On agricultural commodities' extreme price risk

Maarten R. C. van Oordt[1] · Philip A. Stork[2,3] ⓘ · Casper G. de Vries[3,4]

Abstract

We show how fat tails in agricultural commodity returns arise endogenously from productivity shocks in a standard macroeconomic model. Using nearly ninety years of data, we show that the eight agricultural commodities in our sample exhibit fat-tailed return distributions. Statistical tests confirm the heavy-tailedness of price spikes for agricultural commodities. We apply extreme value theory to estimate the size and likelihood of price spikes in agricultural commodities. Back-testing verifies the validity of our risk assessment methodology.

Keywords Commodity prices · Extreme value theory · Heavy tails · Risk management

AMS 2000 Subject Classifications 60G70 · 62G32

1 Introduction

The severe drought in the US during the summer of 2012 coincided with price increases of corn, soybeans and other field crops by more than 50%. Figure 1 provides an illustration of this period by showing the contract price for the purchase of one bushel of corn and soybeans to be delivered at the end of 2012. Prices remained relatively stable during the first few months of 2012, while the level of precipitation in the Primary Corn and Soybean Belt was not much different from its level in

✉ Philip A. Stork
p.a.stork@vu.nl

[1] Bank of Canada, Funds Management and Banking Department, 234 Wellington Ave, Ottawa, ON, K1A 0G9, Canada

[2] School of Business and Economics, Vrije Universiteit Amsterdam, Boelelaan 1105, 1081 HV, Amsterdam, The Netherlands

[3] Tinbergen Institute, Gustav Mahlerlaan 117, 1082 MS, Amsterdam, The Netherlands

[4] Erasmus School of Economics, Erasmus University Rotterdam, Burgemeester Oudlaan 50, 3062 PA, Rotterdam, The Netherlands
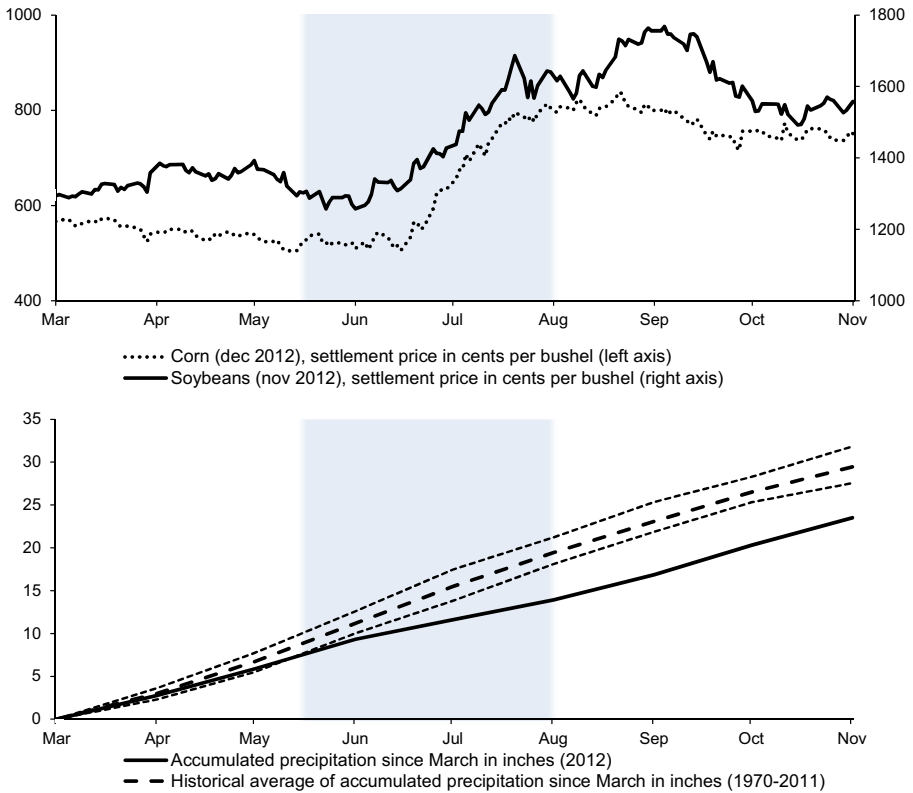
Fig. 1 Futures prices and precipitation in 2012. Note: The upper panel shows prices of futures contracts for delivery at the end of 2012 from Thomson Reuters Datastream. The vertical grey area indicates a period of severe drought. The lower panel shows monthly area-weighted precipitation data for the Primary Corn and Soybean Belt. The bandwidth around the average corresponds to the historical lower and upper quartile. Precipitation data come from the U.S. Climate Divisional Database and are obtained from the National Climatic Data Center of the National Oceanic and Atmospheric Administration

previous years. June and July were however exceptionally dry months (bottom panel of Fig. 1).[1] As this drought was prolonged and growing in severity, prices increased rapidly. Price levels flattened only once the level of rainfall returned to its historical average in the late summer.

