WP/11/64

IMF Working Paper

Reviving the Competitive Storage Model: A Holistic Approach to Food Commodity Prices

Yanliang Miao,

Weifeng Wu, and Norbert Funke

INTERNATIONAL MONETARY FUND

 WP/11/64

IMF Working Paper

Reviving the Competitive Storage Model: A Holistic Approach to Food Commodity Prices

Prepared by Yanliang Miao, Weifeng Wu, and Norbert Funke*

Authorized for distribution by Lorenzo Giorgianni, Abebe Selassie, and Norbert Funke
March 2011

Abstract

We revive in this paper the empirical relevance of the competitive storage model by taking a holistic approach to food commodity prices. We augment the seminal Deaton and Laroque (1992, 1996) model by incorporating more comprehensive and realistic supply and demand factors: output and demand trends, shocks to the yield, and time-varying interest rates. While the computational burden increases exponentially, the augmented model succeeds in replicating all four key patterns of food commodity prices. Our simulation and comparative statics also show that (i) the long-run declining trend of food prices may come to a halt or even reverse due to the shifting balance between supply and demand; (ii) short-run price fluctuations are mainly attributable to sizeable, though low-probability, shocks to output such as inclement weather; and (iii) the impact of monetary policy, though small in normal times, is nonlinear and asymmetric, and can become large if the real rate passes a certain threshold.

JEL Classification: C63, E21, O13, Q11

Keywords: Food commodity prices, Competitive storage model, Rational expectation equilibrium

Email: ymiao@imf.org, wwu19@jhu.edu, and nfunke@imf.org

* We would like to thank Andy Berg, Oliver Blanchard, Helene Rey, and Jonathan Wright for useful discussions and Abe Selassie for his support and guidance throughout the project. The main part of the paper was completed while Weifeng Wu was an intern and then a short-term visiting fellow at the Fund. The views expressed here are solely the authors and do not necessarily reflect those of the International Monetary Fund.

Contents

1 Introduction

"To the ruler, the people are heaven; to the people, food is heaven."

-An ancient Chinese proverb

Despite the centrality of food to both the people and the ruler, food price developments and availability remain poorly understood. The 2007-08 surges in global food and fuel prices are yet another example of their short-run volatility defying long-run tranquility. Using wheat prices as one of many possible examples (Figure 1(a)), there are at least four noteworthy patterns for food prices: they (i) decline slightly in real terms in the long run[1]; (ii) are very volatile in the short run; (iii) exhibit high autocorrelations—those for wheat are 0.84, 0.67, and 0.54 for the first three orders; and (iv) are characterized by asymmetric price movements, with a skewness of 0.97 and an excess kurtosis of 1.48[2]. The asymmetry manifests itself in two related but distinctive ways: large price hikes are often followed by drops rather than the other way around; and rapid hikes are often followed by prolonged and gradual declines. Similar patterns are observable for maize, rice, and soybeans.

Since at least the 1950s, numerous studies have attempted to understand the characteristics of food commodity prices but have generally not been able to explain all four patterns. Relative long-term stability with a small declining trend has been attributed to both the exercise of market power by Northern manufacturers and the low price elasticity of demand for primary commodities (Prebisch (1959) and Singer (1950)). The argument of Sir Arthur Lewis (1954) that in tropical countries the poverty and unlimited supplies of labor at subsistence wage held down the price of tropical products has also been used to explain long-run price behavior (see Deaton (1999) and Deaton and Laroque (2003)).

[1] Our monthly price data started in 1957M1; using annual data, Cashin and McDermott (2002) have documented a small long-run declining trend dating back to 1862.

[2] Skewness measures the degree of symmetry. A positive number indicates that the distribution is skewed to the right, i.e., the mean being larger than the median for an unimodal distribution. Excess kurtosis measures whether the empirical distribution from data is more peaked or flat relative to a normal distribution, which has an excess kurtosis of 0. A positive excess kurtosis suggests that the distribution is more peaked than is a normal distribution.

On short-run variability and asymmetric price movements, Gustafson (1958) made an important contribution by introducing intertemporal storage arbitrage and supply shocks, which predated Muth (1961)'s idea of rational expectations. Building on Gustafson's model of optimal demand and Muth's concept of rational expectation, Samuelson (1971) showed that optimal and competitive storage would generate a nonlinear first-order Markov process for prices. Confronting for the first time the competitive storage model with the actual behavior of prices, Deaton and Laroque (1992 and 1996) were able to replicate significant price volatility and the skewness and kurtosis of actual prices for 13 commodities. The model was particularly successful in explaining sharp price spikes by explicitly recognizing the nonnegativity constraint of storage and thus building an essential nonlinearity into predicted commodity prices. The model, however, appears to be "incapable of generating the high degree of serial correlation of most commodity prices" (Deaton and Laroque(2003, p. 290). For Deaton and Laroque (1995, p. 28), the failure of the model did not reflect "inability to choose the right parameters in the simulations, but is a general feature of the model." This calls into question the model's empirical relevance and, more fundamentally, whether there is a model that can explain all four patterns of food commodity prices.

In this paper, we take a more inclusive approach in order to revive the empirical relevance of the competitive storage model. We propose an augmented model that integrates a range of supply and demand factors into one competitive storage model with rational expectations. Specifically, we augment the original Deaton and Laroque (1992 and 1996) model by introducing output and demand trends, shocks to the yield, and time-varying interest rates. By minimizing the distance between autocorrelations of simulated and actual prices, we obtain minimal distance estimates of 2.3 percent for the storage cost and 0.19 for the price elasticity of consumer demand, both consistent with estimates from other sources.

While it increases the computational burden exponentially, the augmented model succeeds in replicating all four patterns of food commodity prices. Our simulation and comparative statics exercises also show that (i) food prices declined from the 1970s through the late 1990s because output growth, largely driven by yield technology, had outpaced that of demand; but starting in the mid-1990s, the balance between supply and demand has been shifting with supply trend stagnating and demand trend accelerating; (ii) short-run price fluctuations can be attributed mainly to sizeable, though low-probability, shocks to output, such as significant weather-related shocks captured by the deviation of actual yields from expectations; and (iii) the impact of monetary policy, i.e., changes in real interest rates, is nonlinear and asymmetric and can become large when the real rate turns deep negative.

The findings have several policy implications:

1. Because a rapid surge in food commodity prices, as in 2007-08, is to a large extent the result of a temporary shock, it is likely to recur.

2. A moderate buildup of precautionary storage, as opposed to speculative storage by profit-maximizing investors, could be helpful.

3. The long-term decline in real food prices may come to a halt or even reverse now that signs of major structural changes underlying the supply and demand balance are emerging. Growth in the yield has recently slowed and yet the trend growth of demand continues, and perhaps accelerates. Increasing agriculture R&D and green revolution investments are thus needed to boost productivity. Removal of trade barriers could also help dampen price hikes during a crisis.

In what follows, Section 2 describes individual supply and demand factors that shape food commodity prices and motivate our holistic approach. Section 3 presents the theoretical model incorporating output, consumer demand, speculative demand, shocks to yields, and time-varying interest rates. In Section 4, we confront the theoretical model with observed wheat prices. We first apply a simulation-based minimum distance estimation and then identify the possible contribution of each factor via comparative statics. Section 5 summarizes the main findings.

2 Empirical Motivation: Factors That Affect Food Prices

Factors that could affect food commodity prices can be broadly classified as either supply or demand. Total supply for any given period, a.k.a. the "amount on hand," has two components: the new harvest and the storage from the previous period. While the amount of storage carried forward reflects some optimal decision from the previous period, the new harvest depends on the area planted and the current yield. The yield in turn reflects both technological advances and idiosyncratic factors, such as weather and fertilizer cost. Demand also has two components: consumer demand for current consumption and speculative storage demand driven by the real interest rate, the storage cost, and expected future prices. Using wheat as an example, we address each factor one by one before integrating them into one holistic model in Section 3.

2.1 Supply: Production and Storage

The annual global wheat harvest has consistently trended upward since the 1960s (Figure 1(b)); since the area planted has been fluctuating around 220 million hectares, this translates into a rising yield (Figure 1(c)), defined as total output divided by total area planted[3]. Because of the natural production lag, it seems appropriate to view area planted as the result of the optimal decision of farmers. Fluctuations in area planted likely reflect suppliers' trade-offs between revenue expected from growing wheat and the opportunity cost in terms of growing other crops or using the land for other purposes. The upward trend in yield mostly reflects technological advances, but there are also disturbances, sometimes large ones, to the actual yield. We interpret these as supply shocks, due perhaps to inclement weather or the cost of fertilizers.

To capture both the trend and shock components of the actual yield, we make the

[3] Appendix A contains the sources for data on wheat output, area, yield, stock, and etc.

following simplifying assumptions: (1) market participants both form their expectations on the distribution of next year's yield through a backwards ten-year moving-window detrending regression of log yields[4]; (2) suppliers at the beginning of each harvest year, make optimal decisions about areas to be planted, taking into account expected profits; such area decision is known to every market participant and no area surprises to the market throughout the harvest year[5]. Under the two simplifications, the yield next year will follow a log-normal distribution, with mean and variance defined by the recursive regressions at each point of time; the only uncertainty in our model is on future yield. The realized supply shocks could thus be defined as areas planted times deviations of realized from expected yields. Although the trend in yield is clearly upward, trend growth has slowed since the mid-1990s (Figure 1(c)). Large shocks to yields were also observed during the 1970s. Shocks were moderate in the 1980s but again became volatile after the mid-1990s (Figure 1(d))[6].

Storage as part of total supply also has two parts: precautionary and speculative. Public grain stocks are usually taken for precautionary purpose; they are relatively stable, barring revisions in regulations[7]. Speculative storage, however, is an endogenous decision incorporating expected price changes over time; hence it is much more volatile. Despite data constraints in separating the two motives, Figure 1(e), which includes both, suggests a fair amount of stock volatility. A common (mis)perception is that speculation drives up price and thus increases volatility. Our simulation shows that without storage food prices

[4]Expectation formation clearly depends on how far back market participants remember past harvests. Our choice of the 10-year horizon is a balance of two considerations: we want a horizon long enough to capture more than pure business cycle effects, but we do not want to lose too many observations, particularly those during the oil crisis of the 1970s.

[5]The assumption of no area surprises in the middle of a season seems reasonable. For example, market traders at the Chicago Board of Trade know actual areas planted well before the start of a trading season.

[6]Notice our assumptions imply Bayesian learning or adaptive expectations on the yield process, which in turn implies a time-varying distribution of future yield and harvest. This is different from the i.i.d. stochastic process assumed on future harvest by Deaton and Laroque (1992), which may be a little far away from that depicted in Figure 1(c) and 1(d).

[7]For example, when the U.S. loosened its official food reserve requirement in 1985, the U.S. wheat stock to output ratio plunged.

would become more volatile in both the original Deaton and Laroque (1992) model and our augmented model.

2.2 Demand: Consumption and Storage

We assume that consumer demand contains both a trend and a part with constant price elasticity; the trend part is usually considered to be related to both population and real income. From 1969 to 2008 world population grew by 80 percent and world real income by 230 percent, but global wheat output by only 120 percent. Here we choose to approximate the trend in food consumption demand by population[8]. Approximating the trend by real income would end up with an increasingly large demand-supply gap since 1985 as shown in Figure 1(f), which is not consistent with the declining real food prices shown in Figure 1(a). Besides the trend, the price dynamics also depend on the price elasticity of consumption, which we will assume to be stable throughout our sample period.

The other component is speculative demand, an endogenously determined variable, which is affected by real interest rates, storage cost, and market's expectations about the next period's price.

Interest rates affect food prices through the opportunity cost of holding storage[9], an inherent feature of the competitive storage model. We use the annualized real U.S. 3-month T-bill rate as a proxy for world real interest rates (Figure 2)[10]. An inspection of the figure identifies two periods of particularly large negative real interest rates, 1973-74 and 2003-04, which coincide with two large shocks to yield (Figure 1(d)) and two upticks in wheat prices (Figure 1(a)). While both supply shocks and loose monetary policies are

[8]The story could be different for raw industrial materials, demand for which might be more closely related to real income than population.

[9]The premia required to compensate the speculative risk and inflation risk may also be contributing factors. Compared to risk-free savings, investors will demand a premium to hold storage, the return on which is risky. Investors only observe nominal interest rates, but we make a strong though standard assumption that their ex ante inflation expectations are correct with respect to realized inflations.

[10]To be consistent with the price, here the interest rate is annualized in a corresponding international trade year.

often invoked as fueling commodity booms, a calibration and comparative statics exercise is necessary to decipher the extent to which each drives food price surges.

