University of Hawai`i at Mānoa
Department of Economics
Working Paper Series

Saunders Hall 542, 2424 Maile Way,
Honolulu, HI 96822
Phone: (808) 956 -8496
www.economics.hawaii.edu

Working Paper No. 15-12

Commodity Prices and Volatility in Response to Anticipated Climate Change

By
A. Nam Tran
David Lobell
Michael J. Roberts
Wolfram Schlenker
Jarrod R. Welch

September 2015

Commodity Prices and Volatility in Response to Anticipated Climate Change

A. Nam Tran[1], David Lobell[2], Michael J. Roberts[3], Wolfram Schlenker[4,5], and Jarrod R. Welch[6]

[1]University of Missouri
[2]Stanford University
[3]University of Hawai'i at Manoa
[4]Columbia University
[5]National Bureau of Economic Research
[6]Charles River Associates

September 5, 2015

Abstract

Some predict that climate change will decrease average crop yield and increase yield variability. While the first effect, as well as possible adaptation strategies, have been studied extensively, the second is less well understood and the topic of this paper. A unique feature of commodity crops is that they can be stored between periods, thereby allowing storage to smooth production shocks across time. We pair a rational competitive storage model with a statistical analysis linking global production of the four major commodity crops (maize, wheat, rice and soybeans) and climate forecasts from 16 global climate models. The rational storage model predicts a doubling of average storage levels by 2050, slightly raising average prices to cover higher storage losses, but at the same significantly reducing price variability compared to a storage rule that is optimal under past yield distributions. Storage market responses to future yield variability greatly mitigate potential welfare losses of greater production volatility.

1 Introduction

International food prices have been rising since 2002, with record high spikes in the FAO Food Price Index in 2008 and again 2011. Prices signal resource abundance, with high levels indicating scarcity and high volatility reflecting a misallocation of resources across time. Although large price fluctuations may go unnoticed by consumers in developed countries, they can be devastating for poor consumers in less-developed areas who spend a large share of household income on raw staple grains and oil seeds. [1] Weather anomalies have played a clear role in recent price fluctuations (Abbott et al. [2011]) and mounting evidence indicates that climate change will continue to adversely influence agricultural crop yields and cause greater year-to-year variability (IPCC [2007], Schlenker and Lobell [2010], Urban et al. [2012]). Recent research has begun to consider potential effects of climate change on food prices (Diffenbaugh et al. [2012]), but without accounting for adaptive responses in storage behavior. Some policy makers have advocated for strategic grain stockpiles to address price volatility, but such policies should first account for the likely scale and effectiveness of market responses.[2]

Competitive storage theory stipulates that inventories are driven by weather-related variations in crop yields. Positive yield shocks cause prices to decline, and lead to accumulation of inventories since prices are low relative to expected future prices. Negative shocks cause prices to spike, giving an opportunity for competitive storers to divest their inventories. Thus,

[1] The view that food price volatility can be harmful to the poor may seem contrary to standard economic theory, in which utility is convex in prices thereby implying that consumers prefer variable prices over prices stabilized at their arithmetic mean (Waugh [1944], Massell [1969], Samuelson [1972]). Following this basic theory, it has been recently argued that price volatility is irrelevant with regard to poverty and food security — that only high price levels are important (Barrett and Bellemare [2011]). However, poor consumers are typically underinsured, and even a short period of undernourishment can have sustained effects on health, productivity and survival (Dawe and Timmer [2012]). Other research suggests high food prices might cause civil instability (Lagi et al. [2011]). Higher volatility increases the odds that critical price thresholds are exceeded.

[2] In September of 2012, French president Fracois Hollande launched a global campaign to create strategic stockpiles of agricultural commodities aiming at taming food price volatility. `http://www.reuters.com/article/2012/09/11/us-france-grain-hollande-idUSBRE88A1C220120911`. Accessed: July 17, 2013.

competitive storage markets act to smooth consumption in the face of randomly fluctuating supply, although we might imagine situations where the competitive equilibrium would not quell volatility as much as desired.

In a similar vein, storage can act to smooth consumption in the face of cyclical or secularly rising or falling demand or supply. Historically, yields have trended up, just slightly faster than population and demand growth. If the secular decline in price were generally anticipated in advance, it would have generally discouraged storage relative to the stationary case. In practice, although real prices have trended slightly downward over the long run, the trend is gentle enough that commodity price modeling often treats the system as stationary, as if secular increases in demand and supply have been anticipated and equal.

Looking forward, it is becoming harder to justify the idea that supply will necessarily grow at a pace that is greater than or equal to demand. Our aim here is to therefore examine competitive storage and price dynamics within a system where supply does not grow as quickly as demand, and supply simultaneously grows more volatile. We use this model to investigate two key questions: (1) To what extent would an anticipated secularly diminishing trend in yield show up in prices today via anticipatory storage adjustments? (2) How well can market-induced storage adjustments quell the influence of rising yield variability on price variability?

We calibrate anticipated changes in crop yield distributions by forecasting from a recent global statistical analysis of weather and crop yields of the world's four most prevalent staple food commodities — maize, rice, soybeans and wheat — using projections from 16 climate models (Lobell et al. [2011]). These projections do not account for some kinds of long-run potential adaptation. And surely there are a range of possibilities that lie outside the span of the sixteen climate models and weather draws we consider. Still, we believe this calibration provides some empirical basis for the scale of gradual impacts we might expect, and lets us see how competitive storage markets might account for those changes as they happen.

Given our projected changes in yield distributions, we solve for dynamic competitive storage market equilibria, which gives a unique solution for optimal storage and planned production as a function of food availability for each year along an adjustment path. These functions implicitly determine price as a function of the stochastic yield outcome, and thereby allow us to simulate the time paths of price, consumption, production and storage outcomes, as well as producer and consumer surpluses. These time paths of prices and surpluses are compared to those simulated using a stationary yield distribution, as well as a case using the non-stationary distribution but without adaptive changes in the storage rule.

This is the first attempt to calibrate a complete stochastic, dynamic equilibrium path for food commodity prices under climate change, taking into account secularly divergent growth of demand and supply, plus changing production volatility. In this way our approach also differs from previously used storage models that analyze stable demand and stationary yield distributions (Deaton and Laroque [1995], Williams and Wright [1991], Cafiero et al. [2011]).

As with any forecasting exercise, this analysis has limitations. Advances in agricultural technology have led to marked increases in productivity in the last half-century, and we do not account for adaptation of technologies other than storage. Also, our main results exclude CO_2 fertilization effects, however we consider results with crop-model-based CO_2 effects in a sensitivity analysis. While all of these would tend to counteract the downward pressure on agricultural yield and upward pressure on prices, we have already seen yield growth begin to slow for some crops, for example, rice (Dobermann et al. [2004], Auffhammer et al. [2006]). On the demand side, our analysis also abstracts from trends such as population and economic growth. In sum, these assumptions are equivalent to an assumption that supply-side and demand-side changes would, in the absence of climate change, balance each other out, an assumption that is implicit in much of the earlier literature on commodity prices. In consideration of these caveats, our primary intention is not to draw specific conclusions about price forecasts in the distant future, but rather to emphasize salient features of an adaptive dynamic pricing system that are robust to a wide range of parameters and expectations.

2 Projections of Future Crop Yields

Projections indicate that climate change will continue to exert downward pressure on the global mean yield and simultaneously increase yield variability relative to a situation without climate change (Figure 1). These projections are based on yield-climate response functions for corn, rice, soy, and wheat estimated in Lobell et al. [2011]. These projections remove the prevailing trend in yield due to technological change and should therefore be interpreted relative to a situation without climate change, not necessarily relative to global yields realized in the year 2000. Since historical prices appear to be stationary or nearly so, the baseline assumption is that continued technological advance will equal demand growth from rising population and income.

To derive yield-climate response functions, historic yield shocks (deviations from a quadratic time trend) were regressed against a quadratic function of temperature and precipitation. A block bootstrap procedure was used to estimate the parameter distributions. We combine the first 100 bootstrap replications for each of the four crops with 120 years (1961-2080) of weather output from 16 GCMs, thus constructing 1,600 hypothetical 120-year time-series for each crop.[3] Using the conversion factors from Williamson and Williamson [1942], we convert the yield values to calories and aggregate them to obtain global caloric yield. We express this in units of number of persons fed for a year per one hectare of harvested land, assuming an average daily diet of 2,000 calories, or 730,000 per year; total production will be expressed in terms of billions of persons fed for a year.

The variance of the predicted values is lower than the actual variance because it excludes residual variance. To correct this bias, we add the residual variance from the regressions to the projections. Starting with global yield data for 1961-2009 from the FAO statistical database,

[3] We hold year fixed at 2000, so our yield data are weather-induced deviations from average yield in 2000.

we detrend both series, fitting log-yield to a quadratic time trend. We then use the difference in trend values for the year 2000 to adjust the mean value for each year of predicted values. Then we add the difference between the actual variance and the predicted variance for the overlapping years (1961-2009) ($SD_{observed} - SD_{predicted}$.) This calculation assumes that the residual variance remains constant.[4]

Using adjusted yield variances, we generate a yield distribution for each year from 2000-2080 ($T = 0$ for the year 2000), using a discrete approximation of a normal distribution with mean equal to the mean across the 1,600 predicted observations for the corresponding year, and variance equal to the mean of the rolling variance of the deviations for each of the 1,600 series. The projections indicate mean yield decreases almost 7%, from about 10.75 persons per year per hectare in 2000 to about 10 in 2050, while the standard deviation in the baseline case increases 15%, from about 0.233 in 2000 to about 0.269 in 2050.

