimated from the regression

$$\Delta p_{ym} = \gamma + \Delta s_m + u_{ym} = \gamma + \lambda \Delta z_m(m^*) + u_{ym} \tag{B16.1.4}$$

Here, $\Delta z_m(m^*)$ equals -1 if $m = m^* + 1$ or $m = m^* + 2$ and $\frac{1}{5}$ otherwise. The study estimates by performing a grid search choosing the value for m^* that gives the maximum R^2 fit statistic. The illustration of sawtooth seasonality in figure B16.1.1, panel b, is for tomato price seasonality in Lira, Uganda.

- *Sawtooth structure.* This imposes an asymmetric variation in prices, a big drop at harvest, and a gradual recovery afterward. This variation fits most of the single annual harvest crops and locations.

Monte Carlo simulations show that parsimonious seasonal models are likely to be preferred to the standard dummy variable procedure for estimating the extent of seasonality when data samples are short or seasonal processes are

poorly defined. These are typical circumstances for developing country food crop price data. These procedures substantially reduce the bias resulting from the use of dummy variable estimators of the seasonal gap. Their limitation is that they will perform poorly for crops in which there are two harvests per year.

To discriminate between these empirical specifications, the preferred estimate is obtained by a three-step procedure:

- The estimates of the trigonometric and sawtooth specifications are compared with those of the dummy variable model. If the F test rejects both models, the dummy variable estimates are retained. This step helps for example to select the better (more flexible) model for location-crop pairs with two or more harvests.
- If the F test rejects one but not both parsimonious procedures, the nonrejected parsimonious model is taken as an acceptable simplification of the dummy procedure, reducing the bias in the seasonal gap estimates.
- Finally, if the F test fails to reject the trigonometric and sawtooth models, one of them is selected based on fit, as measured by the R^2 statistic.

Given different crops and agricultural settings, the preferred empirical approach will vary. Typically, the dummy variable model fits better when the seasonal pattern is well defined and does not conform to the sinusoidal or sawtooth patterns. This finding is generally true in cases when there are two harvests in the annual agricultural cycle. Using these rules, of the 1,053 location-food crop pairs, the dummy variable specification is preferred in 168 instances (many of which are in equatorial Uganda, where double cropping is common). The trigonometric specification is preferred in 625 instances, and the sawtooth specification in the remaining 260 instances. Although in a proportion of the cases, the success of the trigonometric model reflects a genuinely sinusoidal pattern, in other cases, in which the seasonal pattern is weakly defined or the data set is very short, the trigonometric specification may be chosen solely on the grounds of parsimony.

The Results: Seasonality Is Still Very Much Present

Because the sample size varies mainly by country, the seasonality estimates for the different commodities can be partially purged from potential overestimation by regressing the 1,053 estimated gaps for each commodity-location pair on the commodity type, the nature of the market (retail or wholesale), and a set of country dummies. The average estimated seasonal gap for each commodity is reported in table 16.1 (controlling for the nature of the market and country effects), together with the share of locations in which the null of no seasonality is rejected.

The study asks six questions:

- What is the extent of seasonality?
- How much of the overall price variation is due to seasonality?

Table 16.1 Average Estimated Seasonal Gap and Seasonal R^2 , by Food Crop

<i>Food crop</i>	<i>Seasonal gap (%)</i>	<i>Seasonality significant (%)</i>	<i>Seasonal R^2</i>
Tomatoes	60.8	64.0	0.21
Plantain/matoke	49.1	66.7	0.32
Oranges	39.8	50.0	0.16
Maize	33.1	93.2	0.25
Bananas	28.4	39.1	0.13
Teff	24.0	100.0	0.15
Beans	22.9	81.7	0.21
Sorghum	22.0	48.2	0.15
Millet	20.1	41.3	0.16
Cassava	18.8	26.9	0.08
Rice	16.6	68.2	0.17
Cowpeas	17.6	27.8	0.09
Eggs	14.1	64.0	0.18
Average	28.3	59.3	0.17