Similarly, storage cost affects the trade-off between immediate selling and holding stock. Since Gustafson (1958) it has been often modeled as a fixed marginal physical cost. Most recently, Cafiero et al. (2010) garnered some evidence from past studies that storage cost might be indeed constant over time[11]. Samuelson (1971) and Deaton and Laroque(1992, 1995, and 1996) in contrast specified it as a constant proportional decay of the numeraire commodity, which implies time-varying storage cost due to time-varying price dynamics, which in turn depends on the amount being stored across periods. In this paper we follow the route taken by Deaton and Laroque (1992).

2.3 Other Factors

Several other factors may also influence food commodity prices. For example, the exchange rate against the U.S. dollar is often cited because all major commodity transactions are settled in U.S. dollars. It is not rare to come across arguments that accuse the weak U.S. dollar of being the culprit for the 1974-75 and 2007-08 food and oil crises. At this stage and for tractability reasons, we leave these factors to further studies. We also assume a constant price elasticity of consumer demand. Its exact magnitude along with the decay rate is the free parameter to be estimated, and we shall examine how sensitive the results are to these assumptions.

[11] For details, see Cafiero et al. (2010) footnote 3 and references therein.

3 How Commodity Prices Are Determined

3.1 The Prototype Competitive Storage Model

How would all these factors—trends and shocks in yield, elasticity of consumer demand, cost of storage, and real interest rates—come together and shape food commodity prices? We start with a prototype model in Deaton and Laroque (1992) that incorporates competitive storage into the usual consumer supply and demand dynamics and introduces the concept of stationary rational expectations equilibrium (SREE). Specifically, total demand includes not only demand for current consumption—which depends only on the current price—but also speculative demand, which is a function of both current and the expected next period price[12]. Total supply is the sum of new crops and storage from the previous season. Given supply and demand, the market in each trading season functions as follows:

- There will be a consumer shadow price p_t^c, that equates the total supply with the demand for current consumption.

- Accordingly there is also a speculator shadow price function p_t^s, that renders speculators indifferent about whether to take stocks or stay idle. If the current price exceeds p_t^s, no one will take new storage, and vice versa. This shadow price is not a fixed number but a decreasing function of the total amount stored: given the output distribution and market price function next period, the less the stock-taking this period, the higher the expected price next period, which in turn requires a higher break-even shadow price p_t^s this period, other things being equal. But this inverse relationship between speculator shadow price and stock-taking cannot go on forever because there is a nonnegative constraint for storage. Thus, there is a maximal speculator shadow price denoted p_t^*, corresponding to zero stock-taking in period t.

[12]Notice that speculators do not need to form expectations on prices beyond the immediate next period because they can adjust their position freely then. For instance, if P_t is low, $\mathbb{E}_t[P_{t+1}]$ high, and $\mathbb{E}_t[P_{t+2}]$ low again, they will carry storage from t to $t+1$; but the storage decision at $t+1$ will be based on P_{t+1} and $\mathbb{E}_{t+1}[P_{t+2}]$ and has nothing to do with $\mathbb{E}_t[P_{t+1}]$ and $\mathbb{E}_t[P_{t+2}]$.

- Two types of market equilibrium exist and depend on the relationship between the consumer shadow price p_t^c, and the maximal speculator shadow price p_t^*: (i) if $p_t^c \geq p_t^*$, then no new storage will be taken and the market price will be solely determined by consumer demand; (ii) if $p_t^c \leq p_t^*$, then it will be profitable for speculators to take new stock until speculative demand drives the price up to the break-even point.

Let t denote the trading season, p_t the actual price, z_t the new harvest, I_t the inventory and x_t the supply available for this trading season. Assuming that storage is subject to a constant linear depreciation rate δ[13], we then have the accounting identity $x_t = z_t + (1-\delta)I_t$. Let's also denote consumer demand function as $D(p_t)$, inverse demand function $P(x_t)$, and market price function $f(x_t)$. We also assume that the cross-period real interest rate is r and that speculators are risk-neutral with the discount rate β being $1/(1+r)$.

Formally, the equilibrium price p_t can be defined as a nonarbitrage condition:

$$p_t = \text{Max}\Big[\beta(1-\delta)\mathbb{E}_t[p_{t+1}], P(x_t)\Big] \tag{1}$$

where

$$p_t = f_t(x_t) \tag{2}$$

$$p_{t+1} = f_{t+1}(x_{t+1}) \tag{3}$$

$$x_{t+1} = z_{t+1} + (1-\delta)I_t \tag{4}$$

$$I_t = \text{Max}[x_t - D(p_t), 0] \tag{5}$$

The key feature of this model is to view price (both today's and what is expected tomorrow) as a function instead of a specific number. The expected price next period is not a single number; it is explicitly expressed as a function of availability, which in turn is subject to the uncertainty of harvest yield. The problem is thus inherently recursive—

[13] The depreciation rate is best viewed as an average rate. Nonlinear storage cost could arise if there is warehouse or freight capacity issue. The depreciation also includes possible price discount; for example the leftovers may be priced at 10 percent less than the new crop.

for any given market price function at period $t+1$, it is possible to obtain the market price function at period t. Under certain conditions, i.e., storage is costly ($\beta(1-\delta)<1$) and new harvest z_t independently and identically distributed (c.f., Theorem 1 of Deaton and Laroque (1992)), the sequence of market price functions $\{f_{T-s}(x) : s = 0, 1, 2, \cdots\}$, where T represents a terminal period, will converge to a function, i.e., the SREE in Deaton and Laroque (1992)[14]. Notice that the SREE market price function is independent of the trading time, the central feature of the Deaton and Laroque (1992) model. Combined with the i.i.d. assumption on output process, the SREE implies that the only linkage between two consecutive years is the previous storage; everything else either has identical value or follows the identical stochastic process. So the generated price series is stationary and it makes sense to compute the statistics such as mean, variance, skewness, kurtosis, etc., just like one would do to a cross-sectional sample.

Solution to this model, however, is complicated by the uncertainty of future output and the nonnegative constraint of storage. Fortunately, numerical methods like the endogenous grid points algorithm could help deal with the two inconveniences and lessen the computational burden (for details, see Appendix B).

To get a sense of the convergence of the pricing function, we further assume that consumer demand takes the form of constant price elasticity, i.e., $D(p) = p^{-\rho}$, and that the output follows an i.i.d. log-normal distribution which will be relaxed in the augmented model. Baseline parameters based on the wheat price sample of 1969-2008 are presented in Table 1[15]. Using these baseline parameters, the convergence of the market price function is presented in Figure 3.

The kink point (x^*, p^*) in Figure 3 is the key to understand the price dynamics the model generates. As mentioned, p^* is the maximal shadow price for speculators to take

[14] In mathematical terms, the model defines a contraction mapping, and the SREE is the fixed point of the contraction mapping.

[15] The values for r and σ are the sample averages from 1969 to 2008; the values for δ and ρ are arbitrarily chosen for now as they are free parameters to be estimated by the model.

Table 1: Market Price Function Convergence: An Example

Calibrated Parameters		
Description	Parameter	Value
Real Interest Rate	r	1.1%
Discounting Rate	$\beta = 1/(1+r)$	0.989
Std Dev of Shocks to Log Output	σ	0.037
Mean of Shocks to Log Output	μ	0
Depreciation Rate	δ	0.05
Consumption Price Elasticity	ρ	0.5

new storage given their expectations of the next period price function, the opportunity cost of taking stock, and the depreciation rate of storage. When p_t^c is above p^*, all demand is for current consumption and the market price function is the inverse of the consumer demand function. Otherwise, the price is determined by both consumer and speculative demand.

In the presence of speculative storage, two features of the price dynamics are worth noting. Firstly, the asymmetry and peakedness of price movements are caused by the nonnegativity constraint. Positive output shocks will always be absorbed by speculators buying low and stocking this period and expecting to sell high next period. A severe negative output shock, however, may not be completely absorbed once storage is depleted. After stocking out, the price could rise even further and faster. The pace and magnitude of such an asymmetric price movement depends on the curvature of consumer demand and the distribution of the random output. Secondly, the price series generated by Deaton and Laroque (1992) model may not contain any trend, as exhibited by the low autocorrelation coefficients. In fact, if the only source of shock to price, i.e., harvest, is assumed to be stationary, the resulting price movements are also stationary.

3.2 The Augmented Model

The prototype competitive storage model, while generating the asymmetric price movements observed in the actual time series, fails to deliver high autocorrelations and the possible long-run declining trend of real prices. Fundamentally, by imposing various time-invariant assumptions, it does not capture the time-varying properties of supply and demand factors as outlined in Section 2. We augment the model in several ways:

Output process: Harvests z_t were assumed to be i.i.d. in the prototype model, but they really depend on three factors: area planted A_t, trending yield Y_t, and random shock z_t. We estimate the shock series from a 10-year recursive detrending model. With this augmentation, annual harvest $A_t Y_t z_t$ has not only random shocks that are assumed to follow time-varying log-normal distributions[16], but also a deterministic and constantly updating trend.

Consumer demand: We introduce trend growth into consumer demand as well, approximated by population growth and denoted as λ_t^D. Thus, the augmented consumer demand function is written as $\lambda_t^D D(p_t)$ and the inverse demand function as $P(x_t/\lambda_t^D)$. Without loss of generality, we also assume that the time-varying trend has no effect on the price elasticity of consumer demand.

Monetary policy shock: We abandon the constant real interest rate assumption made in the prototype model and update it with the actual U.S. interest rate series. When the real interest rate is lower (i.e., the opportunity cost of holding stock is smaller), more speculative capital flows into commodity markets, and speculators can afford taking storage at a higher price, which either helps sustain or pushes up the price.

Taken together, the equilibrium price is now given by the augmented nonarbitrage

[16] We do not consider an explicit distinction between a permanent and a transitory shock. In fact, because of the moving-window Bayesian updating, a negative shock today will affect the forecasting for tomorrow's trend. The shocks are in effect a mixture of both a permanent and a transitory type.

condition:

$$\begin{aligned}
p_t &= \mathrm{Max}\Big[\beta(1-\delta)\mathbb{E}_t[p_{t+1}], P(x_t/\lambda_t^D)\Big] && (6)\\
&\text{where} && \\
p_t &= f_t(x_t/\lambda_t^D) && (7)\\
p_{t+1} &= f_{t+1}(x_{t+1}/\lambda_{t+1}^D) && (8)\\
x_{t+1} &= A_{t+1}Y_{t+1}z_{t+1} + (1-\delta)I_t && (9)\\
I_t &= \mathrm{Max}\Big[x_t - \lambda_t^D D(p_t), 0\Big] && (10)
\end{aligned}$$

The purpose of the augmentation is to incorporate demand and supply factors both more accurately and more comprehensively. In the prototype model, price function converges to one particular form in the SREE. In the augmented model, however, there is a converged price function (SREE) for each period. This is because of the adaptive expectation on the stochastic yield process, and our period-by-period updating on other deterministic but time-varying parameters like demand trend. As output and demand evolve with time, so do market expectations. But at any given point of time and in the eyes of market participants, the expected price function for today and for the entire future will be the same and hence there is a SREE.

Our introduction of time-varying parameters, trends, and shocks has market participants constantly updating their beliefs on the next trading season in terms of, for instance, output and consumer demand. Each converged SREE is derived on the basis of the belief that the expected future prices functions will be the same as today's. However, we recognize that, because the market always adjusts to the most recent forecasts of demand and supply, the evolution of SREEs represents in effect adaptive expectations. This augmentation comes with the cost of an exponentially growing burden of computation. For example, in our sample we need to compute 40 price functions, each corresponding to one year and coming from a sequence of convergence, while the Deaton and Laroque (1992) model needs to compute only a single price function, which they assume will effectively determine the price in each sample year. The ensuing problem has to do with the memory of computer processors, in particular during the estimation process (For more details, see Appendix C).

4 Matching Theory And Data

The primary challenge of reviving the empirical relevance of the competitive storage model is to confront the theory with the data. Our estimation strategy is to first derive plausible values for free parameters in the model using method of simulated moments (MSM). Armed with these values, we then generate a food price series and compare it with the actual series. We also conduct comparative statics and shed light on the relative importance of individual factors in driving short-run and long-run price movements.

4.1 The Method of Simulated Moments Estimation

Two free parameters in our model are left for estimation: the price elasticity of consumption ρ and the storage depreciation rate δ. As is common in the literature, the estimation is to compare simulated with actual prices. Since it is not possible to write down an analytical form for the price functions, we estimate them by numerical methods. Specifically, for each candidate pair (ρ, δ), we would be able to simulate a sequence of realistically calibrated market price functions, and by plugging the actual shocks into these functions we would generate a series of artificial prices for 1969 to 2008. By comparing artificial with actual prices, we can pick the pair that best meets certain moment criteria, as is commonly done in the method of simulated moments literature, for example Gourinchas and Parker (2002) and Cagetti (2003).