Regression-based predictions of future yield do not account for the positive effect of CO_2 on crops. For this and other reasons, it is possible to anticipate future yield distributions with higher mean or lower variability than our estimates. It is thus useful to investigate the two effects separately as well as jointly. To consider mean and variance effects separately, we simulate two additional sets of future yield data. The first holds the coefficient of variation fixed at the 2000 level and projects only the predicted mean yield; the second holds mean yield constant, and projects the increasing coefficient of variation forecast of the model.

3 Storage and Commodity Prices

To connect projected climate-induced shifts in growing conditions to prices, we use a competitive storage model. The competitive storage model has a long history in the literature on commodity price behavior, starting with seminal work by Gustafson [1958]. This arti-

[4]Another way to do this would be to scale the forecasted yields by ($SD_{observed}/SD_{predicted}$). The additive adjustment is more conservative so we focus the presentation on these results.

cle lays the foundation for the theory of optimal storage and demonstrates that in seeking opportunities for intertemporal arbitrage, competitive storage can smooth the effects of production risk. The resulting time path of storage and commodity prices maximizes social welfare, measured as the present value of producer and consumer surplus (Scheinkman and Schechtman [1983]). Later, Williams and Wright [1991] presents a thorough exploration of the economic implications of the storage model. The authors extend the storage model in multiple directions, exploring the differences between storage models for raw materials versus finished goods, the interaction between storage and trade, implications of market power, the incorporation of new information, government intervention, and assert that the qualitative features of the price behavior of several important commodities are consistent with the model. However, in each of these cases the analysis focuses on comparative statics of alternative stationary environments. In our study, rather than compare alternative climate equilibria, we focus instead on the broader dynamic implications of the gradual adjustment path from one climate to another.

Although the theory of competitive storage has been widely recognized and embraced (Williams and Wright [1991]; Deaton and Laroque [1992, 1995, 1996]; Cafiero et al. [2011]) some have challenged its empirical validity. Deaton and Laroque [1992, 1995, 1996] pioneered econometric methods for estimating the storage model using only price data. These studies present a wealth of empirical evidence demonstrating that speculative storage replicates key features of commodity prices, like long periods of stable prices interspersed with relatively short periods of volatile, upward spikes. A key criticism is that while storage causes prices to become more autocorrelated, the stationary model is incapable of generating the high degree of autocorrelation typically observed in commodity price dynamics. Cafiero et al. [2011], aided by almost two decades of increased computational power, reassess the relevance of the storage model and demonstrate that with a more precise estimate of the stockout price, the storage model can explain the most prominent features of the dynamics of commodity prices, including high autocorrelation, although their analysis does involve a less elastic demand curve. There is also a large standard error for autocorrelation coefficients with time series of reliable historical length (50-60 years).

The competitive storage model assumes the efficient market hypothesis (an absence of arbitrage) can reasonably characterize commodity markets. In some cases, episodes of price volatility may seem unrelated to fundamentals of market supply and demand. However some uncertainties, like current inventories, future demand and trade policies, can be difficult to quantify, and these factors are typically assumed known. In the standard model, uncertainty derives only from stochastic weather and production. There are also a few documented cases where large-scale investors have attempted to corner commodity markets.[5] Despite these ambiguities, UNCTAD [2009] argues that the financialization of commodity futures trading has attracted traders who may not trade exclusively based on supply and demand fundamentals, and have increased the likelihood of what they call behavioral overshooting in commodity markets. Overshooting would presumably cause less, not more, autocorrelation in commodity prices, which seems to contradict evidence described above. Bobenrieth et al. [2012] present a model that demonstrates that in the presence of supply and demand shocks in a rational storage model, cash prices of storable commodities can behave in a bubble-like fashion, even with fully observable inventories and no uncertainty about future demand.

We use a competitive storage model because it provides a coherent baseline to which long-run fundamentals should eventually be drawn. We also consider a naive model in which storage rules are held constant while current markets, conditional on storage, are assumed to clear. In other words, we consider models in which climate change is both anticipated and unanticipated by storers. To focus on large-scale fundamentals, we also consider a single competitive world market for storable food commodities and costless trade.

The market consists of three primary agents: producers, consumers and storers. While

[5] The Hunt brothers' infamous foray into the silver market in 1979 and 1980 is perhaps the most well known attempt to corner a commodity market. We know of no documented cases for consumable food commodities.

price is endogenous to the model, individual agents are price takers. In any period t, the amount of food initially available, z_t, is comprised of production in period t plus the depreciated carry-over from period $t-1$:

$$\begin{aligned} z_t &= y_t l_{t-1} + (1-d)x_{t-1} \\ &= c_t + x_t \end{aligned} \tag{1}$$

where y_t is the realized yield, l_{t-1} is the planting decision from the previous period, x_{t-1} is the amount of inventory left from the previous period, c_t is consumption, and d is the decay rate.

Current market price is a strictly decreasing function of consumption:

$$p_t = F(c_t). \tag{2}$$

And area harvested is a strictly increasing function of expected marginal revenue:

$$l_{t-1} = G(E(p_t y_t)). \tag{3}$$

Yield is a random process partially determined by weather (w_t):

$$y_t = h(w_t) + \varepsilon_t, \tag{4}$$

where $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$.

With the existence of storers, the discounted expected future price will not exceed the current price by more than the cost of storage. If the expected future price is too high, storers will exploit the opportunity to store more, increasing the current price and decreasing the expected future price. If the expected future price is too low, storers will not carry inventory over into the next period because doing so will result in economic loss. The arbitrage conditions for the competitive storage market are:

$$p_t + k \geq \frac{1-d}{1+r} E_t(p_{t+1}), \quad x_t \geq 0, \tag{5}$$

where r is interest rate and k is the positive constant per-unit storage cost.[6]

When there is no storage, all available stock is consumed, so inverse demand will be a function of the amount on hand:

$$p_t = F(z_t). \tag{6}$$

However, when storage is positive, consumption is less than amount on hand; current price is therefore greater than $F(z_t)$ and is a function of the discounted expected future price. By combining (5) with equation (6), we can express price as

$$p_t = max \left\{ \beta E_t(p_{t+1}) - k, F(z_t) \right\}, \tag{7}$$

where $\beta = \frac{1-d}{1+r}$.

Given assumptions of costly storage and i.i.d. harvests, others have established the existence of a stationary rational expectations equilibrium (SREE) for the storage model (Scheinkman and Schechtman [1983], Deaton and Laroque [1992]). Using this framework, it is relatively straightforward to analyze the comparative statics of two separate, stationary yield distributions. Rather than focus on comparing two stationary cases, we examine what happens to storage behavior and price volatility when the market experiences a gradual shift in the yield distribution, similar to what we might expect from climate change. Specifically, we assume that historical weather realizations are i.i.d. and at time $t = 0$ climate starts to

[6] In this study we do not consider a *convenience yield* which, at low inventory levels, might effectively make storage costs negative.

shift gradually, moving towards a new equilibrium distribution and that will stabilize at the new stationary distribution at time T. Consistent with the literature on agricultural impacts of climate change (e.g. Schlenker and Roberts [2009], Urban et al. [2012]), we assume that $h(\widetilde{\omega}) < h(\omega)$, and $\widetilde{\sigma_w^2} > \sigma_w^2$. We also assume that $h(\cdot)$ remains unchanged, and that σ_ε^2 is constant. Combined with the other assumptions inherent in the model (e.g. constant demand) this assumption is equivalent to the assumption that all non-climate-related changes (e.g. technological change, population growth and income growth) balance each other out.

3.1 A Stationary Storage Model

Because no closed-form solution exists, solving the storage model requires the use of stochastic dynamic programming. We follow the steps used in Deaton and Laroque [1996] and Cafiero et al. [2011] that utilize a flexible functional form to approximate the control variable, price (p), as a function of the amount on hand z. We adjust the numerical procedure to allow for a supply response in the storage model. In the next section we extend the process further to account for a nonstationary yield distribution.