Such extreme movements in agricultural commodity prices are anything but uncommon. For instance, between August 2007 and March 2008, the price of wheat almost doubled, but before the end of 2008, the wheat price returned to its original level. Another example is the price of corn, which fell a massive 55% in the

---

[1]The National Climatic Data Center uses the expression Primary Corn and Soybean Belt to specify an agricultural belt around Illinois, Indiana and Iowa. The area-weighted precipitation in the Primary Corn and Soybean Belt in June and July 2012 was 4.6 inches (cumulative), or almost half the historical average of 8.3 inches between 1970 and 2011. Since 1895, such low levels of precipitation in June and July were only recorded twice before, in 1936 and 1988.

second half of 2008 alone. These examples illustrate the highly volatile behavior of agricultural commodity prices.

The extreme price movements raise questions regarding the shape of the tail distribution of changes in agricultural commodity prices. Mandelbrot (1963a, b), using Houthakker's cotton price series, was the first to suggest that the tails of the distribution of commodity price changes diminish by a power instead of an exponent, as is the case under the more common (log)normal assumption.[2] If the tail of a distribution diminishes by a power, then the probability of variable $\tilde{x}$ exceeding threshold u, if u is large, is distributed as

$$\Pr(\tilde{x} > u) \sim Cu^{-\alpha}, \tag{1}$$

where $C > 0$ and $\alpha > 0$ are referred to as, respectively, the *scale* and the *shape* parameter.[3] The distribution is named after Pareto who discovered that the upper tail of the income distribution follows a power law. Distributions with tails that obey the functional form in Eq. 1 are classified as *heavy-tailed.* Tails which follow a power law are in the end always fatter than *thin-tailed* distributions that decrease at an exponential rate.

Mandelbrot explained the origin of the heavy-tailedness of agricultural commodity prices by advancing that the physical world is full of heavy-tailed phenomena.[4] Such exogenous heavy-tailed shocks could trigger the heavy-tailedness of commodity price changes. But how a power law tail may arise endogenously in the market for agricultural commodities has not been formally investigated.

The main contribution of this article is twofold.

First, we show how fat tails in agricultural commodity prices may arise endogenously in a standard macroeconomic model. In the model, commodity price spikes occur as a result of adverse productivity shocks, due to, e.g., hurricanes, diseases and droughts.[5] In themselves these productivity shocks need not be (although they may be) heavy-tailed. This is the extra kick that our economic analysis provides. We show how the power law spikes observed in agricultural commodity prices can arise endogenously in the economy as productivity shocks feed through the system, even if the productivity shocks come from a thin-tailed distribution or a distribution with a bounded support.

Second, we provide new empirical evidence on how to use extreme value theory (EVT) to measure the extreme price risk of agricultural commodities using nearly ninety years of data. Differently from previous studies, we employ formal statistical tests to confirm the heavy-tailedness of price spikes for agricultural commodities. We use back-testing procedures to provide empirical evidence on the accuracy of

---

[2]Mills (1927) was one of the first to discuss the non-normality of commodity returns as he reported higher kurtosis, implying more extreme returns.

[3]If $x \sim y$, it means that $\lim x/y = 1$.

[4]See, e.g., Newman (2005) and Salvadori et al. (2007) for a number of natural hazards that follow a power law distribution, including various drought measures, flood levels, the magnitudes of earthquakes and the scale of wars. Several of the above events influence agricultural prices in one way or another.

[5]Knittel and Pindyck (2016) show that fundamental factors are important in determining commodity prices, while speculation tends to have relatively little effect.

the risk measures. We show that the non-normality of the return distribution strongly influences the level of the risk measures. Our empirical estimates provide a good indication of the size of the risks as measured by widely used and easily understood risk measures.