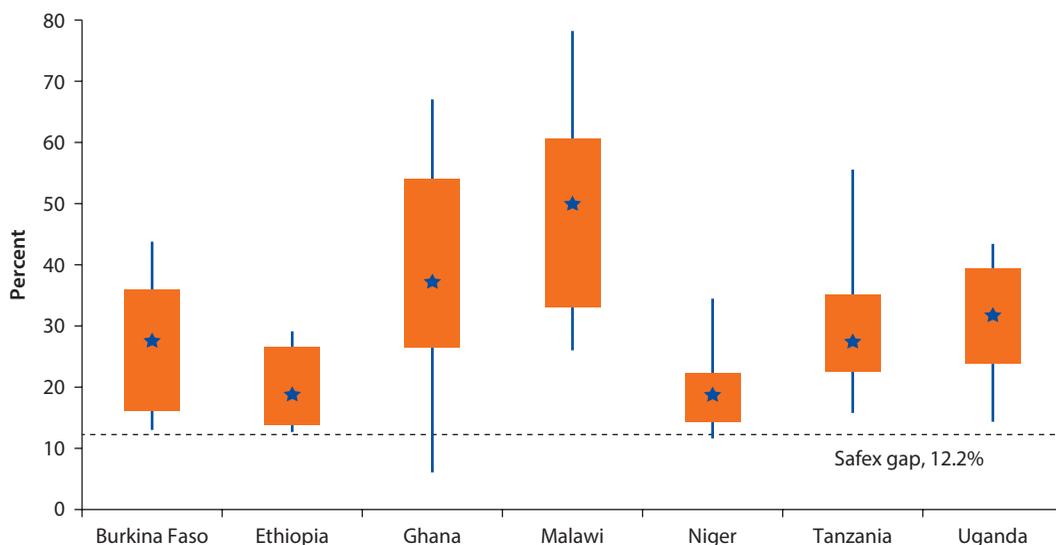
Note: The table reports the regression estimates of the average seasonal gap in wholesale markets, the proportion of locations for which the preferred gap estimate is based on coefficients that are significant at the 95percent level, and seasonal R^2 by crop. The averages reported in the bottom row of the table are the unweighted averages across crops.

- Is measured seasonality excessive?
- Are seasonal price variations widespread?
- Are country effects important?
- Does price seasonality translate into seasonality of consumption?

What Is the Extent of Seasonality?

Fruits and vegetables are the most prone to seasonality. Fruits and vegetables have the highest gaps, as intuitively expected. Their production is highly seasonal, and they are highly perishable. Cassava and eggs, which are produced throughout the year, are among the commodities with the lowest seasonality (first column in table 16.1).

Maize also displays substantial seasonality. Among staples, the clearest evidence for seasonality is in maize prices (33 percent), for which seasonality is about twice as high as that of rice (17 percent). The higher seasonality of maize among the cereals is expected, given its lower storability and greater postharvest loss compared with millet and sorghum. With Africa being a growing importer of rice (which is becoming more important in urban diets), rice markets are more closely linked with the international markets. Part of African rice production is also irrigated. Figure 16.1 provides a visual summary of the distribution of the seasonal gap for maize in the seven countries. The vertical lines measure the

Figure 16.1 Maize Price Seasonal Gaps in Africa Are Well Above the International Benchmark

Source: Gilbert, Christiaensen, and Kaminski 2017.

Note: End points of the vertical lines indicate the maximum and minimum; the endpoints of the boxes, the 80 and 20 percentiles; and the stars, the median seasonal price gap across the different markets within each country.

range of the seasonal gap across markets in each country, which is the distance between the largest and smallest gaps. The rectangles demarcate the interdecile range between the 20 and 80 percent points in the gap distribution. The median gap is indicated by a star. Malawi has the highest median gap, but the Malawian gap distribution has substantial overlap with the Ghanaian and, to a lesser extent, the Tanzanian and Ugandan distributions.

What Is the Contribution of Seasonality?

How much of overall price variation is explained by seasonality? The final column of table 16.1 reports the seasonal R^2 statistics, which show the proportion of the monthly variation in food prices attributable to seasonality. Among crops, plantain/matoke and maize show the largest contribution (0.32 and 0.25, respectively), and cassava and cowpeas the lowest seasonal R^2 values (0.08 and 0.09, respectively). Across countries, seasonality appears to explain around 17 percent of overall price variability. It increases to 27.7 and 21.3 percent in Niger and Burkina Faso, respectively, where agriculture is mainly rainfed and highly seasonal. Although the bulk of intra-annual price variability is not related to seasonal fluctuations, for some crops (maize) and countries (especially in the Sahel), its contribution appears nonetheless non-negligible.

Are Seasonal Gaps Excessive?

The study compares estimated seasonal gaps in these countries with the gaps observed in two international markets: the Johannesburg futures market (SAFEX), providing the reference price for white maize in Southern and East Africa,

and the Bangkok spot rice price. The estimated seasonal gaps are 12.2 percent for SAFEX white maize and 5.1 percent for Bangkok rice. Typically, maize price seasonality is significantly greater than this (figure 16.1). The unsurprising conclusion is that maize prices in Sub-Saharan Africa show substantial seasonal variation, and this variation is on average two and a half times as large as that on world markets. The extent of regular seasonal variability in rice prices is around half that of maize prices, but is on average three times the size of the seasonal variability in world rice prices. There is substantial excess price seasonality in some of Africa's key staples.

How Widespread Is Seasonality?