Assume that inverse consumption demand and land supply response functions have the following constant elasticity forms:

$$p_t = F(c_t) = \alpha_d c_t^{\gamma_d}, \tag{8}$$

$$l_t = G(p_{t+1}y_{t+1}) = \alpha_s E_t(p_{t+1}y_{t+1})^{\gamma_s}; \tag{9}$$

with $\gamma_d < 0$ and $\gamma_s > 0$ We can rewrite Equation (7) as:

$$p_t(z_t) = \max\left\{\beta E_t(p_{t+1}(y_{t+1}l_t + (1-d)x_t) - k, F(z_t)\right\}, \tag{10}$$

and from Equation (1) we have:

$$x_t = z_t - F^{-1}(p_t(z_t)). \tag{11}$$

To further simplify the model, the yield distribution is discretized into N points, each with a corresponding probability $Pr(y_{t+1}^n) = Pr(y_{t+1} = y_{t+1}^n)$.[7] Thus, we rewrite (7) and (9) as:

$$p_t(z_t) = max\left\{\beta\sum_{n=1}^{N} p_{t+1}(y_{t+1}^n l_t + (1-d)x_t) \times Pr(y_{t+1}^n) - k, F(z_t)\right\}, \tag{12}$$

and

$$l_t = \alpha_s(\sum_{n=1}^{N}\left[p_{t+1}(y_{t+1}^n l_t + (1-d)x_t)y_{t+1}^n\right] \times Pr(y_{t+1}^n))^{\gamma_s}, \tag{13}$$

where $\sum_{n=1}^{N} Pr(y_{t+1}^n) = 1$. Given parameters of the demand and supply equations, solving (12) for price as a function of amount on hand, $p(z)$, is an iterative process as follows:

1. Set a range of amount on hand values $[\underline{z},\ \overline{z}]$.[8]

2. Guess a set of initial values for (12), using a linear spline to interpolate values between the N index points of $z_t \ \epsilon\ [\underline{z},\ \overline{z}]$.[9]

3. Given these initial values, calculate storage x using equation (11).

4. Using x from the previous step, solve the system of non-linear equations (13) to find the corresponding land area.

5. Compute price p using the right hand side of equation (12).

[7] The discretization process is done using the CompEcon Toolbox for Matlab by Miranda and Fackler [2007]

[8] The z_t are bounded by $[\underline{z_t}, \overline{z}]$ where $\underline{z_t}$ is the lowest harvest possible in the model. While $\underline{z}$ can be set equal to the lowest production with zero storage, $\overline{z}$ needs to be carefully chosen to be greater than the highest possible amount on hand.

[9] While a higher number of points would yield a more accurate approximation of the price function, it is computationally costly. Cafiero et al. [2011] use 1000 equally spaced points so that they can find a more precise measurement of the cut-off price. We instead use 150 unequally-spaced points, with the majority of the points clustered around the no-storage equilibrium of production. We find no significant difference in the cut-off price using a finer grid of 1000 equally spaced points, and our method of 150 unequally spaced points.

6. If the price from step 5 is equal to that in step 2, stop the iteration process. Otherwise, repeat steps 3-5, using the values generated in step 5. Repeat until the vectors of prices from step 2 and 5 equal each other.

Using this basic structure for solving the stationary case, we proceed to modify it for the non-stationary case in the next section.

3.2 Adaptive Storage with Climate Change

Storage theory indicates that when a shock is anticipated in the system, price and storage will move together at the same direction. With an anticipated contraction in supply, storers – motivated by intertemporal arbitrage – will accumulate inventories in anticipation of higher future prices. In doing so, they reduce quantity supplied for current consumption and thereby drive up current price. Normally this increase in the current price would limit storers' incentive to accumulate inventories. Yet, because the future price is expected to keep increasing in response to a continuing contraction in supply, more may be withheld from the market which lead to increasingly higher prices during the transition period. The accumulation of storage will only cease in response to unanticipated shocks, or when amount on hand reaches the stationary level appropriate for the new environment. From that point, price increases along the adjustment path will gradually become smaller.

We use backward induction, and numerical methods like those described above, to solve the non-stationary storage model and examine how price and other variables will adjust to an anticipated, gradual, and long-term shift in the yield distribution. We first solve for a terminal condition – the storage equilibrium in year T, assuming that the market equilibrium subsequently remains stationary. We then proceed backwards, first solving the storage model from period $T-1$, proceeding backwards one period at a time, back to period 0, the first period in which the market reacts to the news of an impending climatic shift. The storage amount in period $T-1$ is:

$$\begin{aligned} x_{T-1} &= z_{T-1} - c_{T-1} \\ &= z_{T-1} - F^{-1}(p_{T-1}(z_{T-1})). \end{aligned} \tag{14}$$

Equation 12 is modified to express the relationship between $p_T(z_T)$ and $p_{T-1}(z_{T-1})$ over the range of z:

$$p_{T-1}(z_{T-1}) = max\left\{\beta\sum_{n=1}^{N} p_T(y_T^n l_{T-1} + (1-d)x_{T-1}) \times Pr(y_T^n) - k, F(z_{T-1})\right\}. \tag{15}$$

It can seen from Equation 14 and Equation 15 that the storage demand function in period $T-1$ can be found once we know $p_T(z_T)$ and l_{T-1}. Recall that the land function, accounting for shifting yield distributions, can be implicitly solved using the following equation:

$$l_{T-1} = \alpha_s(\sum_{n=1}^{N}[p_T(y_T^n l_{T-1} + (1-d)x_{T-1})y_T^n] \times Pr(y_T^n))^{\beta_s}, \tag{16}$$

where y_T^n is the n value of yield in year T. We use the following iterative approach to find $p_{T-1}(z_{T-1})$:

1. Pick l_T as the starting values for land. (Note: this is found from solving the stationary case at time T.)

2. Solve for $p_{T-1}(z_{T-1})$ using equation (15) and l_T.

3. Solve for inventory and land using equations (14) and (16).

4. If the land value replicates itself, stop and use it to solve for $p_{T-1}(z_{T-1})$. Otherwise, repeat steps 2-4, using this new land value.

Once $p_{T-1}(z_{T-1})$ is found, we use it to find $p_{T-2}(z_{T-2})$ and so on until we get to $p_0(z_0)$. In the end, we have a system of equations describing how the time paths of storage behavior and prices would reflect the market reaction to news of a gradually shifting yield distribution.

3.3 Non-Adaptive Storage with Climate Change

When there is no storage adaptation, the storage rule does not change over time. We begin by solving the storage model in the baseline year where there is no climate change, and assume that the resulting relationship between storage and amount on hand remains fixed. We then apply the same storage rule to find solution of the storage model when supply has shifted completely to the new environment. Using this newly found storage function in the last period T, holding storage rule fixed, we then use backward induction to back out storage functions in year $T-1$. We keep solving backward until period 0.

A formal description of the solution is as follows.

We first solve the storage model in the baseline case with no climate change. Denote this scenario as b, the price and land function can be re-written as:

$$p(z_b) = max\left\{\beta\sum_{n=1}^{N} p(l_b y_b + (1-d)x_b) \times Pr(y_b) - k, F(z_b)\right\}, \text{and}$$
$$l_b = \alpha_s(\sum_{n=1}^{N} p(l_b y_b + (1-d)x_b \times Pr(y_b^n))^{\beta_s}.$$

Solving these two equations yield storage values, x, over the range of z. We use storage and amount on hand to build the storage rule for the baseline case without climate change. Using this storage rule, the storage demand function after supply has completely shifted out can be found using:

$$p_T(z_T) = max\left\{\beta\sum_{n=1}^{N} p_T(y_T^n l_T + (1-d)\mathbf{x_b}) \times Pr(y_T^n) - k, F(z_T)\right\},$$
$$l_T = \alpha_s(\sum_{n=1}^{N} [p_T(y_T^n l_T + (1-d)\mathbf{x_b})y_T^n] \times Pr(y_T^n))^{\beta_s},$$

where y_T represent the same yield distribution for all years from T to infinity. Note that because storage rule is held fixed, there are effectively two state variables in the system, storage and amount on hand. After finding $p_T(z_T)$, we follow the same iterative approach as described earlier to find $p_{T-1}(z_{T-1})$ and price functions for the previous years. Specifically, given $p_T(z_T)$, solving the two equations:

$$p_{T-1}(z_{T-1}) = max\left\{\beta \sum_{n=1}^{N} p_T(y_T^n l_T + (1-d)\mathbf{x_b}) \times Pr(y_T^n) - k, F(z_T)\right\}, \text{and}$$

$$l_{T-1} = \alpha_s(\sum_{n=1}^{N}[p_T(y_T^n l_{T-1} + (1-d)\mathbf{x_b})y_T^n] \times Pr(y_T^n))^{\beta_s}.$$

yields $p_T(z_{T-1})$.

3.4 Adaptive Storage with Unanticipated Climate Change

One objective is to examine how much storage and commodity prices might be affected today if the market anticipates gradually more severe production impacts from climate change going forward into the future. The challenge is to decompose this *anticipatory* effect of climate change from current effects of climate change as they happen. To decompose these effects, we consider a model that accounts for climate change but does not account for the anticipatory effects of climate change, and can thus infer the anticipatory effect as the difference between the full forward looking model described above and the model without anticipation. The model without anticipation still has storage and consumption functions that change each year with climate change. The difference is that changes in the storage and consumption rules are not themselves anticipated. This scenario requires solving an infinite-horizon pricing function in each year, but then applies that infinite-horizon solution for only one year. That is, each year storage and pricing decisions are made assuming climate will remain fixed going forward, but each year the climate nevertheless changes slightly.