Given that the primary focus is on the high autocorrelation of food commodity prices, one natural criterion would be to minimize the distance between simulated and actual autocorrelation coefficients with the objective function taking the form[17]:

$$\min \text{Gap}(\rho, \delta) = \sum_{i=1}^{i=3} \left(\text{AC}(i) - \widehat{\text{AC}}(i)\right)^2 \tag{11}$$

where $\text{AC}(i)$ is the i^{th} order of autocorrelation of actual prices and$\widehat{\text{AC}}(i)$ that of simulated prices, which in turn are generated by a particular (ρ, δ) pair. The pair that minimizes the objective function in Eq. (11) constitutes our MSM estimate.

Our estimation exercise yields a constant decay rate of 2.3 percent and a price elasticity of consumption of about 0.19 over a possible range of parameter values (Appendix C). The estimate of a low food consumption elasticity is in line with those published by the Economic Research Service at the U.S. Department of Agriculture. The average price elasticity for breads and cereals, the subgroup most closely related to wheat, is 0.30, and for advanced and emerging market economies it is normally within the range of 0.10 and 0.30[18]. The low decay rate is also comparable to the converged estimates in Cafiero et al. (2010)[19]. The corresponding value of the minimized distance is about 0.008, which implies that the maximal deviation of any of the first three order autocorrelations from their counterparts of the actual data will be about 0.09. As an analog to standard error estimation, if we restrict the values of the objective function to be within 0.01, i.e., a maximal deviation of 0.1 in autocorrelation, the estimated decay rate would vary between 1.8 and 2.8 percent and the elasticity between 0.12 and 0.22. The whole area in Figure 9(f) represents pairs that could achieve a gap smaller than 0.01.

[17]Another convenience is that we adopt autocorrelations as the sample moments here, which like the coefficient of variation and skewness, are independent of any measurement units. This avoids the problem of conversion since the actual prices are measured by US$ per metric ton, while the simulated ones are just indexes.

[18]The ERS at USDA provides estimated price elasticity of consumption for nine major consumption groups and eight food subgroups across 114 countries; more details are accessible at its `homepage` (clickable).

[19]Anecdotal evidence from public warehouses also suggests that the annual storage cost should be around 2-5 percent depending on wheat prices. In a survey of country elevators, Kenkel (2008) estimated that the annual variable cost (including moisture and shrinkage, electricity, and fumigation) of storing one bushel of wheat was 0.119 in 2005 and 0.180 in 2008. Given the wheat price of 3.42 and 7.50 per bushel in 2005 and 2008, this implicates an annual depreciation rate of 3.5 and 2.4 percent respectively.

Adopting the MSM estimator here has at least two benefits. First, this method is particularly convenient in addressing cases where sample moments are complex non-linear functions of the parameters of interest. Second, the objective function as defined offers an intuitive appeal that enables one to assess directly how a particular pair of parameters performs in terms of matching the first three orders of autocorrelations not only in relative order but also in absolute magnitude.

The empirical strategy here is different from Deaton and Laroque (1992). They adopt the GMM framework, where moments conditions are derived from the efficient market hypothesis. However, two obstacles prevent us from using this framework. Let's denote the actual sample prices as $\{p_t : t = 1, 2, \cdots\}$; solving the model at time t will result in two indexed series: $\{p_t^{\text{Sim}} : t = 1, 2, \cdots\}$, the simulated series by plugging actual shocks into the model, and $\{p_{t+1}^{\text{Pred}} : s = 1, 2, \cdots\}$, the one-period-ahead prediction series. The first problem is related to p_{t+1}^{Pred}: similar as in the prototype model we assume that at any point people will form rational expectations about future price functions, but now we also assume that they will update their knowledge on the distribution of future yield based on ongoing yield history, and hence there will be an adaptive expectation. Under this setting, our simulation will end up with p_{t+1}^{Pred} being very similar to p_t^{Sim}. Therefore, assuming that $\{p_t^{\text{Sim}} : t = 1, 2, \cdots\}$ matches well with $\{p_t : t = 1, 2, \cdots\}$, the forecasting error series $\{\mu_{t+1} = p_{t+1} - p_{t+1}^{\text{Pred}} : t = 1, 2, \cdots\}$ will be quite close to $\{p_{t+1} - p_t : t = 1, 2, \cdots\}$[20]. But the latter can not be a mean-zero series, given the apparently asymmetric price movement in the data. This voids the moment conditions from the efficient market hypothesis[21]. The second problem is the conversion between the real price US$ per metric ton and our simulated index series. In theory there are various conversion options, such as making the mean or the beginning point of two series equal, but none will be particularly helpful here.

[20] This is because p_{t+1}^{Pred} resembles p_t^{Sim}, and the latter by assumption is very close to p_t.

[21] The Deaton and Laroque (1992) model, on the other hand, generates a time-invariant forecasting series: p_{t+1}^{Pred} stays the same for $t = 1, 2, \cdots$, which by choosing the forecasted price as the mean of the sample prices, could help ensure the $\{\mu_{t+1} : t = 1, 2, \cdots\}$ is a mean-zero series. However the implication that the market holds a constant prediction all the time from 1968 to 2008 is not realistic.

Because of the badly behaved forecasting errors and the associated moments conditions, the usual GMM framework ends up with strange point estimates, e.g., the decay rate degenerating to negative. To put this more intuitively, the issue here is not whether one is careful or not in measurement, but that the ruler itself, i.e., the moment conditions in the prototype GMM exercise, is not correctly calibrated.

In practice, parameters estimated by the GMM approach lack stability. Using the same equations and data but only with finer grid points, Cafiero et al. (2010) derived much smaller and more plausible decay than did Deaton and Laroque(1992 and 1996). In addition, they also showed that the estimated constant decay rate varies with the number of grid points being used. Specifically, increasingly finer grid points in their model leads to decreasing estimates of the constant decay rate. In some cases, e.g., sugar, the decay rate degenerates to zero and in other cases, e.g., cotton and cocoa, the rate converges to around 0.05 at around 1000 or finer grid points.

Equipped with the MSM estimates of the two free parameters, we can now compare simulated with actual prices (Figure 4); and the summary statistics are presented in column "Actual" and "IV" in Table 2. Our model succeeds in generating several patterns of actual price time series: 1) the declining long-run trend of real prices, including a slight reverse of the trend after the mid-1990s; 2) large variations, with the ups-and-downs in the two time series in general matching each other, although with varying magnitude; 3) high autocorrelations—the feature seriously missing in Deaton and Laroque(1992 and 1996) but robustly replicated in our model mainly because of trending output and demand; 4) asymmetric price movements with close to 1 skewness and positive excess kurtosis. The model does not generate extreme price hikes with a magnitude similar to that of 1973-74, however. This is reflected in the reduced excess kurtosis: 0.7 for the simulated and 2.6 for the actual price series.

4.2 Comparative Statics

To separate the impact of each contributing factor, we now conduct comparative statics analysis. We plug in our new estimates for (ρ, δ), allow one factor to be time-varying as observed in the actual data, and fix other factors at their baseline value in Table 1. The impact of each factor is presented in figures recording both the evolution of market price functions and the evolution of (x^*, p^*), the threshold point that corresponds to investors' break-even stance in a market price function. The horizontal axis represents per-capita availability (supply) and is normalized to 1 for the simulation starting at year 1969; the vertical axis represents prices with the price in base year 1969 likewise normalized to 1.

4.2.1 The Effect of Output Trend

Other things being equal an upward output trend as observed in the data implies downward pressure on prices. This is confirmed in Figure 5(a) and 5(b), where p^* falls as output grows in the past four decades. Specifically, the magnitude of price changes is in proportion to that of output growth: real prices were significantly reduced when output showed good progress between 1970 and 1990 but were only marginally lowered when output stagnated after the mid-1990s. To facilitate the understanding of price dynamics, we also plot p^* on the left axis and x^* on the right axis, both against the output trend index in Figure 5(c). With the growth of per capita output, the market price function becomes increasingly relaxed: investors will not start taking stock until total supply reaches increasingly high levels (an increasing x^*) or until current prices hit close to bottom (a decreasing p^*).

4.2.2 The Effect of Demand Trend

This exercise singles out the effect of consumer demand trend over the past four decades (Figure 6(a)). The qualitative impact is as expected: market price functions were continuously pushed upward by growing demand. For any given supply, rising price functions imply increasing prices and more stringent market conditions: p^* keeps increasing and as a result, the market price also keeps rising for a given market supply.

The quantitative impact of output and demand trend is dominating: the sheer magnitude of their impact dwarfs the effect of all other factors combined. Everything else being equal, demand trend alone would cause p^*s, the price at the corresponding stock-taking points, to rise from 0.8 in 1970 to 24 in 2008, a thirty fold price run-up in less than four decades. Similarly, supply trend, if left alone, would have kept the 2008 price a paltry one fortieth of the 1970 price (p^* of 0.017 in 2008 and around 0.70 in 1970). The two balancing forces have so far kept food prices in check. Unlike the supply trend which is decelerating, however, demand trend is actually accelerating. Our simulation documents an increasingly large impact on price of consumer demand and yet a decreasing impact of supply trend over the past four decades. This result is a reflection of the shift of the underlying balance between supply and demand.

4.2.3 The Effect of Yield Shocks

In addition to a predictable trend, the yield is subject to unexpected shocks, assumed to be log-normally distributed and derived from 10-year recursive regressions. Compared to the long-run trend, shocks to yields are of less magnitude (see Figure 1(c) and 1(d)) and thus expected to have relatively less impact. Figure 6(b) compares two extreme cases: a maximal yield shock that was forecasted at 1978 for the year 1979, a minimal shock forecasted at 2007 for the year 2008. It shows that other things equal, a higher risk in the yield process will induce a higher price function[22]. This result is intuitive because when yields are subject to a larger shock, the expected next period price will go up due to the convexity of market price functions; hence speculative demand will begin at a higher

[22]The surge in food price also has something to do with precautionary storage by both governments and cautious consumers. When consumers realize that the output may be subject to worse than expected supply shocks, they will take storage to insure against such risks. This additional demand may push prices even higher.

current price level, indicating a tighter current market condition.

4.2.4 The Effect of Interest Rates

Cheap money is often deemed to fuel speculative behavior. The U.S. real interest rate indeed varied considerably between 1969 and 2008, from as low as a negative 3.84 percent (1974) to as high as a positive 5.19 percent (1983). The price impact that can be attributed to pure interest rate movements, however, is small in normal times especially if compared to that of demand and output trends. For reasonable cross-year interest rate variation, the magnitude is also much smaller than that of yield shocks. However, its effect may become disproportionately large if real interest rate passes a certain threshold. Our estimate of the threshold lies at -2.3 percent, and under our estimated cost of taking storage, the condition $\beta(1-\delta) < 1$ holds. If the real rate goes further negative than that, stocking taking will become a one-way bet and price run away[23].

Out of the forty years under study, only in one trading year did the real interest rate go beyond the threshold. That year happened to be 1973-74, a year with -3.84 percent of real interest[24], a price hike of 75 percent, the largest one in our sample, and one of the largest negative yield shocks. Our simulation exercise shows that two-thirds of the price hike may be explained by the negative real interest rate passing the threshold. When the real rate moves from -1.35 percent (1973) to -3.84 percent (1974), price will rise by around 50 percent (Figure 6(c)). However, the effect would have been much smaller if it did not pass the threshold. For example, if the real interests rate is reduced from around zero (2008) to around the threshold level (1979), the price at the corresponding stock-taking

[23] In mathematical terms, this is equivalent to that there is no convergence for the sequence of market price functions; they diverge instead. Hence there is no SRRE for this scenario.

[24] In Figure 6(c), 1974 was labeled as "-3.84%", and the displayed function is the one after 50 periods' of iteration. If more periods are allowed, the (x^*, p^*) point will move left-upwards further. It will eventually diverges to an infinite price regardless of the total supply. This is because at lower than -2.3% real interest rate, speculative demand becomes infinite which induces higher-than-ever prices. Of course our 2.3% depreciation estimate is roughly an average cost over the sample years. If more and more storage is taken, the marginal storage cost will keep rising as well, due to various capacity constraints in processing, warehousing, etc. So eventually the rising depreciation rate will put a brake on explosive stock-takings.

point changes by less than 10 percent[25].

Other aspects of our simulation confirm the nonlinear and asymmetric effect of the real interest rate. When the real interest rate plunges into negative territory and moves further beyond the threshold, the marginal impact is increasingly large; on the other hand, as long as the real interest rate remains positive or close to zero, its impact on prices is almost negligible. Take the year 1983 (5.19 percent) and 2008 (-0.026 percent) as an example: more than 500 basis points reduction in the real interest rate only leads to around 5 percent rise in the price at stock taking point[26]. In normal times, the effect of interest rates could be even smaller as they often move more smoothly and remain positive.