What lays at the origin of heavy-tailed phenomena in the economy is an intriguing question. The literature has proposed several mechanisms that can explain some of the heavy-tailed phenomena in the economy (for a literature review, see Gabaix, 2009). To provide a few examples: Gabaix et al. (2003) and Gabaix et al. (2006) suggest a mechanism where trading by large institutional investors – whose size distribution follows a power law tail – could be responsible for heavy-tailed stock market returns. Reed (2001) suggests that the power law in the cross-sectional distribution of the size of agglomerations, i.e., "Zipf's Law", could be the result of a process where each agglomeration grows with geometric Brownian motion and where the time of observation itself is a random variable that follows an exponential distribution. Toda (2014) and Toda and Walsh (2015) show in a general equilibrium model that this principle could also be responsible for the power law in the cross-sectional distributions for wealth and consumption. Similarly, our paper suggests a mechanism that can explain the origin of heavy tails in agricultural commodity prices.

A good understanding of the most extreme agricultural commodity returns is instrumental in any commodity risk management application. Strong agricultural commodity price movements hurt many of the industry's stakeholders, including producers, processing firms, traders, consumers and institutional investors. Commodities are often included in investors' portfolios because of their diversification benefits and inflation-hedging potential; see, e.g., Erb and Harvey (2006) and Gorton et al. (2012). The economic prosperity of both importing and exporting developing countries often depends on the price development of raw material commodities. More than fifty countries depend on only three or fewer commodities for more than half their total exports.[6] As a result these countries are very vulnerable to price volatility, see, e.g., Deaton (1999) or Balagtas and Holt (2009). The producers, processing industry and traders can be hurt when prices spike, while farmers can be highly vulnerable to commodity price decreases, which translate directly into a loss of income. Extreme commodity price swings may also hurt large institutional investors. Pension funds, hedge funds like CTAs, and fund-of-funds are known to invest heavily in commodity futures contracts. The risk management findings in this study are relevant for these affected parties.

Market participants need a risk quantification methodology to decide upon the optimal strategy. They should be able to answer questions like: How likely is a 10% drop in the corn price over the next week? Suppose my risk tolerance allows me to accept a 25% price increase, how likely is it that the wheat price will exceed this threshold next month? Is either the corn or the wheat price more likely to experience

---

[6]Based on the UNCTAD 1995 Commodity Yearbook. We refer to Bidarkota and Crucini (2000) for an extensive analysis of the relationship between the terms of trade of developing nations and world prices of internationally traded primary commodities.

extreme price movements? What is the expected size of the maximum loss due to price risk during the next decade?

Such questions may be addressed by the use of EVT. The popularity of EVT to assess the risk of an extreme event has increased considerably. For example, EVT has been used to examine the severity of stock market crashes, the pricing of catastrophic loss risk in reinsurance or the extent of operational risk in banks (ECB 2006). EVT is particularly suitable for analyzing rare events when sample sizes are too small for determining the probability, extent or cause of the extreme returns using conventional statistical techniques. The semi-parametric EVT approach exploits the functional regularities that probability distributions necessarily display far from the center.

Interestingly, in spite of its growing recognition, application of EVT in agricultural price risk management has so far been sparse in the academic literature.[7] Kofman and De Vries (1990) estimate the tail parameters for potato futures. Matia et al. (2002) estimate the tail parameters of a large number of general commodities. Their article provides no risk management applications, however. Krehbiel and Adkins (2005) apply EVT to four complex NYSE energy futures contracts to estimate various risk measures. Even so, their analysis is limited to oil and gas contracts, whose return distribution may be very different from those of renewable agricultural commodities. More recently, Morgan et al. (2012) use EVT on weekly data to estimate three different tail risk measures for corn and soybeans. Their thorough study is evidence of the growing interest in this topic.

2 Theory

We use a standard off-the-shelf two-sector macro model to describe how fat tails in agricultural commodity prices may arise endogenously. The economy in this model consists of households and firms. The households maximize their utility subject to their budget constraint by purchasing goods from firms while choosing their labor supply to firms. The firms need labor in order to produce goods. Firms operate in the agricultural or the differentiated goods sector. The agricultural sector is modeled as a competitive sector (Ardeni and Freebairn 2002). The differentiated goods sector is an environment with monopolistic competition in the spirit of Dixit and Stiglitz (1977). Exogenous shocks affect the firms' productivity in both sectors, where productivity is defined as the number of goods that a firm can produce per unit of labor. In the agricultural sector, these shocks can be best thought of as changes in weather and other natural hazards. For the differentiated goods sector, which also captures the services industry, the shocks mostly represent changes in productivity due to innovations.