We evaluate scenario by solving the storage model for each year of the adjustment path as if the forecast yield distribution for that year were expected to remain stationary going forward. Specifically, we solve the two equations:

$$p_t(z_t) = max\left\{\beta \sum_{n=1}^{N} p_t(l_t y_{t+1} + (1-d)x_t) \times Pr(y_{t+1}) - k, F(z_t)\right\}, \text{and}$$

$$l_t = \alpha_s(\sum_{n=1}^{N} p_t(l_t y_{t+1} + (1-d)x_t \times Pr(y^n_{t+1}))^{\beta_s}.$$

where t is any year in the transition process with its own yield distribution y_t.

3.5 Welfare calculation

To estimate the social impact of the price changes, we calculate the resulting producer and consumer surplus in our model.[10] Rather than calculate total consumer surplus which would be $+\infty$ for a constant elasticity demand curve, we calculate the difference in consumer surplus using:

$$CSDIFF = \int_{p_{2000}}^{p_t} F^{-1}(p_t)dp_t,$$

for $t = 2001 - 2080$. For producers, we compute the welfare in time t, which equals revenue in time t minus the production cost in time t:

$$PS = p_t q_t - \Phi(l_t).$$

This standard welfare measure clearly does not account for persistent effects of price spikes on human health or indirect effects like civil conflict or political instability.

[10]The storage arbitrage condition implies a zero surplus for storers.

4 Supply and Demand

Demand and supply parameters, reported in Table 1, were drawn from existing literature. Some demand elasticities were estimated while others were assumed. Most studies in the commodity pricing literature assume perfectly inelastic supply, so in these studies the demand elasticity implicitly accounts for both demand and supply response.[11]

Table 1: Storage model parameters

Parameters	**Values**	
Demand elasticity	-0.04	-0.08
Supply elasticity	0.08	0.04
Decay rate	0.01	
Interest rate	0.02	0.05
Storage cost	0.87	
Non-stochastic price in year 2000 (dollars)	43.29	
Non-stochastic future price in year 2000 (dollars)	49.70	
Non-stochastic land in year 2000 (billion hectares)	0.58	
Non-stochastic consumption in year 2000 (billion people)	6.21	

Estimated values for the combined elasticities (demand minus supply) range from -0.04 to -0.2. Roberts and Schlenker [2013] perform a reduced-form analysis in which they use contemporaneous and lagged weather shocks to identify the price elasticity of both demand and supply; they estimate elasticity of demand to be about -0.05, and elasticity of supply to be about 0.1, which lies in the range of most commodity pricing studies. As a baseline we use a demand elasticity of -0.04 and a supply elasticity of 0.08 because -0.12 is the middle of the range of estimates of combined elasticities from the storage literature, and this split between supply and demand preserves the relative magnitudes of the estimates in Roberts and Schlenker [2013]. Elasticities on the lower end of this range are needed to rationalize observed price volatility and autocorrelation [Cafiero et al., 2011, Gouel and Legrand, 2015]. Additional results for the entire range from -0.04 to -0.2 for the combined elasticities are

[11] For example, using econometric methods developed in various papers by Deaton and Laroque, as well as a more recent maximum likelihood estimator by Cafiero et al. [2015], which use only price data, it is not possible to separately identify supply and demand.

reported as sensitivity checks. We set the parameters of demand and supply functions such that they are in equilibrium (with zero storage) at the observed price and quantity values in the year 2000, and the elasticities are equal to the chosen values at that point. We also follow the previous storage literature in choosing the parameters for depreciation and storage cost: we set storage cost to be equal to 1% of the price equilibrium in the static setting. Decay rate and interest rate are assumed to be 0.01 and 0.02 respectively. [12]

The constant parameter in the demand equation is set so that, given the observed price of 43.29 in year 2000 and demand elasticity of -0.04, the quantity consumed equals the consumption level of 6.21 in year 2000.[13] The constant parameter for the supply equation is set so that given a supply elasticity of 0.08 and the producer expected marginal revenue, land area equals that observed in year 2000.[14]

5 Model Validation

To obtain a sense of how well this model can rationalize historical prices, we solved for the dynamic competitive equilibria assuming fixed demand and a stationary yield distribution with the baseline parameters. We then fed historical global yields through the model to simulate a stochastic equilibrium time path of prices. Global yield shocks were calculated as aggregated deviations from crop-specific yield trends, which serve as a proxy random weather surprises. We compare these simulated prices to the actual historic caloric price index.

To our knowledge such a comparison of simulated and actual prices has never been at-

[12]Cafiero et al. [2011] find that the estimated food storage decay rate is close to zero once the pricing function is estimated more precisely with a finer grid.

[13]Price is the dollar amount in year 2000 to buy 730,000 calories, which equals the sum of 2000 per day for 365 days. Quantity consumed is expressed as billions of people, assuming each consumes 730,000 calories annually. These annual consumption quantities are not meant to be realistic; we simply use these units so that numbers can be easily understood in global perspective.

[14]The expected future price in year 2000 equals to 49.70, which equal the production weighted price of the four grain futures contracts that was scheduled to be delivered in December of 2001 and was . The covariance between yield and price is set to -0.49, which is the covariance between detrended yield and price using FAO crop data between 1961-2009. Land in year 2000 equals 0.58 billion hectares.

Table 2: Comparison of Actual and Model-Simulated Prices

	Full Sample		*Restricted Sample*	
	Actual Price	Simulated Price	Actual Price	Simulated Price
Mean	68.18	56.64	62.19	57.03
Std. Dev.	21.53	9.07	12.93	9.54
Correlation	-0.01		0.20	
Corr $\Delta \log p_t$	0.32		0.53	

The table compares actual and model-simulated caloric prices. The full sample includes prices from 1961 through 2009. The restricted sample excludes two periods that experience prominent demand-oriented shocks, 1972-1975 and 2005-2008, which correspond to the Russian Wheat Deal and major policy-induced ethanol expansion.

tempted using a competitive storage model. Instead, because identifiable production shocks are not typically identifiable in commodity markets, and because, in reality, both demand and supply may experience unexpected variability, estimation and calibration exercises typically compare moments of the price distribution, not the actual time path of prices. It is therefore important to keep in mind that this simulation assumes stable, fully anticipated demand, with price variability following only from random, transitory supply shocks. The simulation does not account for demand shocks.

Because the data start in 1961, model parameters are set so that the intersection of supply and demand match observed price, consumption, land and expected revenue of the year 1961. In order to account for the fact that there is no trend in our baseline storage model, observed price and yield are also de-trended by removing a quadratic time trend and then centered on the trend level of 2005. We then set the simulated price in the first year equal to the observed price in 1961 and then find the associated amount on hand in 1961 using the inverted storage equilibrium (i.e., amount on hand is expressed a function of price). Given the amount on hand and price, we subsequently find consumption, production, land, and inventories in the first year. Then, using inventories from 1961 and observed yield in 1962, we can find all the other variables for the year 1962. We continue to solve forward for each year until we reach 2009. Simulated and historically observed prices from 1961 to 2009 are plotted figure 1.

Unsurprisingly, model simulated prices do not closely match historical prices. The main