4.2.5 The Effect of Depreciation Rates

We have assumed, as in the literature, a constant decay rate. With the upgrade of the food storage and distribution system to major grocery chain stores and increasing bulk purchases by consumers, there is a likelihood that storage cost may go down. In the model, reduced storage cost will induce more storage taking and thus increased price smoothing, other things being equal. To confirm this in the simulation, storage cost is allowed to vary between 0 and 15 percent. Two effects are immediate: First, an increasingly lessened depreciation would induce investors to take stock increasingly earlier as the intervention point x^* becomes smaller (Figure 7(a)). Second and more central to our model prediction, reduced storage cost indeed smoothes prices. The market price function curve becomes smoother with reduced storage decay. Consider for example the Magenta line (0 storage cost) and the Purple line (15 percent storage cost) in Figure 7(a), the overall price process under the Magenta line is much smoother than that under the Purple line, indicating less

[25] The real interest rate in 1979 happened to be -2.31 percent, which enables us to examine the threshold effect via simulation.

[26] A caveat in interpreting the result: the only thing this exercise reveals is that, without any trend in supply and demand, the SRRE in the prototype model will change little when the real interest rate plummets from 5% to 0%, other things being fixed as in Table 1. This is different from the exercise of looking at the effect of changing interest rate in 2008 from 5% to 0% while fixing everything else at their 2008 values including supply and demand trend.

dramatic price movements when the per capita availability varies.

Overall, the effect of changing storage cost is not as large as other factors already considered, unless there is a dramatic industrial innovation.

4.2.6 The Effect of Consumption Elasticity

The price elasticity of food demand is a key determinant of the computed market price; and Figure 7(b) shows the sensitivity of the SREE with respect to this parameter. For more inelastic demand (lower ρ), both consumer demand and market price will be more convex, the hikes and plunges of prices more frequent, and the threshold of stock taking x^* lower. This is quite different from the more elastic case ($\rho = 0.50$) that would imply a less violent price movement.

4.3 Overall Effects

Figure 8(a) presents the overall effect of combined output and demand trends, yield shocks, and real interest rates to the benchmark model using the newly estimated (δ, ρ) and other time-varying parameters from actual data. For the five snapshot years, 1970, 1980, 1990, 2000, and 2008, by examining market price functions and the associated (x^*, p^*) we find that 1) from 1970 to 1980 and 1990, market conditions continued to improve: per capita output grew and price dropped; 2) from the mid-1990s to 2008 market conditions deteriorated and the 2008 market condition was only marginally more relaxed than that of 2000. The tightening of the market condition since the mid-1990s, in our view, is best characterized by constant and perhaps accelerating growth in consumer demand and intrepid growth in output. That shift of balance between demand and supply, coupled with negative interest rate shock and inelastic demand, may have produced the 2007-08 food crisis. The market environment in 2008-2009 became somewhat less stringent because there was a positive yield shock in 2008 after negative ones in 2006 and 2007. There is, however, a risk of rising prices in the medium to long term because we are still in the middle of the low output/demand ratio regime observed since the mid-1990s.

It is worth highlighting the dominant impact of supply-demand (im)balance, measured by the trend output/demand ratio, on long-run price movements. Figure 8(b) presents this story more vividly: we show the series of x^* against the per capita output/demand ratio. Note that x^* is the threshold level of supply beyond which speculative investors will enter the market. A higher x^* with a corresponding lower p^* thus implies a more relaxed market condition. It is clear that output/demand ratio closely resembles the threshold per capita supply, which in turn is the mirror image of the threshold price series. Hence over the long-run the price trend is dominated by per capita output, the key determinant for which is the output/demand ratio. In the short run, however, factors such as yield shocks will also affect per capita output, and there is some small divergence between output/demand ratio and threshold per capita output.

A final point may be worth mentioning: the storyline we presented here closely resembles the permanent income hypothesis first advocated by Friedman (1957). A vast literature has since shown that the effect of permanent income on consumer behavior is far more important than transitory income shocks and risks in future interest rates. Our model in essence also tells a story of consumption and saving; the difference is that the decision maker here is not a microeconomic agent but the "invisible hand" in a centralized market. In fact the nonarbitrage condition (Eq. (1)) appears quite similar to the Euler equation in precautionary savings studies (c.f., Carroll (2004)). It is thus not surprising that the output and demand trend, i.e., the trend of per capita output, speaks much louder than any other factor in the comparative statics exercise. Just like permanent income plays a key role in determining optimal consumption, here the output/demand ratio, becomes the predominant factor in price movements. This is also why the augmented model could solve the question left open by Deaton and Laroque (1992), who acknowledged that their generic model is not capable of generating the autocorrelations seen in actual prices.

4.4 The Role of Storage

With the augmented model, we have successfully matched the four patterns of food prices. However, one question remains open—what role, if any, does storage play in driving the volatility and serial autocorrelations of food prices? To address this question, a prior question needs to be answered first, namely, whether the competitive storage model would still be relevant if per capita supply as augmented contains a trend. In particular, if the trend is upward people would expect a lower price in the future, which in turn would make speculators much less interested in stock-holding today. In other words, does our augmentation mute the voice of speculative storage?

Our answer is no. Storage-taking still happens because on a year-to-year basis, trends move little and other factors such as yield shock and interest rate volatility dominate. In some cases the market would still expect a higher price next period, which makes storage attractive. The fact that a speculator's planning horizon is just one season ahead means that the long-term trend may not affect the optimal storage decision in the current period.

What is the exact role of storage then? In the model, competitive storage is endogenously determined by risk neutral speculators. It is demand in this period and yet supply in the next, an inherent dynamic feature of the model. Contrary to popular belief, Deaton and Laroque (1996) echoed the findings of many scholars that speculators can smooth prices by buying cheap and selling dear. They also reached the conclusion that speculative storage can substantially increase autocorrelation for prices but not to the high level observed in the data.

To shed some light on these issues, we analyze and compare five settings in the order of increasing complexity and proximity to the reality: (I) The barebone version: without storage and without trends in output and consumer demand; the only shock is an i.i.d. disturbance to output; and interest rate is held constant; (II) Barebone plus trend: adding to the barebone version trends to output and consumer demand; (III) Deaton and Laroque (1992) version: including storage to the barebone model but with neither trends in output and demand nor time-varying interest rates; (IV) Augmented model: integrating storage, trends in output and demand, and time-varying interest rates into the model; (V) Augmented plus oil: a possible ad hoc oil price adjustment[27]. For each scenario, we feed into the model the same estimated output shock, mostly reflecting weather shocks to the yield, and compare the summary statistics based on the simulated and the actual observed prices. The result is shown in Table 2.

In the barebone setting (model I) serial correlations are very small, which is expected since there is neither autocorrelated shocks nor intertemporal storage. There is little asymmetric price movement as well with the skewness and kurtosis being the smallest among the five settings. Adding a trend (model II) significantly increases autocorrelations, although there is still some distance to matching those observed in the actual data. Including storage (model III) as in Deaton and Laroque (1992) leads to a better match for the first order autocorrelation as well as the skewness measure. But the second or higher order autocorrelations remain small. This is because competitive storage provides a link between availability in this and the immediate next period, but not for outer periods, given that shocks in this model has no persistence at all. And the only comparison relevant for storage decisions are prices between today's and the immediate next period's.

Our augmentation (model IV) leads to the best overall fit with the actual data. Neither trend nor storage alone can explain high serial correlations in the actual prices. In addition, the natural nonnegative constraint, which is inherent in the storage decision, plays the most important role in explaining asymmetric price movements. But without trend and time-varying interest rate, the overall price volatility can be significantly underestimated as in both the barebone and the Deaton and Laroque (1992) settings: the coefficient of variation

[27]The heuristic assumption we make here is a complete pass-through from crude oil price shocks to food prices via its effect on fertilizer and transportation costs. For instance, if oil price jumps (plunges) by 50 percent in a given year, the model simulated food prices will be adjusted upward (downward) by 50 percent accordingly. The full pass-through may seem large if one looks at transportation only, but would be reasonable if we take into account the magnifying effect from crude oil to retail gasoline, increasing food production cost, as well as higher-than-expected inflation.

have been kept artificially low.

Our augmented model, however, can not reproduce the peakedness in the data, e.g., the huge price hikes in a crisis period. We conjecture that other factors not captured in the augmented model may be responsible. The most natural candidate is the oil price, the addition of which (model V) helps improves the fit of excess kurtosis and skewness, at the cost of slightly reduced autocorrelations. The pass-through of oil price shocks to food commodity prices help capture such egregious price movements as in1973-74 and 2007-08. Figure 4 displays simulated prices for the model II, IV and V[28], allowing for a more vivid comparison of them.

4.5 Other Food Commodities

Wheat is a representative of the general food commodity group, and the supply and demand characteristics of maize, rice, and soybeans are similar to those of wheat, as shown in Figure 10. Rice is most similar to wheat: the output/demand ratio is increasing while the real price is falling. Maize price also behaves as expected, except that there may be an additional effect from demand for bio-fuel after the 2000s. Soybeans are slightly different: its output growth clearly outpaces that of GDP and population, which raises the question whether there is a good proxy for the growth in its demand. Soybean output accelerated significantly after the mid-1990s. Anecdotal evidence suggests that it is in part due to the derived demand of emerging market consumers upgrading from staple food into meats and especially pork. Soybeans are the primary feed for pork and other live stock.

[28]For the ease of comparison, we omit model I and III; the simulated prices for both models will be stationary and look like a very flat line as compared to Figure 4.

Table 2: The Role of Storage and the Goodness-of-Fit in Five Settings[a]

Variables	Statistics	Actual[b]	I[c]	II[c]	III[c]	IV[c]	V[c]
Prices	AutoCorr(1)	0.829	0.0276	0.486	0.3179	0.7618	0.6056
	AutoCorr(2)	0.6337	0.1369	0.5339	0.0904	0.6216	0.467
	AutoCorr(3)	0.4945	-0.1364	0.4187	-0.0997	0.5507	0.4323
	CoVariation	0.4928	0.2281	0.3727	0.1736	0.3281	0.516
	Skewness	1.5679	0.786	1.1086	2.2967	1.0258	1.4039
	Excess kurtosis	2.6036	0.9482	0.9988	6.1632	0.7105	2.6063
Storage	Mean	N.A.	0	0	6.9632	9.6005	9.6005[d]
	Median	N.A.	0	0	6.1931	7.7559	7.7559[d]
	Minimum	N.A.	0	0	0.0193	0.0012	0.0012[d]
	Maximum	N.A.	0	0	74.648	31.2925	31.2925[d]

[a] This table compares the summary statistics of the simulated prices and storages from five different models with the actual data. The simulation utilizes the estimated depreciation rate (2.3%) and price elasticity (0.19). For other parameters, we use the fixed values from Table 1 or the time-varying counterparts of Table 1 where applicable to the underlying model.

[b] This is based on actual wheat price data.

[c] Model I-V are respectively the barebone version without trend and storage, the barebone version with trend, the barebone version with storage (i.e., Deaton and Laroque (1992) model), the barebone version with trend and with storage (i.e., our augmented model), the augmented model with ad hoc oil price adjustment.

[d] Model V assumes an ad hoc adjustment for the simulated prices and no effect on storage.

5 Conclusion

We attempt in this paper to present a holistic view of food commodity prices by integrating a number of factors such as output and demand trend, yield shocks and real interest rates, depreciation cost, and the elasticity of demand into one rational expectation model for competitive storage. We augment Deaton and Laroque (1992) prototype model with more realistic and time-varying factors, solve it by constantly updating knowledge about these factors as time evolves, and arrive at a SREE in each period. The evolution of the market price functions incorporates the evolution of output, demand, storage, and monetary policy, which all play their expected roles in our simulations and comparative statics exercises.

Our most significant findings are the predominant role of the output/demand ratio, which must be combined with the intertemoral storage, in explaining long-run food commodity price movements and the high autocorrelations observed in actual prices. Short-run price movements are mostly due not to shocks with a large variance but rather to the realization of small-probability events, such as a yield shock larger than two standard deviations. Yield shock distributions as calibrated from the data are quite stable and have small variances. The abrupt short-run fluctuations of food commodity prices are caused by small-probability events, such as severe drought or oil price shocks. Monetary policy plays a limited role in normal times but could have nonlinear and significant impact when the real rate becomes deep negative.