The macro literature has focused almost exclusively on the Dixit and Stiglitz (1977) specification for the differentiated goods demand, see, e.g., Walsh (2017). The

---

[7]We refer to Kat and Oomen (2006), and Wang and Tomek (2007) for thorough studies of the general time series properties of agricultural commodity prices.

familiar Dixit and Stiglitz (1977) specification with endogenous labor supply derives from the following utility function for households

$$U = Z^{1-\theta} \left[\frac{1}{n} \sum_{i=1}^{n} Q_i^{\rho} \right]^{\theta/\rho} - \frac{1}{1+\delta} L^{1+\delta}, \tag{2}$$

where Z is the consumed quantity of the competitive good, the Q_is are the consumed quantities of the differentiated goods and L is quantity of labor. Essentially, all goods are perishable in this one-period model.[8] The parameter ρ, which is constrained to $\rho \in (0, 1)$, determines the "love of variety" of households, which is the extent to which households value diversity of the products they purchase from firms in the differentiated goods sector. Values of ρ closer to zero are associated with a stronger love of variety. Parameter $\theta \in (0, 1)$ determines the importance of the other goods and services relative to the agricultural produce in the consumer's consumption bundle. The higher the level of θ, the smaller the share of income the consumer is willing to spend on agricultural goods. The parameter δ is the inverse of the Frisch (1959) elasticity of labor supply, $\delta \in (0, 1)$. The higher the level of δ, the less responsive labor supply will be to changes in wages.

Households maximize their utility subject to their budget constraint. The budget constraint of the households reads

$$wL + \Pi(Q) = qZ + \frac{1}{n} \sum_{i=1}^{n} p_i Q_i, \tag{3}$$

where w is the wage rate, where q and p_i are the goods prices, and where $\Pi(Q)$ are the profits of firms in the differentiated goods sector which are distributed as dividends to the households.[9]

The production of firms in both sectors depends on their productivity and labor inputs. For their production functions we assume classical Ricardian technologies (e.g., Turnovsky 1974), where

$$Z = BN, \tag{4}$$

and

$$Q_i = AN_i. \tag{5}$$

The N and N_i are the labor inputs. The A and B are the respective productivity coefficients that determine how many goods the firms in both sectors can produce per unit of labor. Both A and B are random variables. Shocks to A and B are productivity

---

[8]Our objective is to analytically derive the tail distribution of agricultural commodity prices. We do not explicitly model the role of storage. The reason is that specifications of models with storage generally do not allow for analytical solutions due to the introduction of nonnegativity constraints; see, e.g., Wright and Williams (1982) and Williams and Wright (2005). An exception is the model of Aiyagari et al. (1989), but their model imposes a finite endpoint for the statistical distribution of the commodity prices by assuming an upper bound (see p. 40). This upper bound rules out the possibility of a power law tail ex ante. As a consequence, our model provides a description of the price distribution of perishable goods, rather than that of commodities that decay at a very low rate.

[9]The quantities of the differentiated goods, the Q_is, are normalized by the number of differentiated goods, n. This notation is analogous to the common continuous good notation often used in the theoretical macro literature.

shocks in the sense that changes in A and B translate into identical changes in the amount of goods a firm can produce per unit of labor. In the case of variable A, which measures productivity in the differentiated goods sector, such shocks are the in macroeconomics familiar supply side total factor productivity shocks. The shocks to variable B in the agricultural sector capture the shocks to agricultural productivity that are inherent to nature.

Firms in the economy maximize profits. The market for the agricultural product is assumed to be perfectly competitive, and, hence, the aggregate profit of all firms in the agricultural sector,

$$\Pi_\theta(Z) = qZ - wN, \tag{6}$$

equals zero in equilibrium. The producers of the differentiated products aim to maximize their profit functions

$$\Pi_i(Q_i) = p_i Q_i - wN_i. \tag{7}$$

They do so by exploiting their pricing power while ignoring their pricing effect on consumer income $wL + \Pi(Q)$ and on the price index of the differentiated goods,

$$P = \left(\frac{1}{n} \sum_{i=1}^{n} p_i^{\rho/(\rho-1)}\right)^{\frac{\rho-1}{\rho}}. \tag{8}$$

Finally, to determine the general price level we assume a simple quantity-type relationship (Fisher 1911),

$$M = wL, \tag{9}$$

where M is defined as the total quantity of money.