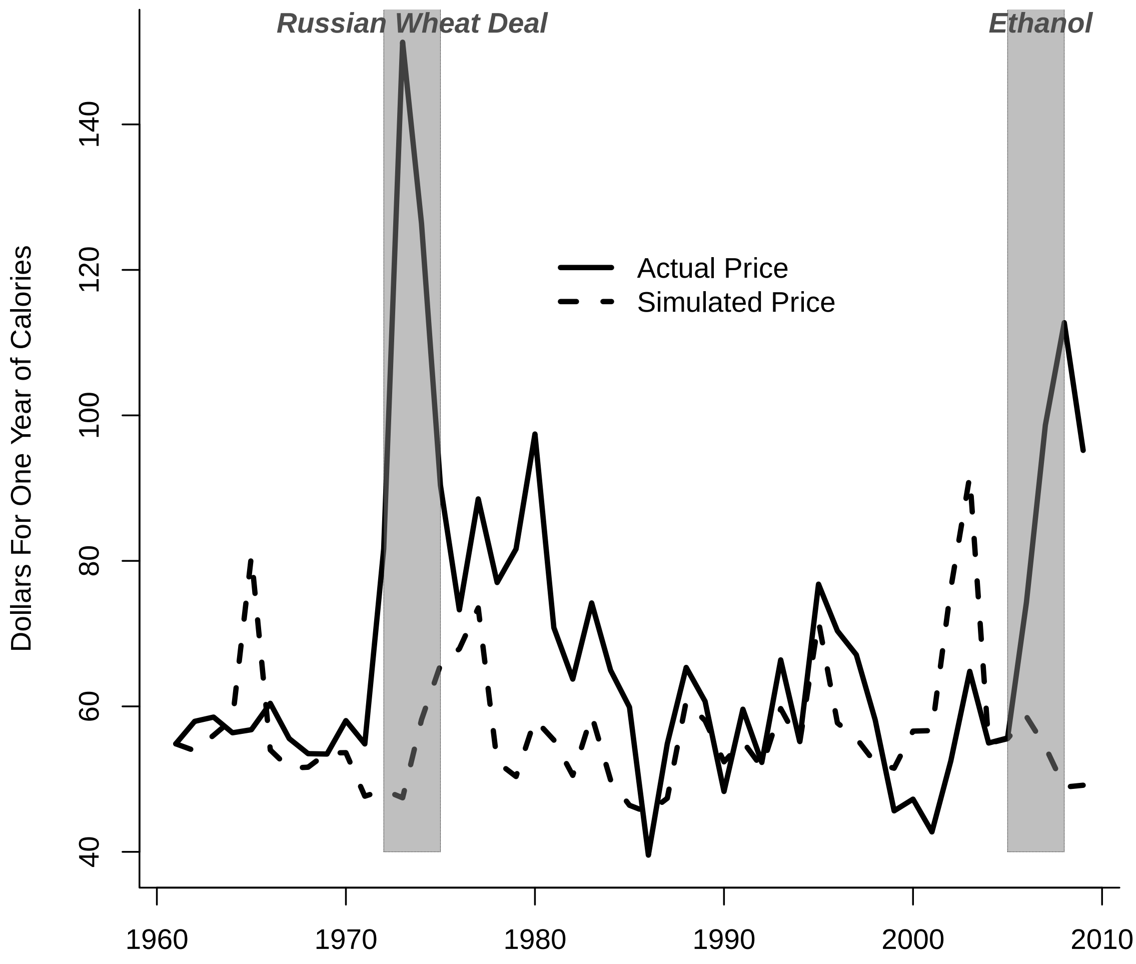

Figure 1: **Actual and model-simulated prices from 1961-2010.** The graph shows actual, de-trended average caloric price and simulated price derived from feeding de-trended historical production values through our baseline model. Because the model assumes stable demand and partially-random (weather-driven) supply, we highlight three major demand-related events that could cause major departures between our model: (1) the Russian Wheat Deal of 1972; (2) the Asian financial crisis; and (3) the ethanol boom which followed from generous subsidies and mandates.

reason for this departure likely stems from demand shocks and perhaps other non-weather related shocks to food commodity prices, such as changes in governmental policies and trade restrictions. One notable demand-related shock was the Russian Wheat Deal in the early 1970s, during which the United States sold a large share of its accumulated inventories to the Soviet Union. Other commodity prices shocks occurred at the time, including a large oil price spike related to the Iranian embargo, which may have influenced *prospective* future supply of staple food grains, since fuel and fertilizer are critical agricultural inputs. These shocks would not be well captured by crop yield fluctuations. A more recent demand-related factor followed from the ethanol mandate and ensuing boom in corn-ethanol production in the United States. Around the same time, China, India and other countries experienced rapid and sustained growth, which could be a key demand-related source for agricultural and other commodity prices rises [Trostle, 2008, Kilian and Hicks, 2013].

Although other kinds of demand and anticipated supply shocks likely affected commodity prices and storage decisions, model-simulated prices show a reasonably strong association with actual prices when we exclude the early 1970s and the late 2000s (Table 2). It is important to keep in mind that early real-world departures from the theory, like the 1970s demand-related shocks, affect the whole future path of prices via storage, so ought not to expect a perfect fit even if the model were perfect reflection of actual markets outside of these two excluded time periods. Thus, although the model cannot perfectly account for historical prices, we find the correlative evidence between the simulation and reality to be encouraging with regard to its ability to account for transitory yield shocks that are mainly attributable to weather. After all, our objective is to discern how markets and storage might anticipate and adapt to changing weather-related factors, holding all else the same.

6 Projecting Future Prices

We follow the steps outlined above to find storage demand and price functions for every year from 2000 to 2080. Next, we simulate two hypothetical price time-series. For the stationary yield distribution with no climate change, we draw 201 random yield observations from the distribution for the year 2000. We use a starting value of no storage so that in period one, the amount on hand is equal to the harvest. Given the amount on hand, we can subsequently solve for price, consumption, production, land, and inventories in the first period using the storage demand function assuming no yield shift in the model. Then, with inventories from the first period and yield in the second period, we can find all the other variables in the second period, and continue sequentially to the final period. In order to avoid bias associated with selection of the starting value, we discard the first 120 observations. The remaining 81 observations give us the hypothetical price path from year 2000 to year 2080.

The price path and other variables depend on inventories and amount on hand, and these are influenced by random production shocks. To estimate the whole distribution of potential outcomes along the path, we simulate this price path 10,000 times to obtain an empirical distribution of prices for each period of each scenario.

We consider four scenarios. The first is a baseline, which assumes a stationary yield distributions around year 2000. For each of the three the non-stationary cases, we use the value of amount on hand in year 2000 from the baseline scenario as our starting value. We then simulate a price series similarly to the baseline case, but this time using the relevant price function for each year of the non-stationary case, as described above. We use these solutions to compare hypothetical price distributions under three different sets of assumptions: 1) no storage adaptation; 2) storage adaptation without anticipation of future climate change; and 3) storage adaptation with anticipated future climate change.

Because we assume the yield distribution remains constant after 2080, and climate or other changes may continue beyond that time, the time path of simulated data close to 2080 might be dubious. We therefore focus discussion on the time period between 2000 and 2050.

7 Results

Figure 2 shows the effects on average prices, price volatility, and storage under four different assumptions: 1) The dynamic storage model, with both declining yield and increasing variance as projected by our yield response functions and climate models; 2) The dynamic storage model with trending yield, and constant CV; 3) The dynamic storage model with constant mean yield, and increasing CV; and 4) A baseline case with no climate change, and a stationary storage model for comparison.

As a result of declining future yield and increasing variance, mean price increases from 47 dollars in 2000 to 87 dollars, while storage increases along the entire transition path.[15]

Inventory growth substantially ameliorates the influence of greater yield volatility on price volatility. When both yield level and volatility change, storage increases buffer shocks to such an extent that the resulting CV of price actually declines relative to the stationary case. This surprising result follows because inventories are accumulated both from higher yield volatility and the declining trend in mean yield.

When we consider yield trend and yield CV separately, we find storage increases more when yield is trends down than when yield CV is increases. Most of the increase in mean price is also due to the declining yield trend. While the CV of price declines, accumulation of inventories is not enough to reduce the absolute variance of prices. If there exist critical thresholds above which price causes greater harm to human health or political stability, then the risk of such events still increases markedly, especially in later years.

In Figure 3 we present results that how anticipation of changing yield distributions affects

[15] Note that the model does not replicate the level of observed inventories due to the possibility of a stock-out, which never happens in reality. Part of this is an aggregation problem combined with transactions costs of moving commodities from one location to another (Brennan et al. [1997]). There is also the influence of convenience yields, which we do not consider.

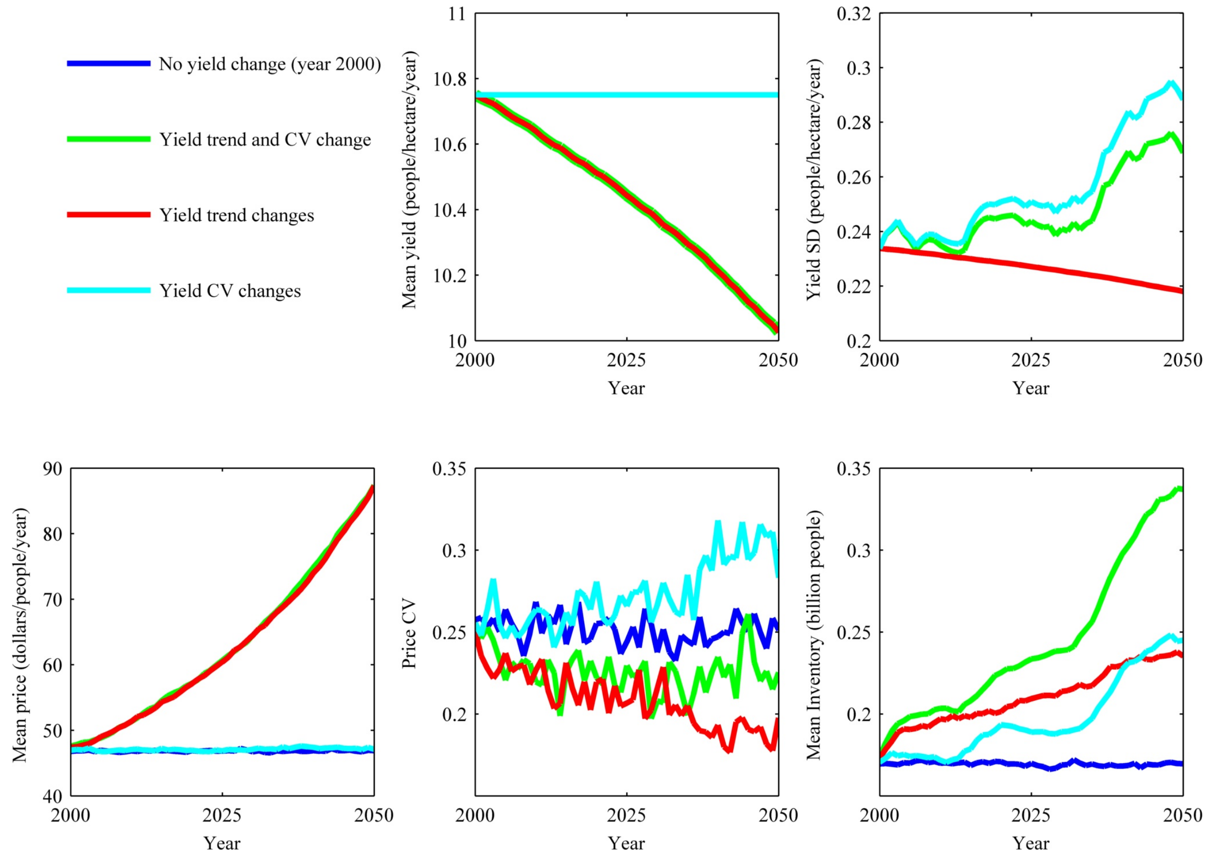

Figure 2: **Mean and variability of yield and price, and mean inventory under various scenarios** The top two panels depict the projected time-series plot of mean yield, and the standard deviation of yield under four different scenarios. The bottom three panels display the corresponding level of inventory, and resulting mean and coefficient of variation of price. While our yield projections and storage model span the years 2000-2100, the focus of our analysis is on the transition path, so we focus on the years 2000-2050 in the figures.