We acknowledge that a few factors that are not reflected in our model may help explain the occasional large price hikes. One might be oil/energy prices that affect the cost of food production through their impact on fertilizers, pesticides, and transportation, as well as unexpected inflation. Another might be the deterioration of world trade during food crises. For tractability reasons, our model assumes that all global output will be available for global consumption. With food prices soaring high, to maintain stable domestic supply, many countries worldwide have imposed bans on exporting certain staple foods, or raised their regulatory restrictions, which eventually could only push the prices higher. Thirdly, the model does not account for substitution among different food commodities.

Finally, though it might be tempting to apply our augmented model to energy prices, this cannot be done in a straightforward way because we made use of several characteristics that are peculiar to food commodities. Compared to oil prices, food staples might be unique in at least two aspects. The first is the lower and fairly stable price elasticity of consumer demand because staple foods are necessary to human nutrition and few substitutes are available. Secondly, it is relatively easy to identify an instrument to model the suppliers' response to food prices. In this paper we simply approximate the supplier's optimal decision by a simple variable observable to both the market participants and the econometricians: area planted; however, similar concise instrument cannot be found for oil and industrial raw materials[29]. Its use allows us to avoid two complicated elements of modeling: (1) how suppliers will optimally respond to prices, and (2) how their optimal response will affect current prices. In our model, the area decision for $t+1$ is known by the market at t, and it will affect the current price through the indirect channel of total supply at time $t+1$ and hence speculative demand at time t rather than directly by changing time t total supply. Given these specificities, more careful modeling strategy is needed to analyze oil prices using an augmented speculative storage model.

References

Cafiero, C., E. S. Bobenrieth, J. R. Bobenrieth, and B. D. Wright (2010): "The Empirical Relevance of the Competitive Storage Model," *Journal of Econometrics*, In Press, Corrected Proof.

[29]The impact of the OPEC on oil prices cannot be underestimated, but there is no similar food commodity cartel. Thailand wanted to launch the OREC (Organization for Rice Exporting Countries) in 2008, but the plan did not materialize.

CAGETTI, M. (2003): "Wealth Accumulation Over the Life Cycle and Precautionary Savings," *Journal of Business and Economic Statistics*, 21, 339–353.

CARROLL, C. D. (2004): "Theoretical Foundations of Buffer Stock Saving," *NBER Working Paper*.

——— (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91, 312–320.

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

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

DEATON, A. AND G. LAROQUE (1992): "On the Behaviour of Commodity Prices," *Review of Economic Studies*, 59, 1–23.

——— (1995): "Estimating a Nonlinear Rational Expectations Commodity Price Model with Unobservable State Variables," *Journal of Applied Econometrics*, 10, 9–40.

——— (1996): "Competitive Storage and Commodity Price Dynamics," *Journal of Political Economy*, 104, 896–923.

——— (2003): "A Model of Commodity Prices after Sir Arthur Lewis," *Journal of Development Economics*, 71, 289–310.

FRIEDMAN, M. A. (1957): *A Theory of the Consumption Function*, Princeton University Press.

GOURINCHAS, P.-O. AND J. PARKER (2002): "Consumption Over the Life Cycle," *Econometrica*, 70, 47–89.

Gustafson, R. L. (1958): "Implications of Recent Research on Optimal Storage Rules," *Journal of Farm Economics*, 40, 290–200.

Kenkel, P. (2008): "Grain Handling and Storage Costs in Country Elevators," *Working Paper at Department of Agricultural Economics, Oklahoma State University*, available at http://agecon.okstate.edu/coops/files/GrainHandlingandStorageCostsinCountryElevators.pdf.

Lewis, A. (1954): "Economic Development with Unlimited Supplies of Labor," *Manchester School*, 22, 139–191.

Muth, J. F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29, 315–335.

Prebisch, R. (1959): "International Trade and Payments in an Era of Coexistence: Commercial Policy in the Underdeveloped Countries," *American Economic Review*, 49, 251–273.

Samuelson, P. (1971): "Stochastic Speculative Price," *Proceedings of the National Academy of Sciences*, 68, 335–337.

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

Appendices

A Data

Table 3: Data Sources

Variables	Sources
Nominal Commodity Prices	IMF, `Commodity Price System` (clickable)
Nominal Interest Rate	US Federal Reserve Board, `H15 Releases` (clickable)
US CPI	US Bureau of Labor Statistics, `CPI Table` (clickable)
Global Wheat Production	USDA, `Wheat Year Book` (clickable)
Other Food Production	UN Food and Agriculture Organization, `FAOStat` (clickable)
World Population	US Census Bureau, `International Database` (clickable)
World Real GDP	USDA, `Economic Research Service` (clickable)

There is a convention on annual food commodity prices. Usually spot price data in the Chicago Board of Trade market are available at monthly frequency, but for annualized food prices, one has to differentiate an international trade year from a calendar year. According to the U.S. Department of Agriculture, July 1 approximates the wheat harvest in many northern hemisphere countries; and the international trade year is in accordance with the time when new harvests arrive in the US market. We calculate the annual price as the average monthly price over the 12 months in an international trade year, which for wheat is July 1-June 30.

Therefore, the "2008 food crisis" refers to the price hikes in responses to low harvests in 2006 and 2007 harvest years. Notice the negative supply shock in each year were not particularly large; however, the back-to-back nature of consecutive bad harvests may imply a very low level of storage before the 2008 new crops. Since July 2008 when the good harvest due to a positive supply shock gradually reached the market, the world saw a decline in the prices of staple from their highs between late 2006 and early 2008[30].

B The Endogenous Grid Points Algorithm

The equilibrium price function is solved by backward iteration until convergence using the endogenous gridpoints algorithm as in Carroll (2006). To make the model tractable, we first relax the non-negative storage constraint. The reduced formula governing the price process is

$$f_t(x_t) \quad = \quad \beta(1-\delta)\mathbf{E}_t\big[f_{t+1}\big(z_{t+1}+(1-\delta)(x_t-D(f_t(x_t)))\big)\big] \tag{12}$$

For each function $f_{t+1}(x)$, the formula implicitly defines $f_t(x)$, which in turn could be used to define $f_{t-1}(x)$, and so on. The theoretical results of Deaton and Laroque (1992) show that eventually the sequence of price functions will converge, which is the stationary rational expectations equilibrium we are interested in. We start with a terminal condition[31] that $f_T(x) = P(x)$, and then we can conduct backward iterations to generate a sequence $\{f_t(x) : t = T, T-1, T-2, ...\}$. However, how fast the convergence will take place depends on the numerical solution.

One natural approach is to (1) approximate the stochastic z_{t+1} by some discreet probability distribution so that the expectation could be computed as a summation of values at these discreet points, as opposed to the excruciatingly slow numerical integration procedures; (2) choose some gridpoints for x_t and calculate the corresponding optimal value of f_t; and then (3) interpolate around these gridpoints to obtain the numerical function

[30]Some would argue that it is due to the global economic slowdown since September 2008. However it is expected that consumption for staple foods should have little income elasticity. A relevant factor could be the liquidity shortage of many traders during the crisis, which constrains their speculative capacity.

[31]We could assume that the horizon $T-t$ is long enough and at this terminal period T (e.g. the end of the universe), nobody will care about a period further and thus storage demand will be 0.

$f_t(x)$. However, the second step involves a root-finding algorithm, which in general is time-intensive. A faster way is to replace step (2) by choosing some grid points for the storage amount I_t, and calculating the corresponding value of f_t, with which one could recover corresponding x_t as $I_t + f_{t+1}^{-1}(f_t)$; since the computer does not need invoke a root-finding procedure, this algorithm as in Carroll (2006) is much more efficient.

This also helps in dealing with the nonnegative constraint, which essentially means the smallest grid point of I_t is zero. If we define the values of x_t and f_t that correspond to $I_t = 0$ as x_t^* and f_t^*, then $\forall x_t \leq x_t^*$ (equivalently $p_t \geq p_t^*$), the market price will be governed by consumer demand alone, and there is no need to adopt the numerical methods any more.

Therefore, the market price function on the domain to the left of x_t^* will be the same as consumer demand, and that to the right of x_t^* will be an interpolated price function that depends on a combination of speculative storage and consumer demand. To ensure convergence we will iterate the steps many times. The archives associated with the text contain a programming folder that describes the implementation details in the `Mathematica` software.

C Method of Simulated Moments

At the beginning the candidate range for (δ, ρ) is the unit square $[0, 1] \times [0, 1]$. We divide this square into 100 grid points, compute our objective function for each point, and display the result in a contour plot: the x-axis is the candidate for the depreciation rate, y-axis is for the elasticity, and the point corresponding to each pair is labeled using color: the darker the region, the smaller the value of the objective function. Our task is to spot the darkest point in this square, which will be our MSM estimates.

There are several rounds of trials, so we narrow our search range gradually. As seen in Figure 9, we eventually arrive at an estimate of depreciation, 2.3%, and elasticity, 0.19, at which point the gap is 0.008.

This minimum-finding process cannot be done automatically by numerical optimizers. Hence step-by-step brute force grid-search is necessary in the search for the pair that best matches the simulated with the actual data. In fact, because of memory constraints, the computational kernel will need to be shut down after a round of trials to free up the memory; and because of the need to compute a series of converged price functions at each trial, the computation takes non-trivial time. With the various tricks in the endogenous grid points algorithm, on a Intel Core-2-Quad processor with 4GB RAM, a trial for a particular pair of parameters will take 3-5 minutes. Overall this process will require about eight hours.