Seasonality is larger than in the international reference market in virtually all the 133 wholesale maize markets and 107 wholesale rice markets examined. There are only two centers where the estimated gap for maize is lower than the SAFEX gap of 12.2 percent (Ho in Ghana and Niamey in Niger), and three where the gap is lower than the 5.1 percent gap in the Bangkok spot market for rice (Santhe, Lizulu, and Neno in Malawi). The occurrence of excess seasonality is widespread. Nonetheless, there is also substantial variation in the extent of seasonality across locations within countries, as in Malawi, Ghana, and Tanzania (for maize and rice). These findings counsel caution against overgeneralization from case studies, and underscore the need for differentiated and targeted interventions.

Are Country Effects Important?

The study shows that 30.4 percent of the variation in the preferred seasonal gap measure is attributable to the crop, 14.5 percent to the (market) location, and only 0.5 percent to the country and 0.4 percent to the market level (wholesale or retail). Country-specific variation is not statistically significant. But maize and especially Malawi are exceptions. Maize price seasonality is particularly striking in Malawi. This country effect is confirmed when comparing maize seasonal gaps across locations in Malawi and Tanzania close to their common border. The prevalence of high seasonal gaps throughout Malawi, together with the sharp drop in the gap moving north into Tanzania, suggests that the high Malawian gaps are the result of political or institutional factors specific to the country, rather than agro-economic factors. To that extent, it should be possible to reduce some of the more extreme instances of seasonal maize price variation in Malawi, including by facilitating cross-country trading, which would also benefit Tanzania.

What Is the Effect on Food Consumption?

Follow-up analysis in Tanzania shows that caloric consumption also displays seasonal patterns, although limited on average, when looking across the country. This seasonal variation in caloric intake is further shown to be linked to seasonal fluctuations in Tanzania's maize and rice prices, indicating that households are on average not fully able to smooth their consumption. The urban poor and rural net food sellers are the most affected, with their caloric intake about 10 percent

higher during the peak month compared with the trough. Food price seasonality has real welfare effects.

The Implications

Policy pointers. Together, the findings indicate that the current neglect of seasonality in the policy debate is premature. Although it is not a major contributor to food price volatility, food price seasonality often proves substantial, with annual peak prices for maize across the countries studied on average 33 percent higher than those during the trough month. Moreover, the peak-trough markup is almost three times as high as the peak-trough markup observed in the international reference market, suggesting substantial excess seasonality. In some countries (especially Malawi) and several markets in the study countries, the gap is even higher. Food price seasonality at times also translates into seasonal variation in caloric intake. This is especially harmful when it affects children in their first 1,000 days of life. The findings draw attention to better access to financial markets for households, more secure storage at the village level, reduction in transport costs, and increased intra-African food trade as important policy areas and possible policy entry points. From a broader measurement perspective, the findings also underscore the importance of correcting for seasonality in food prices when constructing expenditure-based welfare and poverty measures, a largely ignored issue among poverty measurement practitioners so far.

Future research. Especially the results for seasonality in food consumption are based on limited data and are suggestive rather than conclusive. They are also based on a single country. Future work will bring in further survey waves and allow generalization to other Living Standards Measurement Study–Integrated Surveys on Agriculture countries. In the meantime, as long as time series remain limited (10 to 15 years), more use could be made of more parsimonious methods in measuring seasonality. It will also be important to extend the discussion to a wider range of welfare indicators, including indicators of longer-term impacts, such as child growth and nutrition.

Note

1. Other aspects include the supply and demand of labor, or the seasonal recurrence of certain diseases.

Additional Reading

This chapter draws on:

Gilbert, Christopher L., Luc Christiaensen, and Jonathan Kaminski. 2017. "Food Price Seasonality in Africa: Measurement and Extent." *Food Policy* 67: 119–32. doi 10.1016/j.foodpol.2016.09.016.

Kaminski, Jonathan, Luc Christiaensen, and Christopher L. Gilbert. 2016. "Seasonality in Local Food Markets and Consumption: Evidence from Tanzania." *Oxford Economic Papers* 68 (3): 736–57. doi 10.1093/oenp/gpw013.

Other key references:

- Dercon, S., and P. Krishnan. 2000. "Vulnerability, Seasonality and Poverty in Ethiopia." *Journal of Development Studies* 36: 25–53.
- Ellis, F., and E. Manda. 2012. "Seasonal Food Crises and Policy Responses: A Narrative Account of Three Food Security Crises in Malawi." *World Development* 40: 1407–17.
- Ghysels, E., and D. R. Osborn. 2001. *The Econometric Analysis of Seasonal Time Series*. Cambridge: Cambridge University Press.
- Goetz, S., and M. T. Weber. 1986. "Fundamentals of Price Analysis in Developing Countries' Food Systems: A Training Manual to Accompany the Microcomputer Software Program MSTAT." Working Paper 29, Michigan State University, East Lansing.
- Sahn, D. E., and C. Delgado. 1989. "The Nature and Implications for Market Interventions of Seasonal Food Price Variability." In *Seasonality in Third World Agriculture*, edited by G. J. Gill, 179–95. Baltimore: International Food Policy Research Institute and Johns Hopkins University Press.