2.1 Equilibrium price distribution

With the above preparations, we can now obtain the implications for the equilibrium prices.

Proposition 1 *Equilibrium prices result when consumers maximize utility* (2) *and firms in both sectors maximize profits. The equilibrium prices of the differentiated goods are,*

$$p_i = M \frac{1/\rho^{\theta/\delta+1}}{A \left(\theta^\theta (1-\theta)^{1-\theta} A^\theta B^{1-\theta}\right)^{1/\delta}} \text{ for all } i. \tag{10}$$

For the agricultural good the equilibrium price is

$$q = M \frac{1/\rho^{\theta/\delta}}{B \left(\theta^\theta (1-\theta)^{1-\theta} A^\theta B^{1-\theta}\right)^{1/\delta}}. \tag{11}$$

Proof See Appendix A. □

Most macro models consider shocks to M, A and B. Let us focus on the agricultural productivity shocks due to randomness in B.[10] Assuming M and A to be constant, we can write the price of the agricultural good as

$$q(B) = \Theta B^{-\frac{1+\delta-\theta}{\delta}}, \tag{12}$$

where

$$\Theta = M \frac{1/\rho^{\theta/\delta}}{\left(\theta^\theta (1-\theta)^{1-\theta} A^\theta\right)^{1/\delta}}.$$

Note that the power in Eq. 12 is negative. From Eqs. 10 and 11 it transpires that a positive productivity shock to A or B lowers the equilibrium prices p_i and q. This is because when more produce comes to the market, this extra supply can only be absorbed if prices are lowered. This reflects what happens in both well-known real business cycle models and the now standard New Keynesian Model that is the backbone to most current day macro models.[11]

For illustrative purposes, consider the situation where B follows a beta distribution (we relax this assumption later):

$$\Pr\{B < t\} = t^\beta \tag{13}$$

on [0, 1] and $\beta > 0$. The likelihood of extremely adverse productivity shocks under this distribution is lower for higher values of β. Under the assumption of the beta distribution, we can derive the price distribution of the agricultural product as follows. Denote the randomness in q by $\tilde{q}$. Then

$$
\begin{aligned}
\Pr\{\tilde{q} > u\} \quad &= \Pr\left\{\Theta B^{-\frac{1+\delta-\theta}{\delta}} > u\right\} \\
&= \Pr\left\{B < \Theta^{\frac{\delta}{1+\delta-\theta}} u^{-\frac{\delta}{1+\delta-\theta}}\right\} \\
&= \Theta^{\frac{\beta\delta}{1+\delta-\theta}} u^{-\frac{\beta\delta}{1+\delta-\theta}},
\end{aligned}
\tag{14}
$$

with support on $[\Theta, \infty)$. The distribution of equilibrium prices in Eq. 14 has the same functional form as the heavy-tailed distribution in Eq. 1, where $\alpha = \beta\delta/(1 + \delta - \theta)$.

We proceed by relaxing the assumption of the beta distribution. More specific, we derive a general condition for the density function of the productivity shocks $f_B(t)$. Any density function of the productivity shocks that satisfies this condition will imply

---

[10]For our results to hold true, it is not necessary to assume a constant A and M. For example, the heavy-tailedness of the equilibrium price distribution due to natural shocks is preserved if the productivity of the differentiated sector does not collapse completely, which implies that the support of A is bounded away from zero. Further, the heavy-tailedness of the equilibrium price distribution is not affected if the distribution of M has exponential tails, such as the lognormal distribution.

[11]In his popular textbook, Gali (2008) reviews the ins and outs of this model that consists of an aggregate demand and supply equation and a monetary policy rule, which together determine the output gap (output relative to what is optimally attainable) and the inflation rate. A positive shock to productivity leads unambiguously to a decline in inflation (Gali 2008, Subsections 3.4.1.2 and 4.4.1), as in our setting.

a equilibrium price distribution that is heavy-tailed. Let $\mathcal{L}(u)$ denote a slowly varying function, in the sense that

$$\lim_{u \to \infty} \frac{\mathcal{L}(tu)}{\mathcal{L}(u)} = 1 \ \text{ for any } \ t > 0.$$

Then the following proposition provides a general condition that results in the heavy-tailedness of the commodity price distributions. The essential requirement is that zero is in the support of B.