storage, price level and price volatility. The figure shows the difference between 2000 and 2050 of average price, price volatility, and storage under a scenario where storage is assumed to be constant as a proportion of amount on hand, and does not respond to expected climate change. When there is no storage adaptation, mean price increases 112 percent over 50 years while the CV of price grows 75 percent and inventories decline -71 percent. With storage adaptation but no anticipation of future climate change, mean price and price CV increase only 87 percent and 11 percent, respectively. The last scenario where there is storage adaptation with fully rational anticipation of future climate change observes the smallest increase in mean price at 84 percent, a 9 percent decrease of price CV, and the largest change of inventory of 94 percent.

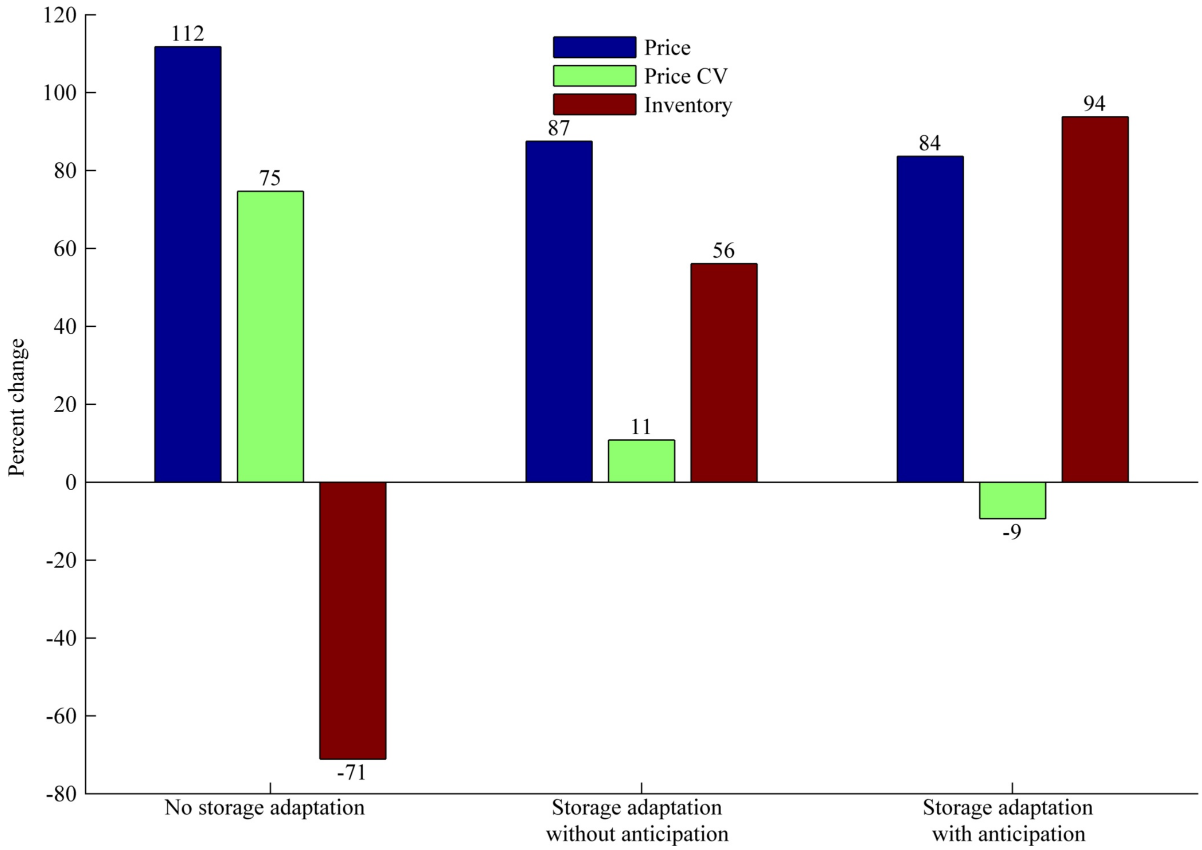

Figure 3: **Price, price volatility, and inventory changes under various storage rule assumptions** Each bar represents the percentage change over 2000-2050, and the number at the top of each bar gives this value. We compare three scenarios: no storage adaptation, storage adaptation to observed climate change, and storage adaptation to both anticipated and observed climate change.

Although price volatility is often measured by its CV, as with yield above above, it may also be meaningful to consider how absolute price volatility changes over time. Figure 4 shows the full distribution of price over 10,000 simulations. As expected, the price distribution in each year is skewed positively, with occasional high price spikes. Although storage increases buffer most changes in price variability, the trend in price is positive, and absolute price volatility increases. Define absolute price volatility as the percentage of price exceeding \$100, approximately double the mean price in the baseline case, we found that absolute price volatility increase more than 7 fold by 2050.

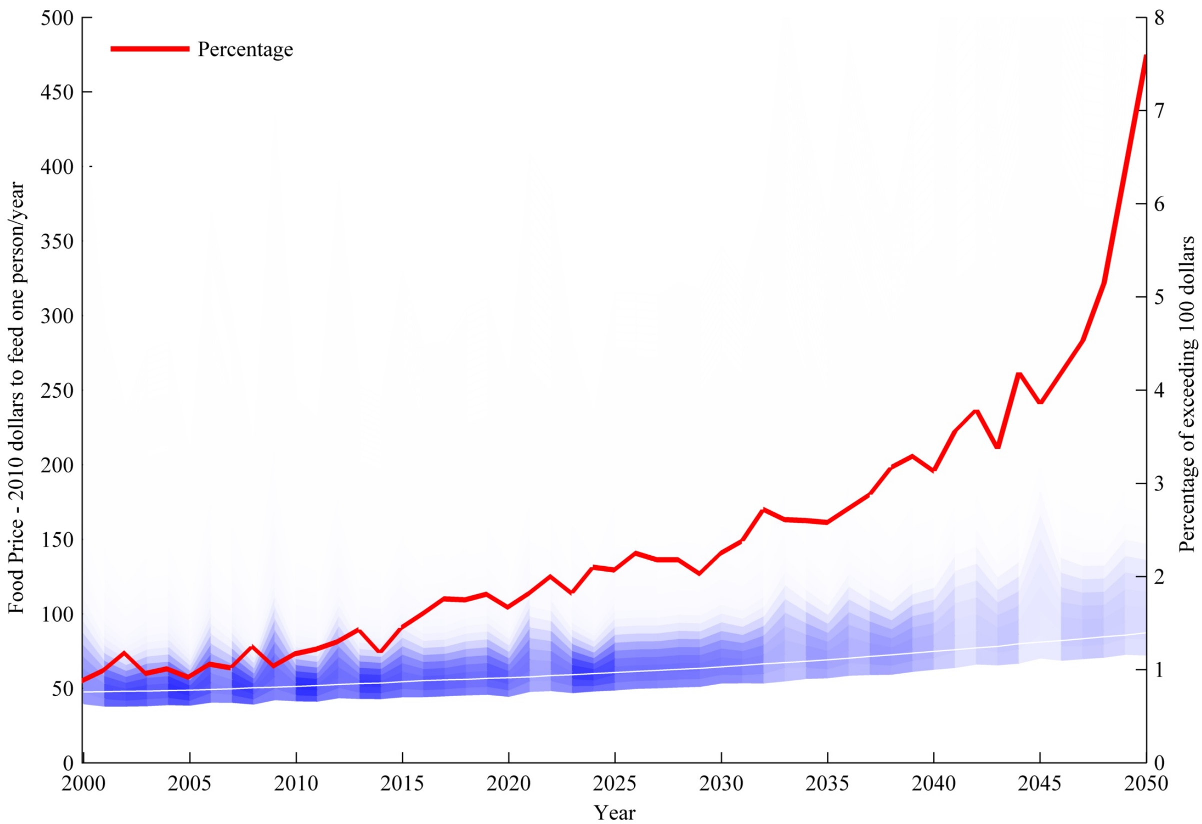

Figure 4: **Entire distribution of price simulations** This plot represents the full distribution of price forecasts. The white line represents the time path of mean price, and the darkness of the blue shading represents the density of the individual price paths around the mean. The red line represents the percent of price observations that exceed $100.