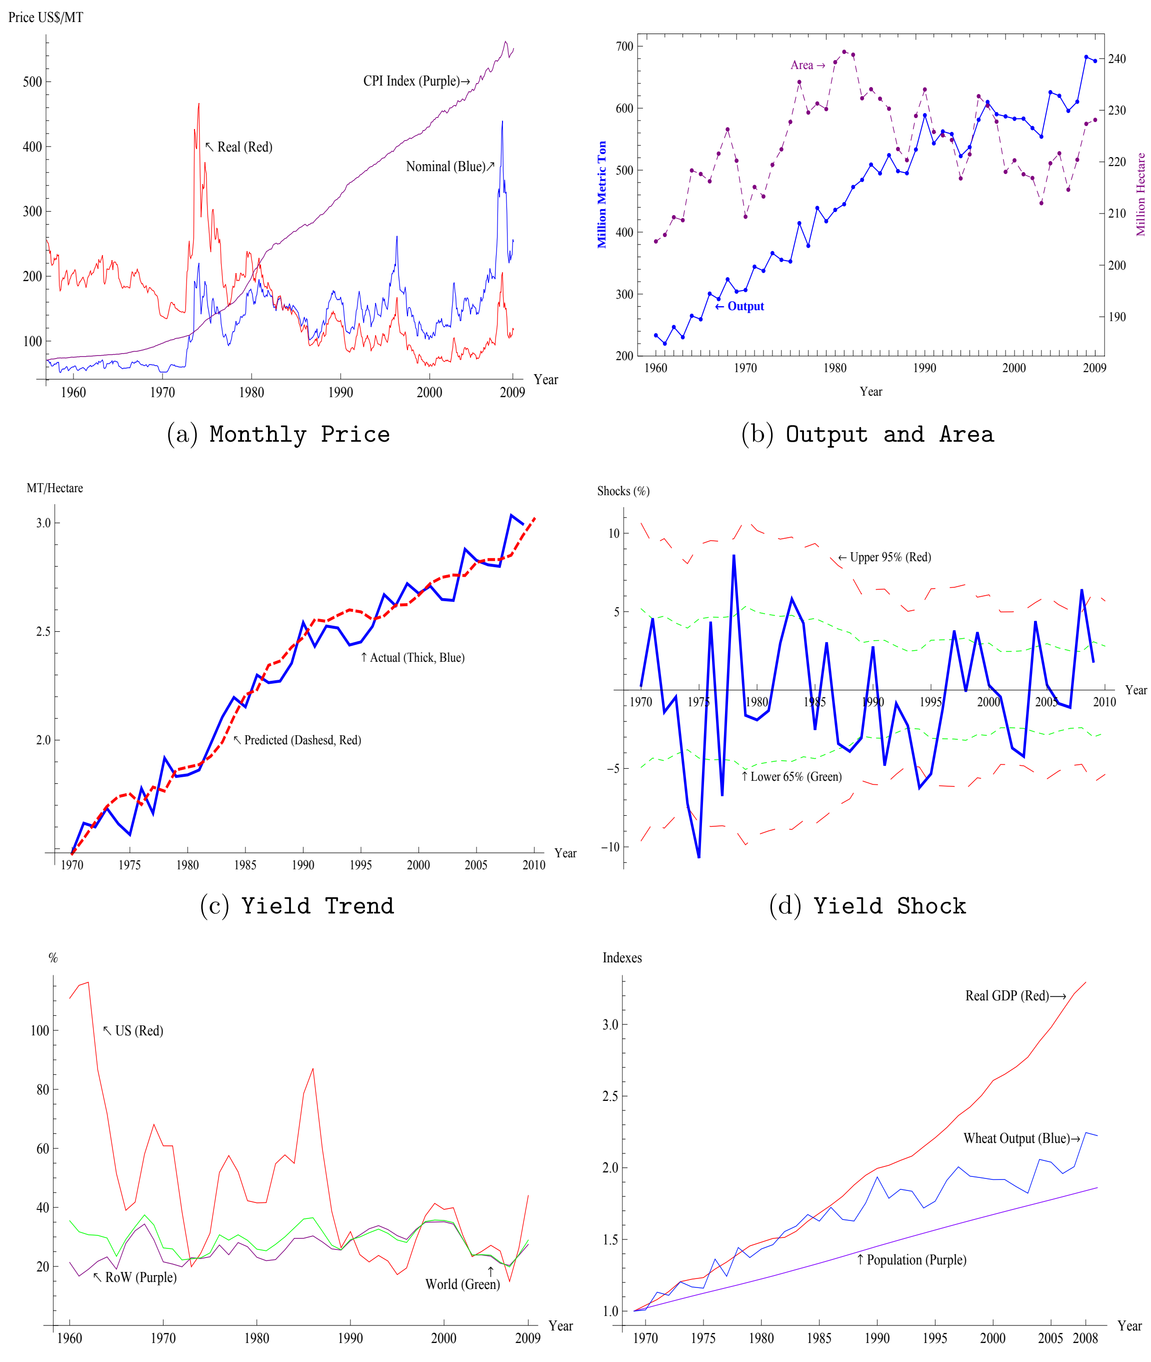

(a) Monthly Price

(b) Output and Area

(c) Yield Trend

(d) Yield Shock

(e) Stock/Output

(f) GDP and Population

Figure 1: Global Wheat Commodity Market: 1960-2009

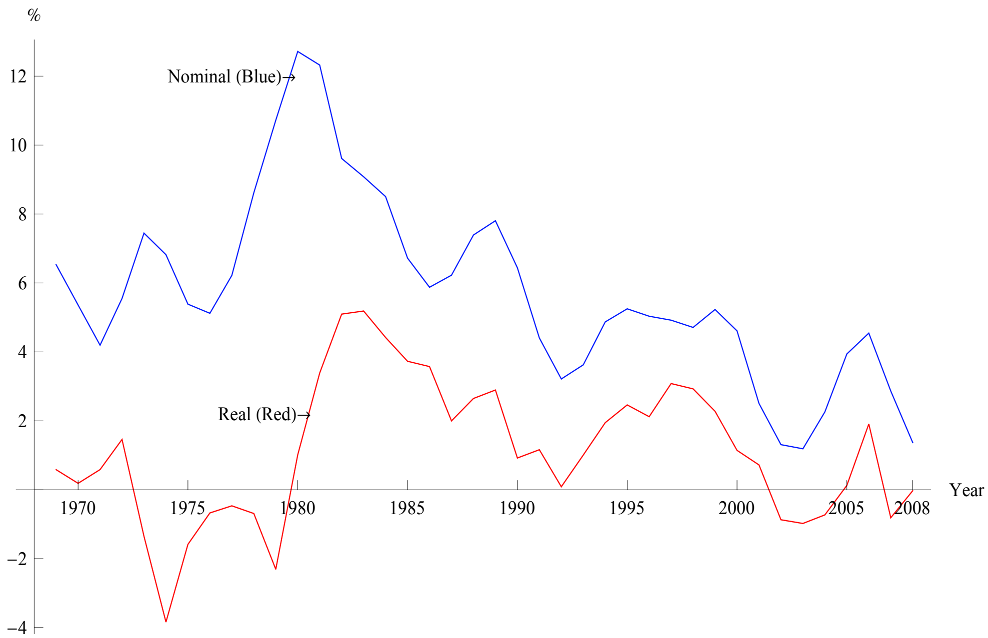

Figure 2: U.S. 3-Month T-Bill Annualized Interest Rate: 1969-2008

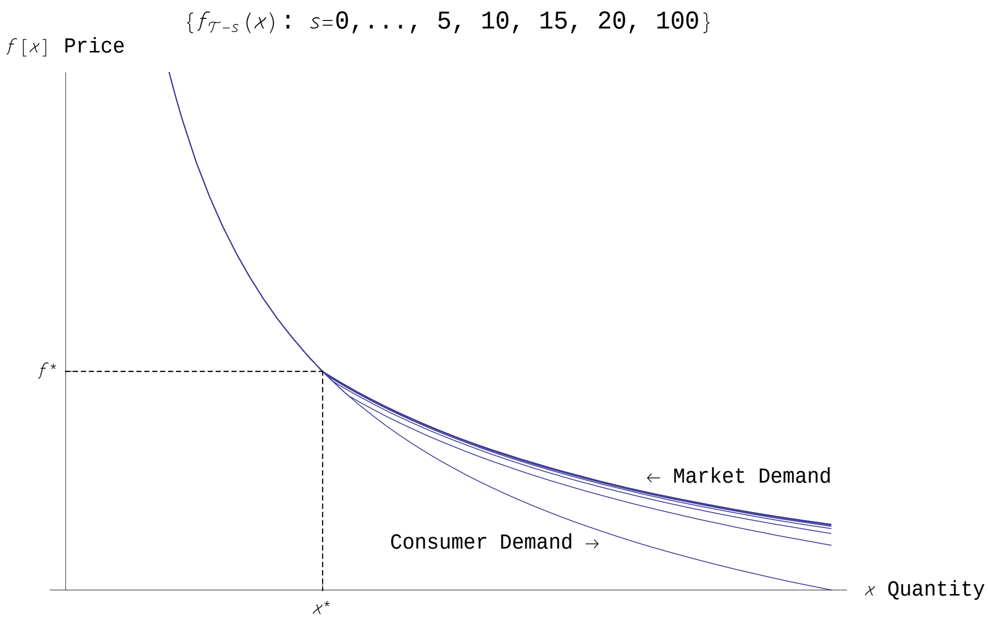

Figure 3: Convergence of Market Price Functions: An Example

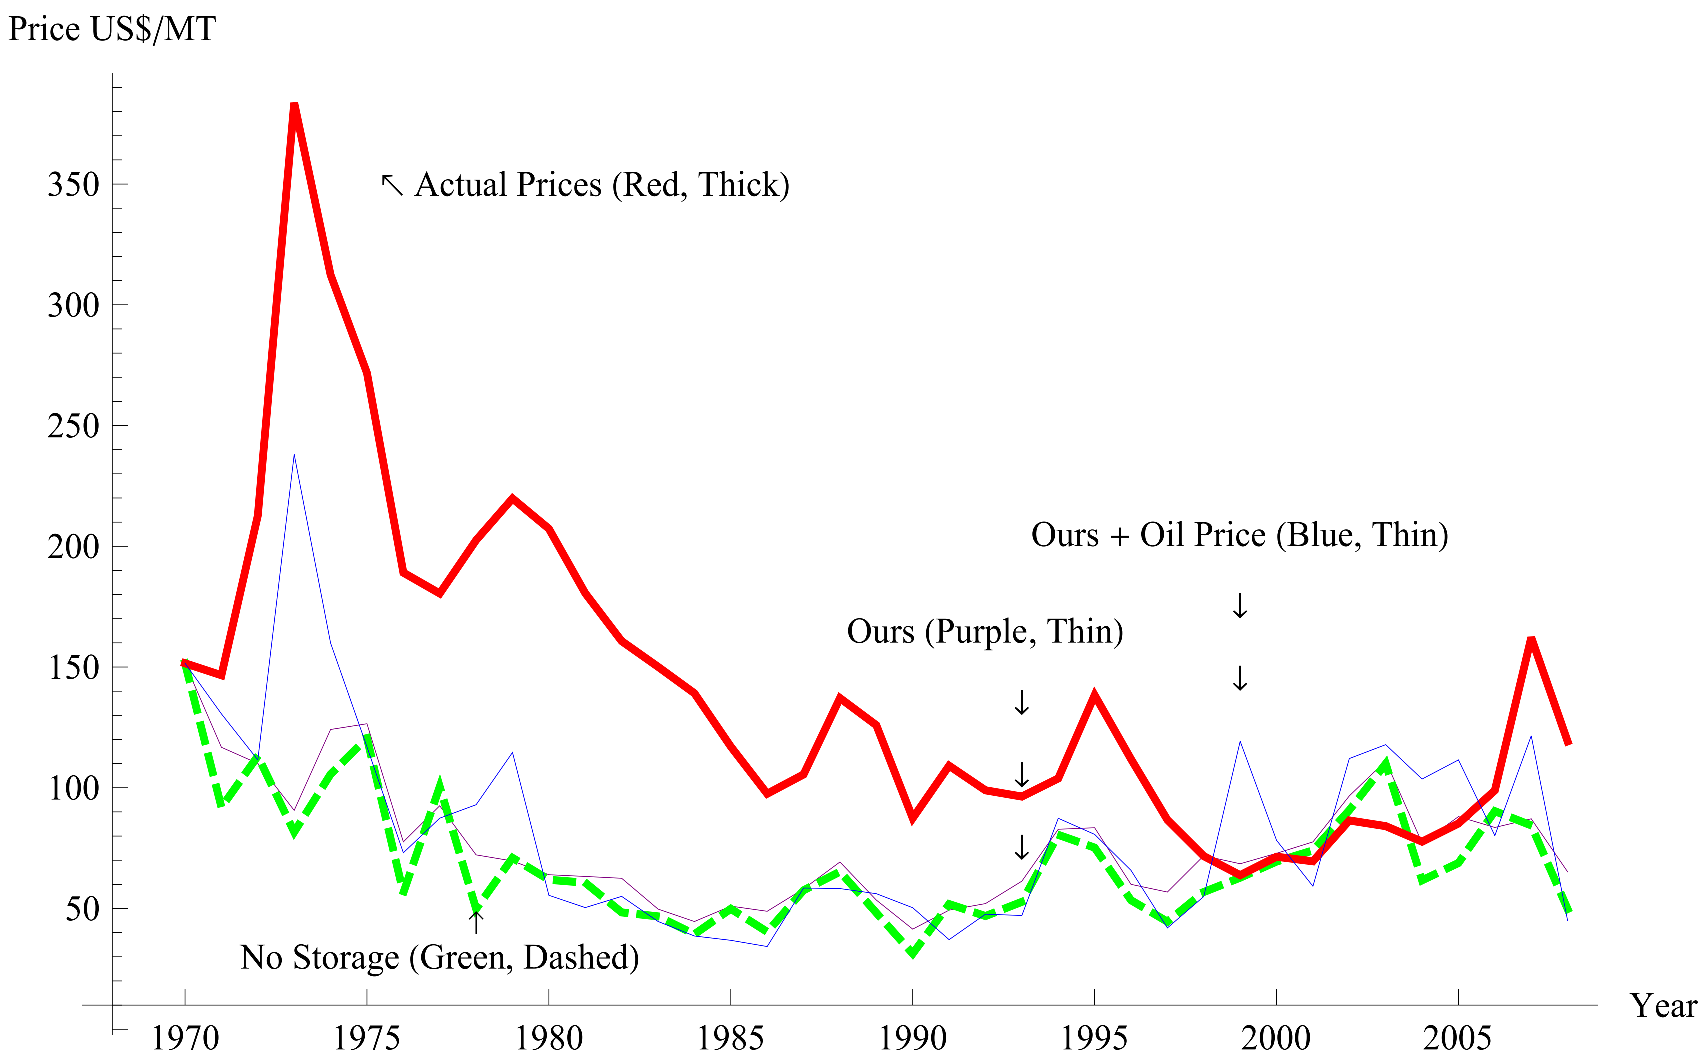

Figure 4: Global Wheat Commodity Market: Actual vs. Simulated

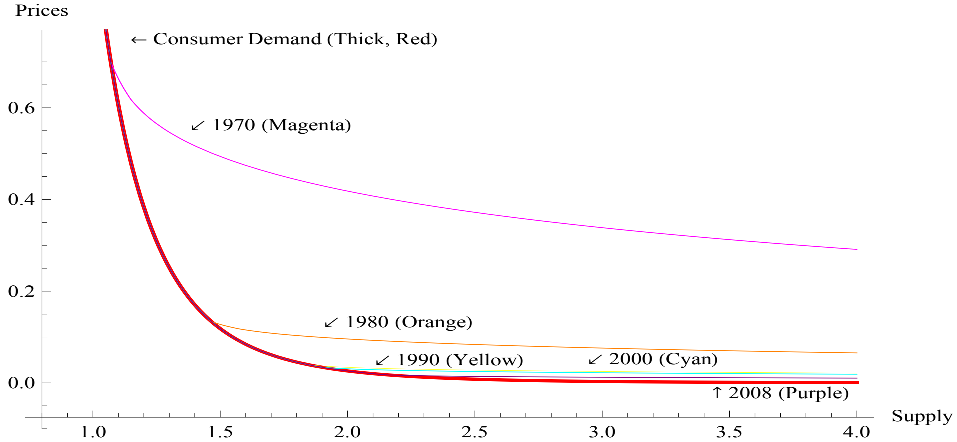

(a) 1969-2008

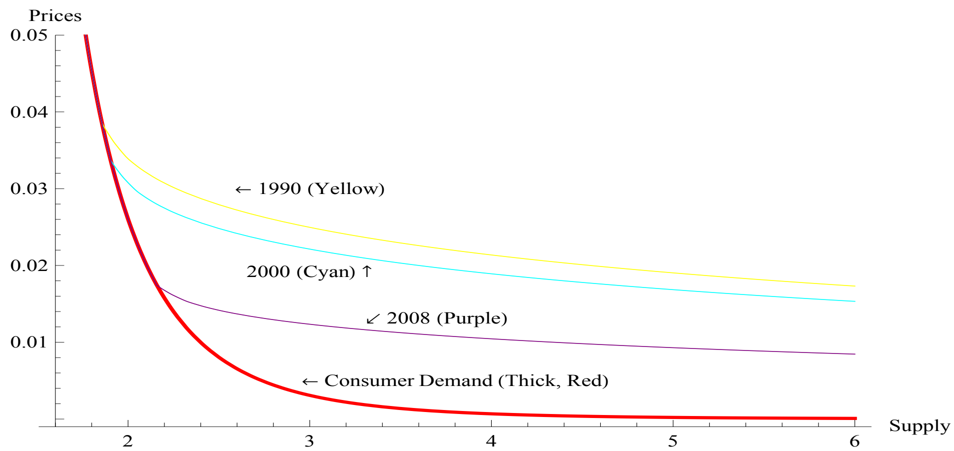

(b) 1990-2008