Proposition 2 *Suppose that the distribution and density of the agricultural productivity coefficient B are continuous. Given the price-productivity relationship for agricultural produce in Eq. 12, we have that*

$$\Pr(\tilde{q} > u) \sim \mathcal{L}(u)u^{-\alpha} \ as \ u \to \infty, \tag{15}$$

with

$$\alpha = \xi \frac{\delta}{1 + \delta - \theta}, \tag{16}$$

if

$$\lim_{s \downarrow 0} \frac{w f_B(sw)}{f_B(s)} = w^{\xi} \ with \ \xi \in \mathbb{R}^+. \tag{17}$$

Proof See Appendix B. □

Whether the condition on the distribution of B in Eq. 17 is satisfied depends on the shape of the density of the productivity shocks, $f_B(B)$, for values close to zero. The reason is that high price levels occur in periods of low agricultural productivity, such as severe droughts. The tail distribution for those low levels of productivity determine the shape of the distribution function for extremely high prices.

Corollary 1 *The shape of the distribution of the agricultural productivity coefficient B close to zero determines whether the equilibrium price distribution is heavy-tailed. The more slowly the density of B converges to zero for extremely low productivity levels, i.e., the lower ξ, the heavier is the tail of the agricultural price distribution, i.e., the lower is α.*

The condition in Proposition 2 holds true with different values for ξ for a broad range of distribution functions with positive support. For the beta distribution, which we considered for illustrative purposes, the condition in Eq. 17 is satisfied with $\xi = \beta$. The condition also holds true for distribution functions that follow very different functional forms. For instance, the standard uniform distribution and standard exponential distribution satisfy the condition in Eq. 17 with $\xi = 1$. More general, the distribution function $\Pr(B < u) = 1 - \exp(-\lambda x^{\beta})$ with $\lambda, \beta > 0$ satisfies the condition with $\xi = \beta$. Other popular distribution functions that satisfy the condition in Eq. 17 are the Chi-squared distribution with l degrees of freedom, which satisfies the condition with $\xi = l/2$; the Gamma distribution with shape parameter l, which satisfies the condition with $\xi = l$; and the (heavy-tailed) Burr (Type XII) distribution with

parameters (c, l), which satisfies the condition with $\xi = c$. All the aforementioned distributions would result in a heavy-tailed equilibrium price distribution.

The condition also shows that not every possible distribution does lead to heavy-tailed prices in our macro-economic framework. An example of a popular exception is the lognormal distribution: Its limit in Eq. 17 converges to 0.[12] Seriously low levels of agricultural productivity are too rare under the lognormal distribution to generate a heavy tail among the occurrences of high agricultural prices in the model.

As follows from Proposition 2, the shape parameter of the tail of the distribution of agricultural prices not only depends on the distribution of productivity shocks, but also on the preference parameters θ and δ.

Corollary 2 *The greater the share of agricultural produce in consumption, i.e., the higher $1 - \theta$, and the higher the elasticity of labor supply, i.e., the higher $1/\delta$, the heavier is the tail of the distribution of agricultural prices, i.e., the lower is α.*

Given the distribution function of productivity shocks, it follows from Eq. 16 that a high value of $1 - \theta$ results in a low shape parameter of the equilibrium price distribution of agricultural goods, α, and hence in a fatter tail. This finding has the following intuition. The importance of the share of the agricultural good in the consumption bundle of the agents is represented by $1 - \theta$, see the utility function in Eq. 2. The larger the role of the agricultural good for the agents' utility, the more extreme price reactions one may expect if supply falls. This is reflected in a fatter tail of the equilibrium price distribution, i.e., a lower α.

It also follows from Eq. 16 that a high value of parameter δ results in a high value of shape parameter α. Adverse productivity shocks (low realizations of B) have a dual effect on the output of the competitive sector. First, given the amount of labor used, an adverse technology shock in the competitive sector directly reduces output, since $Z = BN$. Second, low productivity decreases the equilibrium amount of labor used in the competitive sector, which further reduces output. In particular, the equilibrium labor supply $L = cB^{(1-\theta)/\delta}$, where $c > 0$ is a constant (see Eq. 31 in Appendix A). Therefore, in equilibrium, the change in the production of the competitive good in response to a shock to B is smaller for higher values of parameter δ (lower elasticity of the labor supply), which results in thinner tails of the equilibrium price distribution, i.e., a higher α.