A key reason why these projections may overstate the production and price impacts of climate change is the still-ambiguous role of CO_2 on crop yields. While most scientists believe higher CO_2 concentrations will be beneficial for crop yields, the size of the effect remains controversial. We account for CO_2 by using the predicted CO_2 level under the A1B scenario found in the IPCC's Special Report on Emissions Scenarios (Nakicenovic et al. [2000]) , and use the DSSAT crop model response to CO_2 to estimate the change in crop yield for each year. We report results in Figures 5 and 6. When we account for CO_2, the yield trend slightly increases over time instead of decreasing, and yield CV increases only 14 percent between 2000 and 2050 instead of increasing 23 percent when there is no CO_2 in the model. As a

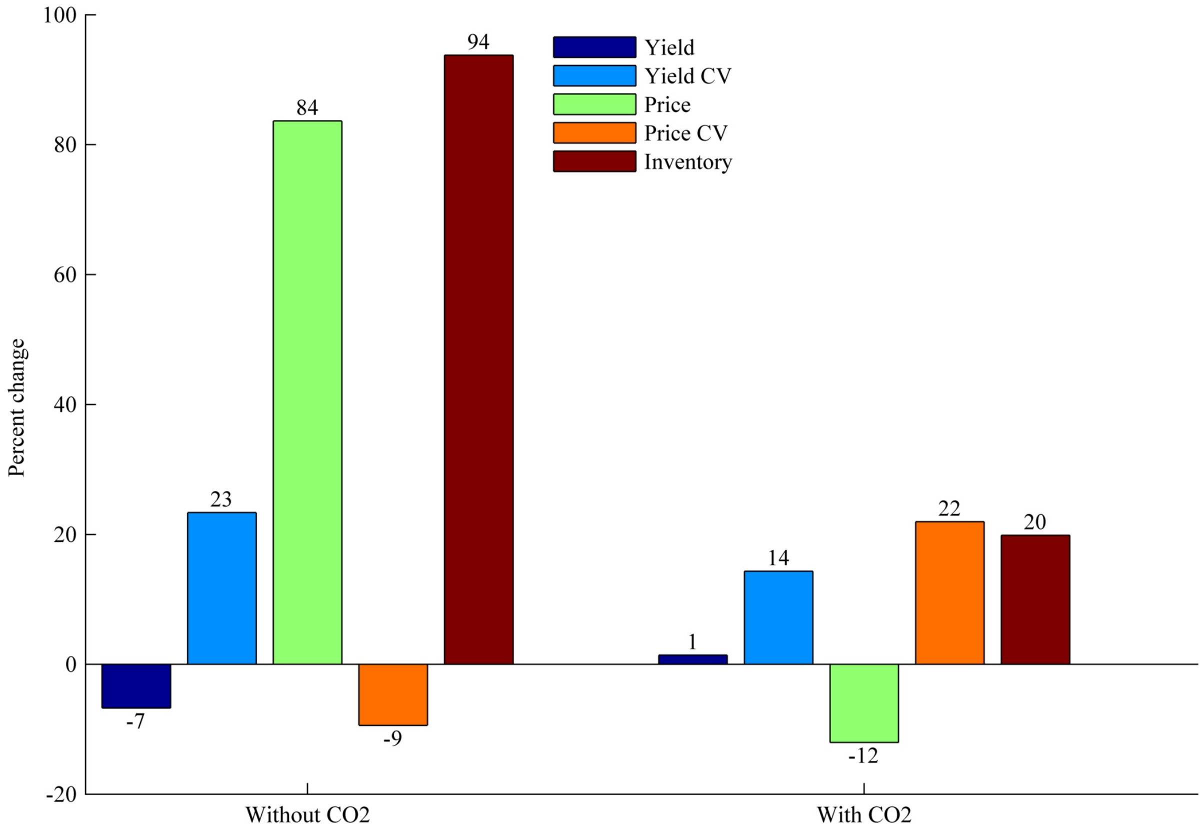

Figure 5: Comparison of percentage change in mean yield, inventory, price, and coefficient of variation of yield and price, with and without CO_2 effects

result of higher yield trend, mean price drops by 12 percent in 50 years. Similar to the case with no CO_2 effects, most of the change in mean price is caused by the difference in yield trend. In addition, because declining mean price reduces storage incentives, the increase of inventory in Figure 6 is mainly due to higher yield variability.

Table 3 summarizes welfare impacts of anticipated future climate change under various scenarios of demand and supply elasticities, interest rate, and CO_2 effect. In addition to the three cases of future yield distribution as described earlier, we add two more cases: a "best" case which uses the 95th percentile yield trend and the 5th percentile yield standard deviation, calculated over the 1600 yield series we projected; and a "worst" case which uses the 5th percentile yield trend and the 95th percentile yield standard deviation.

When there is no effect of CO_2 on crop yield, the results imply a potential loss in the

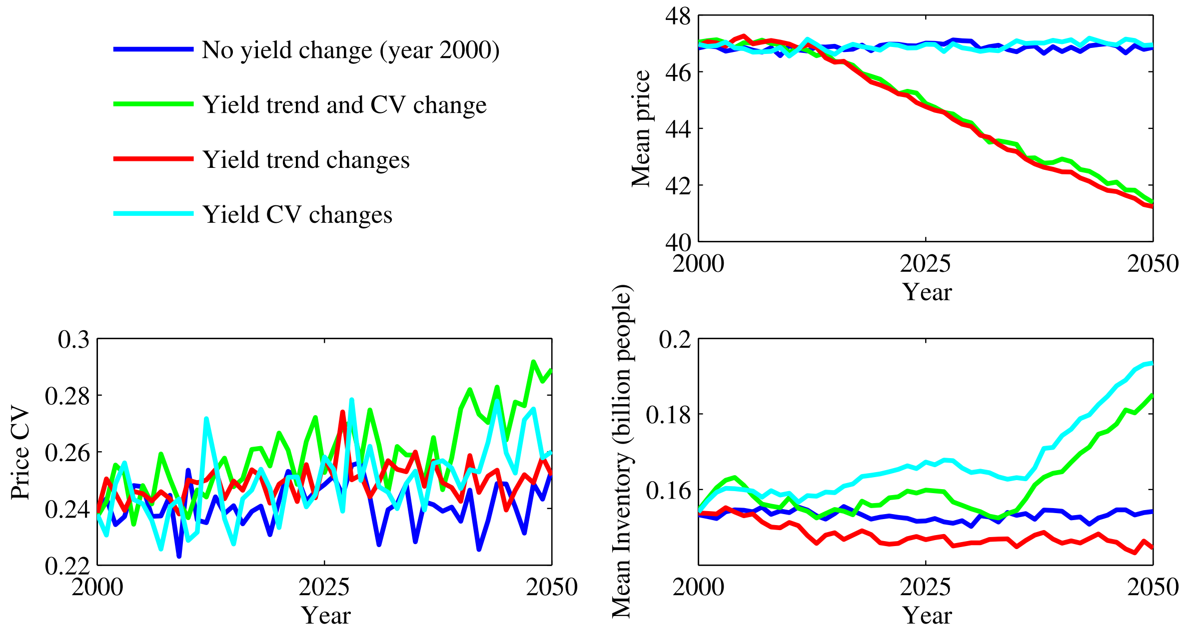

Figure 6: Data simulations under different yield distribution assumptions with CO_2 effects

range of \$0.2-58 billions by the year 2050 across scenarios. In all cases, producers' gains are outweighed by losses to consumers. Most of the changes in welfare are driven by the higher level of food prices. As noted in the introduction, standard welfare measures do not account for the more potentially damaging effects of price volatility. When we account for CO_2 effects, the lower mean price causes a rise in consumer surplus increases that exceeds to the loss to producers in scenarios 1, 2, and 5.[16] The absolute change in welfare is much smaller, ranging from \$0.1 to \$12 billions.

8 Conclusion

It remains ambiguous how much climate change has and will continue to affect agriculture. When we extrapolate historical yield-weather relationships onto projections from prominent climate models, crop yields gradually decline and become more variable. In this paper we

[16]Scenario 1-3: 1 - both trend and CV change, 2 - changing trend, 3 - changing CV. Scenario 4-5: Worst and best yield scenarios respectively.