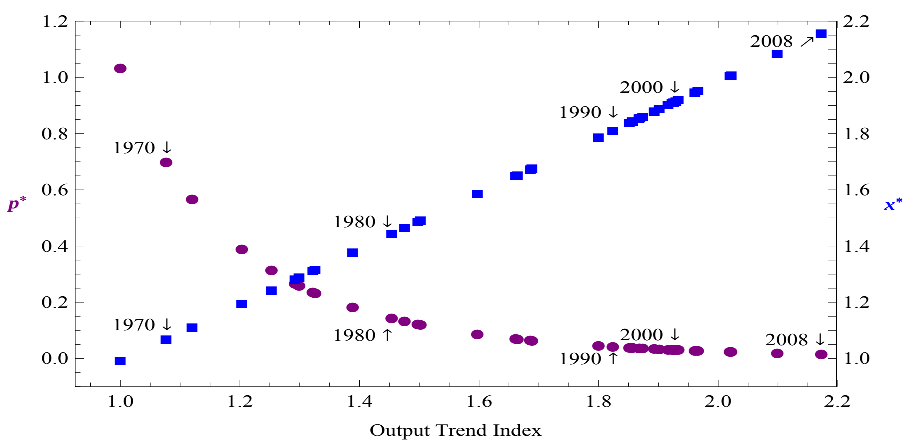

(c) Output Trend and p^*

Figure 5: Simulated Wheat Market Price Functions: 1970-2008 (Output Trend Only)

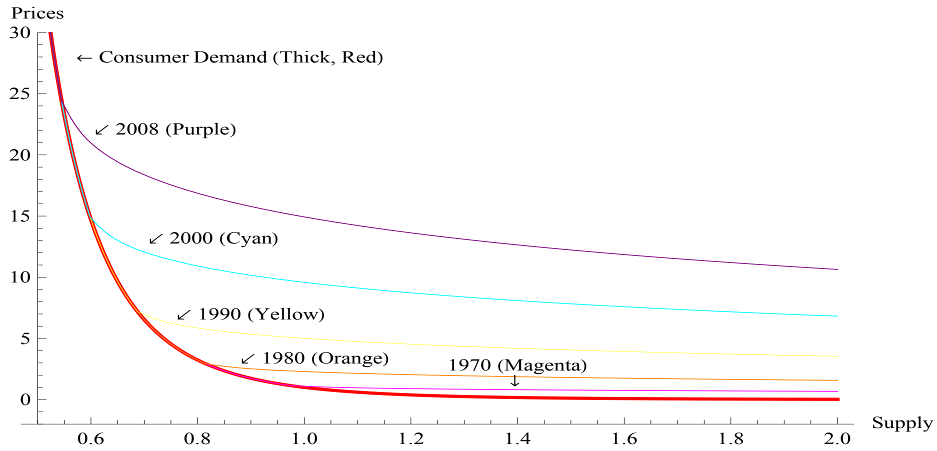

(a) `Demand Trend Only`

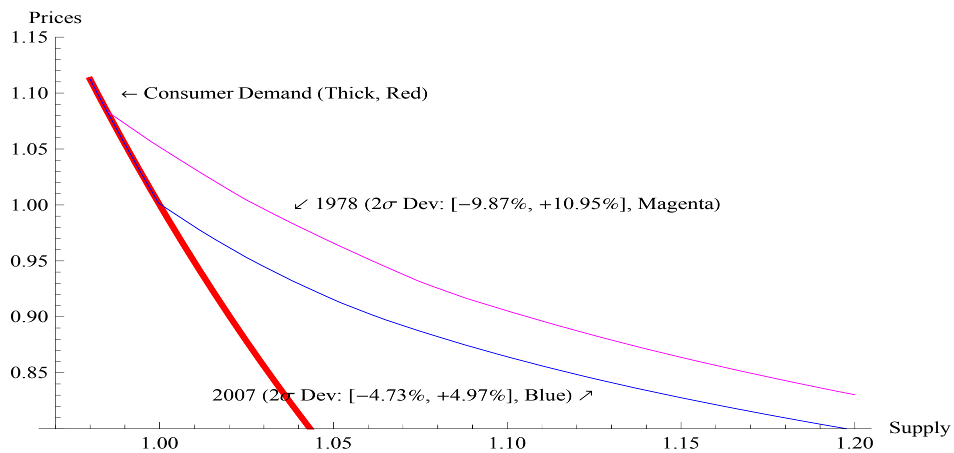

(b) `Yield Shock Only`

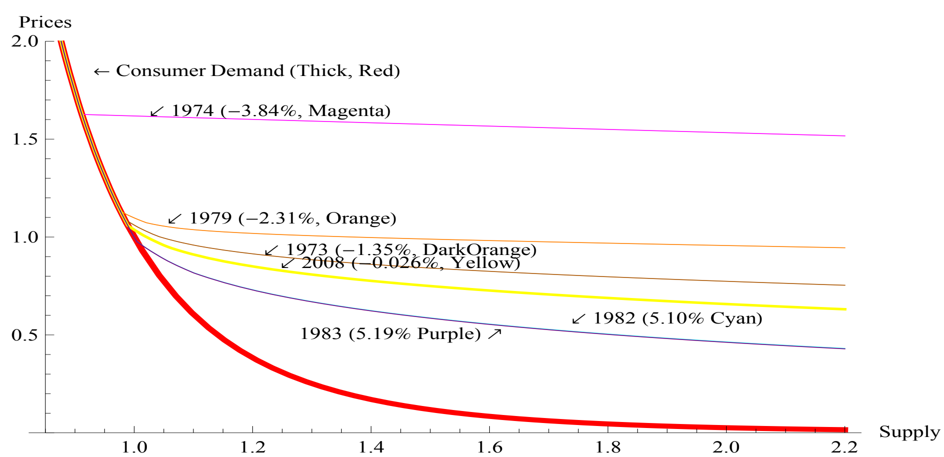

(c) `Interest Shock Only`

Figure 6: Simulated Wheat Market Price Functions: 1970-2008 (Other Shocks)

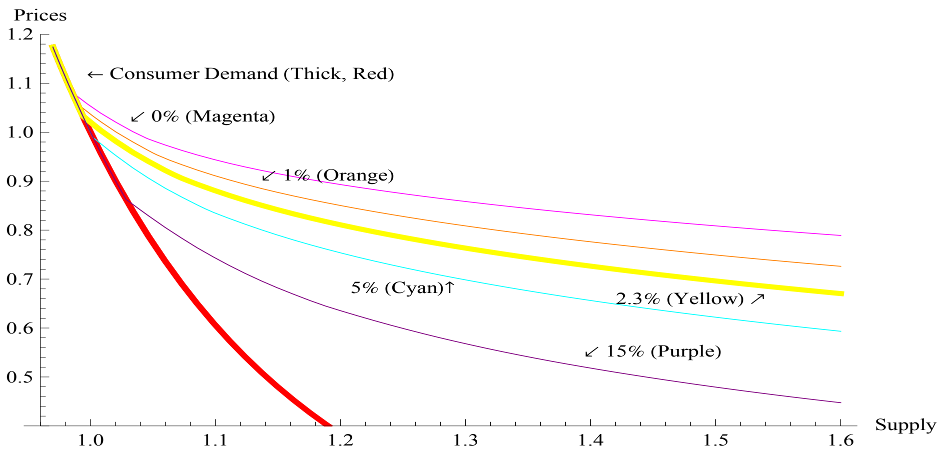

(a) `Depreciation Shock Only`

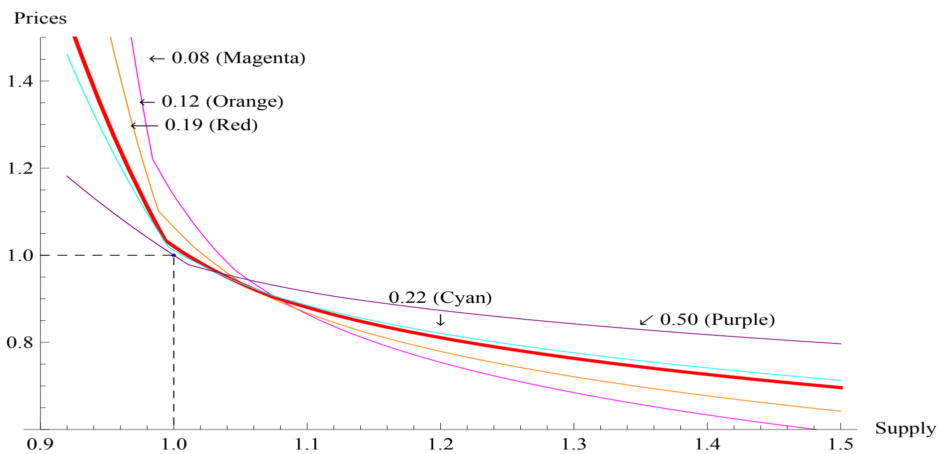

(b) `Elasticity Shock Only`

Figure 7: Simulated Wheat Market Price Functions: 1970-2008 (Other Shocks)