The property of the model that allows the equilibrium price distribution to be heavy-tailed, even if the productivity shocks are not heavy-tailed themselves, is that the equilibrium prices respond in a nonlinear fashion to shortages due to adverse productivity shocks. More precisely, the price of the agricultural product in our model is inversely proportional to the productivity. So, increasingly larger shortages of the agricultural product have an increasingly larger impact on its price in absolute terms. The property that product prices are inversely proportional to productivity is not unique to the environment we study above, and, in fact, a property that is very common to other economic models. Heavy-tailed price distributions can be obtained in

---

[12] The statistical distribution of crop yields has been the topic of a wide body of literature, see, e.g., Ramírez et al. (2003), Harri et al. (2009) and Koundouri and Kourogenis (2011).

those other models in a very similar manner as in the stylized model above. For various reasons, however, price distributions may not be heavy-tailed in a macroeconomic model. For example, one reason in our model would be that productivity shocks are bounded away from zero. In some macroeconomic models, prices are a linear though negative function of productivity, in which case prices would only be heavy-tailed if productivity shocks are heavy-tailed distributed. Finally, the relationship between α and ξ in Eq. 16 suggests a way to calibrate the value for the tail index ξ of agricultural productivity.[13]

3 Empirical methodology

The previous section discussed the plausibility of fitting a power law to the right tail distribution of changes in food prices. Section 3.1 discusses how we test whether the right tail of the distribution of agricultural commodity returns follows a power law. Section 3.2 explains how we fit a power law and estimate risk measures under the assumption that returns are heavy-tailed, while Section 3.3 discusses how we estimate the risk without assuming heavy tails.

Next, we apply EVT to determine the parameters of the power law.

3.1 Testing the power law

We start by testing whether the right tail of the distribution of agricultural commodity returns follows the power law in Eq. 1 against thin-tailed alternatives. We do so by using the generalized extreme value distribution. This distribution describes the behaviour of the largest observation in a sufficiently large number of randomly observed draws. Let $\tilde{y}_m$ denote the maximum of m i.i.d. observations of $\tilde{x}$. The distribution of $\tilde{x}$ is said to be in the domain of attraction of the generalized extreme value distribution, if, for some normalizing sequences of constants $\mu_m \in \mathbb{R}$ and $\sigma_m > 0$, and $\gamma \in \mathbb{R}$, the limit probability

$$\lim_{m \to \infty} \Pr\left(\frac{\tilde{y}_m - \mu_m}{\sigma_m} \geq s\right) = G_\gamma(s), \tag{18}$$

where

$$G_\gamma(s) = \exp(-(1 + \gamma s)^{-1/\gamma}). \tag{19}$$

One distinguishes three types of extreme value distributions depending on whether $\gamma < 0$, $\gamma > 0$ or $\gamma = 0$ (where $(1 + \gamma s)^{-1/\gamma}$ is read as e^{-s} for $\gamma = 0$). If $\gamma < 0$, then the maximum possible value of $\tilde{y}_m$ is finite.

---

[13]Reichling and Wahlen (2012) conclude in their literature review that macroeconomic estimates of the Frisch elasticity typically fall in a range from 2 to 4, which implies a value of δ in the range from 1/4 to 1/2. Moreover, the gross output of farms as a percentage of total consumption expenditures could be used as an estimate of $1 - \theta$. Over the period 2014-18, this number was in a range between 2.8 and 3.8 percent in the United States based on data from the US Bureau of Economic Analysis, suggesting $1 - \theta \approx 0.03$. Plugging these numbers into Eq. 16 would suggest a level of ξ in a range between $\alpha/0.95$ and $\alpha/0.85$ for the United States.

By estimating parameter γ in Eq. 19, we can test whether commodity price returns are heavy-tailed (i.e., $\gamma > 0$) versus thin-tailed and other alternatives (i.e., $\gamma \leq 0$). The density of the generalized extreme value distribution resembles a power law when $\gamma > 0$

$$g_\gamma(s) = \begin{cases} (1 + \gamma s)^{-1/\gamma - 1} G_\gamma(s) & \text{if } \gamma \neq 0, \\ e^{-s} G_\gamma(s) & \text{if } \gamma = 0. \end{cases} \tag{20}$$

Balkema and De Haan (1974) establish the one-to-one relationship between the parameters in the generalized extreme value distribution and the heavy-tailedness of a distribution. All distributions for which the distribution of exceedances beyond a sufficiently large threshold converges weakly in the limit to the power law in Eq. 1 with shape parameter $\alpha > 0$, must be in the domain of attraction of the generalized extreme value distribution with $\gamma = 1/\alpha$, and vice versa (Balkema and De Haan 1974, Theorem 4). By contrast, all distributions for which the tail converges to a thin-tailed exponential distribution must be in the domain of attraction of the generalized extreme value distribution with $\gamma = 0$ (Balkema and De Haan 1974, Theorem 3).