Table 3: Changes in consumer (CS) and producer surplus (PS) between 2000 and 2050 (billion dollars) under various scenarios

Storage model parameters	No CO_2 effects		With CO_2 effects	
	ΔCS	**ΔPS**	**ΔCS**	**ΔPS**
Scenario 1				
Ed$= -.04$,**Es**$= .08$,**r**$= .02$	-247.58	219.40	33.88	-29.95
Ed$= -.04$,**Es**$= .08$,**r**$= .05$	-248.91	219.97	35.54	-31.78
Ed$= -.08$,**Es**$= .04$,**r**$= .02$	-222.76	196.84	29.36	-25.96
Scenario 2				
Ed$= -.04$,**Es**$= .08$,**r**$= .02$	-247.21	219.05	34.74	-30.84
Ed$= -.04$,**Es**$= .08$,**r**$= .05$	-245.34	217.09	35.82	-31.86
Ed$= -.08$,**Es**$= .04$,**r**$= .02$	-220.31	194.76	31.14	-27.54
Scenario 3				
Ed$= -.04$,**Es**$= .08$,**r**$= .02$	-1.87	1.66	-0.66	0.55
Ed$= -.04$,**Es**$= .08$,**r**$= .05$	-1.93	1.43	0.50	-0.72
Ed$= -.08$,**Es**$= .04$,**r**$= .02$	-1.91	1.66	-1.84	1.56
Scenario 4				
Ed$= -.04$,**Es**$= .08$,**r**$= .02$	-496.71	440.18	-83.72	73.92
Ed$= -.04$,**Es**$= .08$,**r**$= .05$	-489.10	431.41	-88.34	78.11
Ed$= -.08$,**Es**$= .04$,**r**$= .02$	-433.01	381.95	-75.69	66.81
Scenario 5				
Ed$= -.04$,**Es**$= .08$,**r**$= .02$	-92.19	81.37	106.30	-94.33
Ed$= -.04$,**Es**$= .08$,**r**$= .05$	-101.13	89.73	108.67	-96.55
Ed$= -.08$,**Es**$= .04$,**r**$= .02$	-89.45	79.15	99.87	-88.25

Note: Scenario 1-3: 1 - both trend and CV change, 2 - changing trend, 3 - changing CV. Scenario 4-5: Worst and best yield scenarios respectively.

solve a competitive storage model to derive how expectations of such changes would affect food commodity prices today and going forward into the future, assuming prices would otherwise remain stationary. These expectations imply an immediate 2% increase in price in 2000, when future climate change is first anticipated, gradually increasing to 84% by 2050. Thus, it is plausible that prices today already reflect anticipated changes in future climate.

Though yield volatility is likely to increase, we find that relative price volatility actually declines. This seemingly counter-intuitive result comes from market-induced storage adjustments in the face of a diminishing yield trend and rising yield variability. Price level and absolute price volatility still increase, however, with the probability of price exceeding $100 increasing more than seven fold by 2050. Although storage adjustments can greatly smooth the transition to a warmer world, they cannot greatly mitigate a secular decline in productivity. Our welfare analysis suggests a decrease in consumer surplus, and a smaller increase in producer surplus, indicating an overall welfare decline. This has significant distributional implications. On a household level, many net food buyers in developing countries are in the lower end of the income distribution (i.e. the urban poor); in this case, the result would be a direct transfer from the poor to the rich. On a national level, this looks like a transfer of income, and perhaps wealth, from poor food-importing countries to already rich, food-exporting countries like the U.S. This is a troubling consideration in the context of international climate negotiations.

Although we lend little credence to the specific forecasted value of prices in the distant future, we see such expectations and pricing effects as plausible. Perhaps equally plausible is the more optimistic projected trend in productivity with CO2-enhanced yields. In both projections, crop yields become significantly more variable. Our findings suggests that growing volatility of crop yields from climate change could be of little concern, however. If markets function well in anticipating climate change then adaptation in storage behavior can accommodate more volatile production at a low cost. The minimal influence of yield volatility has an important caveat: it relies on seamless trade between between regions and countries, and commodity markets that do not anticipate abrupt changes in trade or policy going forward. Recent experience with export bans and other kinds of trade restrictions in the face of production and price shocks indicates that these policies might obstruct market incentives to smooth shocks over time and space.

References

P.C. Abbott, C. Hurt, and W.E. Tyner. What's driving food prices in 2011? *Issue Report*, 2011.

Maximilian Auffhammer, V. Ramanathan, and Jeffrey R. Vincent. Integrated model shows that atmospheric brown clouds and greenhouse gases have reduced rice harvests in india. *Proceedings of the National Academy of Sciences*, 103(52), 2006.

C. Barrett and M. Bellemare. Why food price volatility doesn't matter. *Foreign Affairs*, 12, 2011.

Eugenio SA Bobenrieth, Juan RA Bobenrieth, and Brian D Wright. Bubble troubles? rational storage, mean reversion and runs in commodity prices. In *The Economics of Food Price Volatility*. University of Chicago Press, 2012.

Donna Brennan, Jeffrey Williams, and Brian D Wright. Convenience yield without the convenience: a spatial–temporal interpretation of storage under backwardation. *The Economic Journal*, 107(443):1009–1022, 1997.

Carlo Cafiero, Eugenio S.A. Bobenrieth H., Juan R.A. Bobenrieth H., and Brian D. Wright. The empirical relevance of the competitive storage model. *Journal of Econometrics*, 162 (1):44–54, May 2011.

Carlo Cafiero, Eugenio SA Bobenrieth, Juan RA Bobenrieth, and Brian D Wright. Maximum likelihood estimation of the standard commodity storage model: evidence from sugar prices. *American Journal of Agricultural Economics*, 97(1):122–136, 2015.

David Dawe and C. Peter Timmer. Why stable food prices are a good thing: Lessons from stabilizing rice prices in asia. *Global Food Security*, 1(2):127 – 133, 2012. ISSN 2211-9124. doi: 10.1016/j.gfs.2012.09.001. URL `http://www.sciencedirect.com/science/article/pii/S2211912412000132`.

A. Deaton and G. Laroque. Estimating a nonlinear rational expectations commodity price model with unobservable state variables. *Journal of Applied Econometrics*, 10(S1):S9—S40, 1995.

A. Deaton and G. Laroque. Competitive storage and commodity price dynamics. *The Journal of Political Economy*, 104(5):896—923, 1996.

Angus Deaton and Guy Laroque. On the behaviour of commodity prices. *The Review of Economic Studies*, 59(1):1 –23, January 1992.

N.S. Diffenbaugh, T.W. Hertel, M. Scherer, and M. Verma. Response of corn markets to climate volatility under alternative energy futures. *Nature Climate Change*, 2(7):514–518, 2012.

A. Dobermann, C. Witt, and D. Dawe. *Increasing productivity of intensive rice systems through site-specific nutrient management*. IRRI; Science Publishers, Inc., 2004.

Christophe Gouel and Nicolas Legrand. Estimating the competitive storage model with trending commodity prices. Technical report, University of Paris West-Nanterre la Défense, EconomiX, 2015.

Robert L. Gustafson. *Carryover levels for grains: a method for determining amounts that are optimal under specified conditions*. U.S. Dept. of Agriculture, 1958.

IPCC. Fourth Assessment Report. Technical report, Intergovernmental Panel on Climate Change, 2007. http://www1.ipcc.ch/.

Lutz Kilian and Bruce Hicks. Did unexpectedly strong economic growth cause the oil price shock of 2003–2008? *Journal of Forecasting*, 32(5):385–394, 2013.

Marco Lagi, Karla Bertrand, and Yaneer Bar-Yam. The food crises and political instability in north africa and the middle east. *Available at SSRN 1910031*, 2011.

D.B. Lobell, W. Schlenker, and J. Costa-Roberts. Climate trends and global crop production since 1980. *Science*, 333(6042):616, 2011.

B.F. Massell. Price stabilization and welfare. *The Quarterly Journal of Economics*, 83(2): 284–298, 1969.

Mario Miranda and Paul Fackler. Computational methods in economics-matlab toolbox. *File downloadable from the website http://www4. ncsu. edu/unity/users/p/pfackler/www/ECG790C*, 2007.

Nebojsa Nakicenovic, Joseph Alcamo, Gerald Davis, Bert de Vries, Joergen Fenhann, Stuart Gaffin, Kenneth Gregory, Arnulf Grubler, Tae Yong Jung, Tom Kram, et al. Special report on emissions scenarios: a special report of working group iii of the intergovernmental panel on climate change. Technical report, Pacific Northwest National Laboratory, Richland, WA (US), Environmental Molecular Sciences Laboratory (US), 2000.

Michael J. Roberts and Wolfram Schlenker. Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *American Economic Review*, 103(6):2265–95, 2013.

P.A. Samuelson. The consumer does benefit from feasible price stability. *The Quarterly Journal of Economics*, 86(3):476–493, 1972.

Jose A Scheinkman and Jack Schechtman. A simple competitive model with production and storage. *The Review of Economic Studies*, 50(3):427–441, 1983.

W. Schlenker and D.B. Lobell. Robust negative impacts of climate change on african agriculture. *Environmental Research Letters*, 5:014010, 2010.

Wolfram Schlenker and Michael J. Roberts. Nonlinear temperature effects indicate severe damages to u.s. crop yields under climate change. *Proceedings of the National Academy of Sciences*, 106(37), 2009.

Ronald Trostle. Global agricultural supply and demand: Factors contributing to the recent increase in food commodity prices (revised). Technical report, Economic Research Service, July 2008.

UNCTAD. *Trade and Development Report 2009*. UNITED NATIONS, New York and Geneva, 2009.

D. Urban, M.J. Roberts, W. Schlenker, and D.B. Lobell. Projected temperature changes indicate significant increase in interannual variability of us maize yields. *Climatic change*, pages 1–9, 2012.

F.V. Waugh. Does the consumer benefit from price instability? *The Quarterly Journal of Economics*, 58(4):602–614, 1944.

Jeffrey C. Williams and Brian Wright. *Storage and commodity markets*. Cambridge University Press, 1991.

L. Williamson and P. Williamson. What we eat. *Journal of Farm Economics*, 24(3):698–703, 1942.