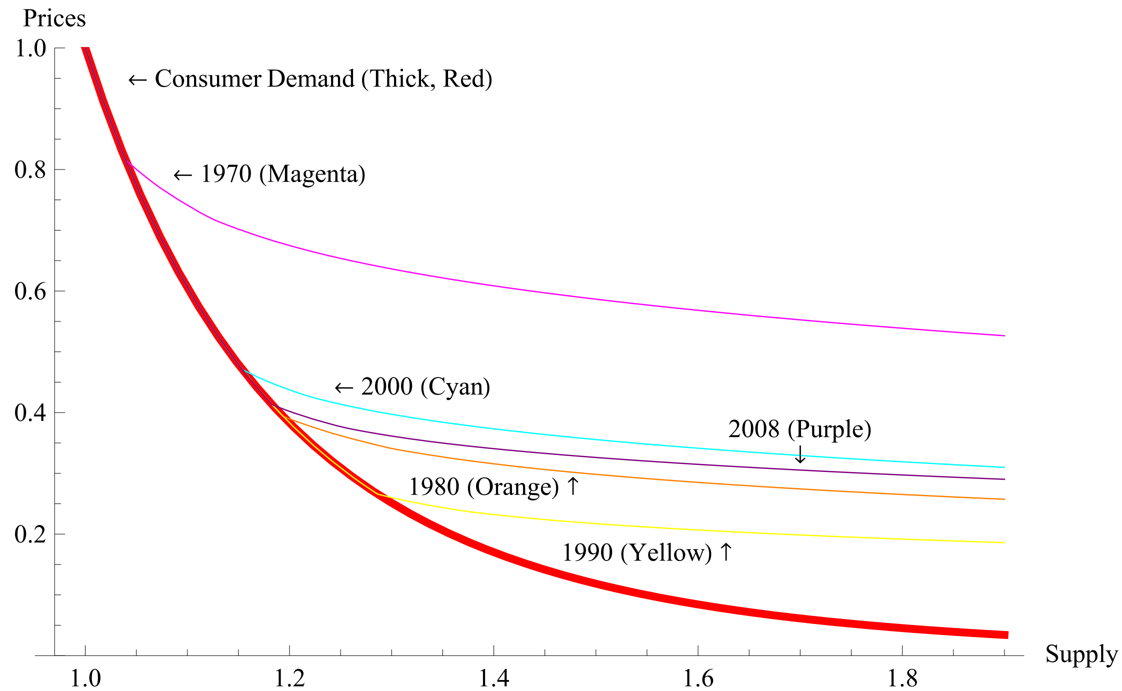

(a) Price Functions

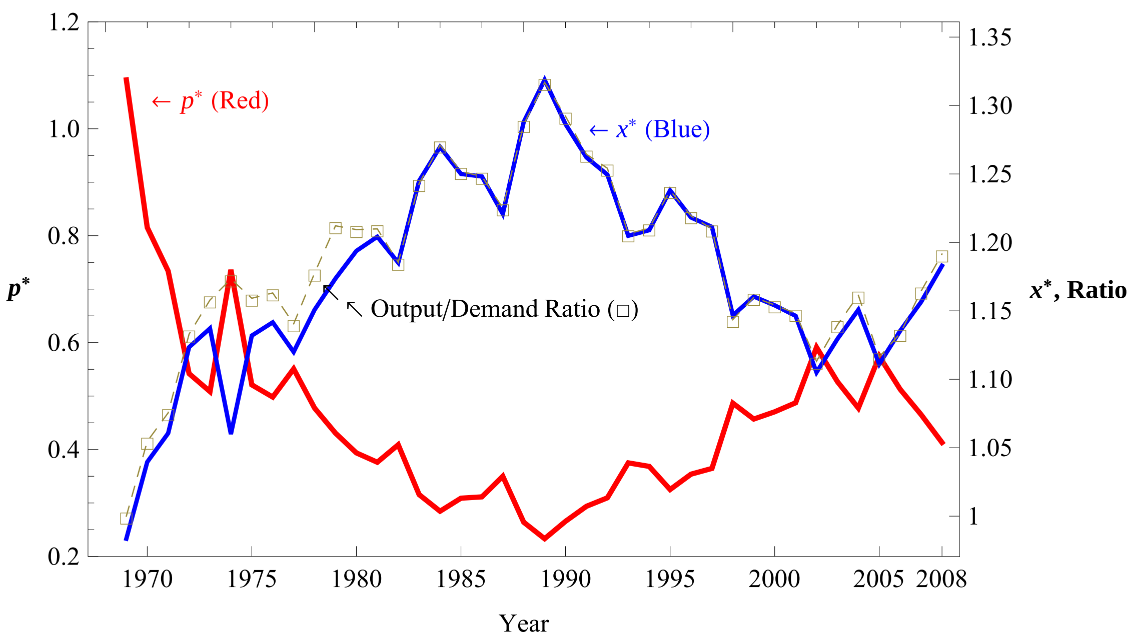

(b) Output/Demand and (x^*, p^*)

Figure 8: Simulated Wheat Market Price Functions: 1970-2008 (Overall Effects)

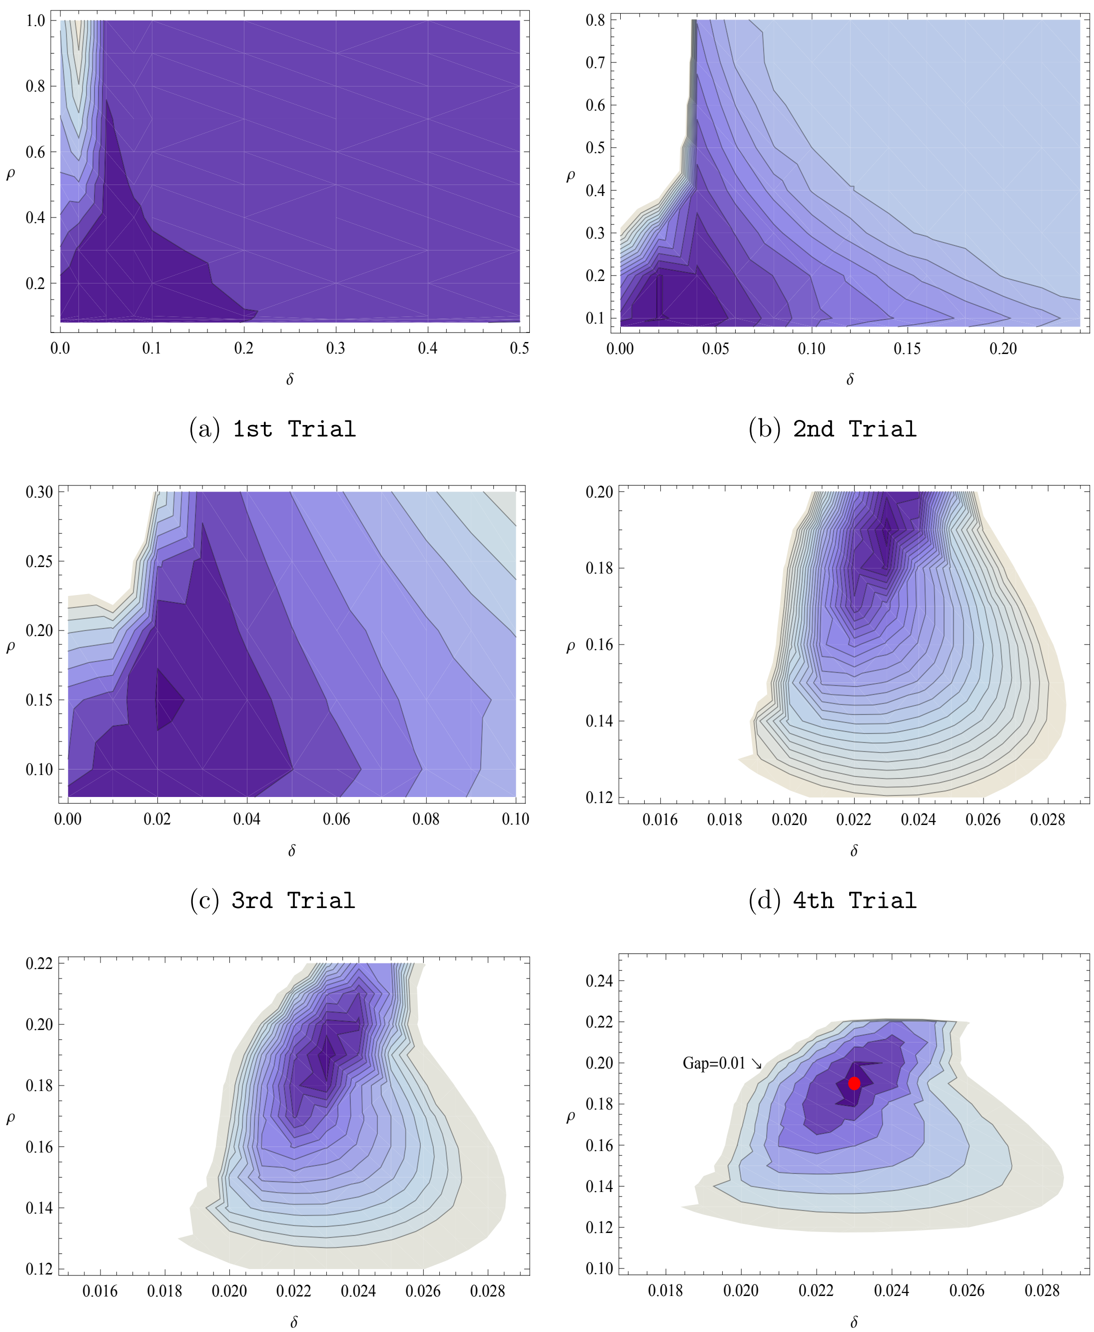

(a) `1st Trial` (b) `2nd Trial`

(c) `3rd Trial` (d) `4th Trial`

(e) `5th Trial` (f) `Final with Std. Err.`

Figure 9: Minimum Distance Estimation of Decay Rate δ and Price Elasticity ρ

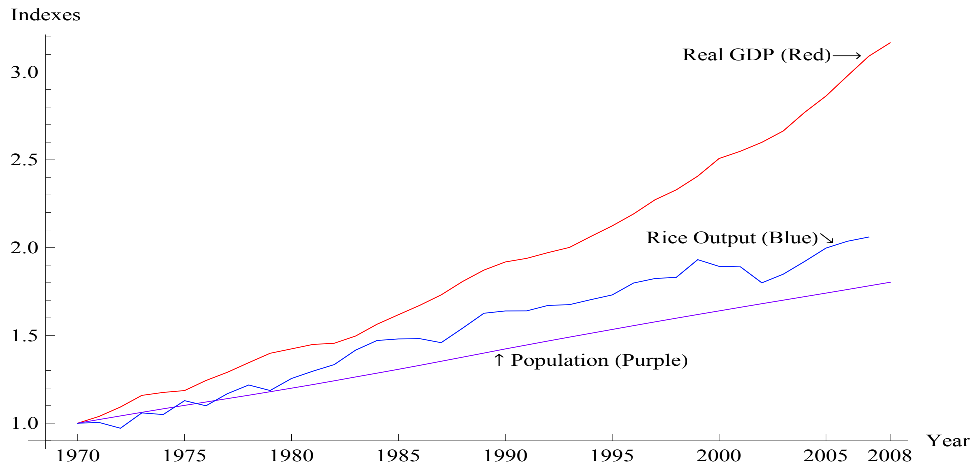

(a) Rice

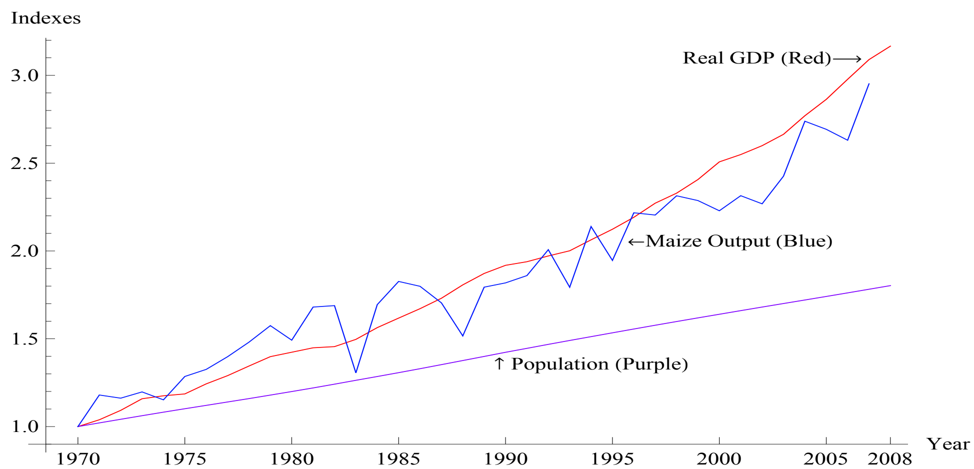

(b) Maize

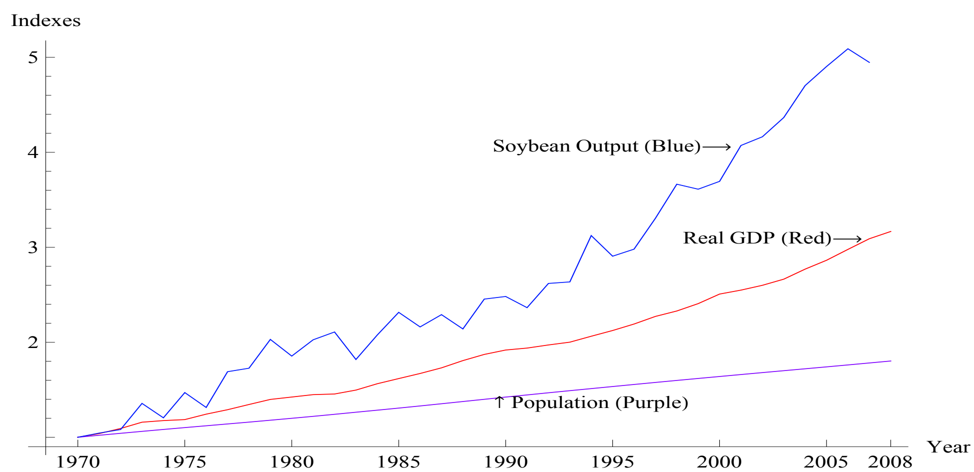

(c) Soybean

Figure 10: Global Supply and Demand for Other Food Commodities: 1970-2008