The literature has developed estimators for γ that allow for testing whether $\gamma > 0$. Since the theoretical model is predicated on the heavy-tailedness of the upper tail of the commodity price return distribution, we first test whether the data do not reject this presumption. To this end, we first use three estimators that allow for γ to fall into all three possible classes. Once we have confidence that $\gamma > 0$, we employ the Hill estimator which only applies if $\gamma > 0$, but is more efficient. The general estimators can be somewhat ambiguous depending on how many observations are taken into account. We will therefore test for the heavy-tailedness using three different estimators: the generalized Hill estimator of Beirlant et al. (1996) ($\hat{\gamma}_{BVT}$), the moment estimator of Dekkers et al. (1989) ($\hat{\gamma}_{DEH}$) and the mixed moment estimator of Fraga Alves et al. (2009) ($\hat{\gamma}_{FGHN}$).

3.2 Estimating risk with heavy tails

After establishing the heavy-tailedness of the right tail of agricultural commodity returns, we proceed with estimating the parameters in the power law. Let a series with n discrete returns be ordered from high to low: $X_1 \geq \cdots \geq X_n$. Let k denote the number of returns considered in the tail of the distribution. Then, following Hill (1975), one can estimate the shape and the scale parameter in Eq. 1 as

$$\frac{1}{\hat{\alpha}} = \hat{\gamma}_H = \frac{1}{k} \sum_{j=1}^{k} \ln \frac{X_j}{X_{k+1}} \tag{21}$$

and

$$\hat{C} = \frac{k}{n} X_{k+1}^{\hat{\alpha}}. \tag{22}$$

One can easily convert the estimated parameters into economically meaningful risk measures, such as Value-at-Risk (VaR). The VaR is defined as the loss level that is exceeded with probability p, that is, for return R_t, $VaR(p) = -\sup\{c : \Pr(R_t \leq$

$c) \le p)\}$. Using the Hill-estimators in Eqs. 21–22 in the Pareto distribution in Eq. 1 gives the VaR-estimator for the right tail as

$$\widehat{VaR}_H(p) = X_{k+1} \left(\frac{k}{np} \right)^{1/\hat{\alpha}}. \tag{23}$$

The VaR plays an important role in the safety-first framework developed by Roy (1952) and Telser (1955). Agents with the safety-first principle of Telser (1955) in their utility functions maximize their expected return, while limiting the probability that a loss larger than some disaster level occurs at some admissible level p.

3.3 Estimating risk without assuming heavy tails

The theory in Section 2 suggests why the right tail of the distribution for commodity prices follows a power law. Different from the right tail of the distribution, the left tail of the distribution has a theoretical endpoint, provided that the price of an agricultural commodity cannot fall below zero. One could therefore object to using a power law, which has an unbounded support, for modelling the tail risk in the left tail. Since the tail risk in the left tail is important for several market participants (e.g., farmers, exporters), we therefore consider two different methods to estimate the downside tail risk.

Our first method is to simply estimate the downside tail risk as if the left tail follows a power law. The basis for following this approach is a Taylor approximation, by a power function; which may work well at medium to higher loss levels. Theoretically, one may expect the risk estimates under the assumption of a power law to be increasingly upward biased in situations where changes in prices get closer to the maximum possible price decline. However, the method could still be useful from a pragmatic point of view as long as it is supported by empirical evidence that the method provides accurate risk estimates.

Our second method is to estimate the VaR under a more general distribution that allows tails with finite endpoints. In particular, we estimate parameter γ in Eq. 19 with an estimator that allows for all three cases $\gamma > 0$, $\gamma = 0$ and $\gamma < 0$. Let $\hat{\gamma}_G$ denote a consistent general estimator of γ. Then the estimate $\hat{\gamma}_G$ can be used in the general quantile estimator developed by Dekkers et al. (1989) to estimate the VaR as

$$\widehat{VaR}_G(p) = \left[X_{k+1} \left(\frac{k}{np} \right)^{\hat{\gamma}_G} - X_{k+1} \right] \times \frac{\hat{\gamma}_H}{\hat{\gamma}_G \min\{1, 1/|1 - \hat{\gamma}_G|\}} + X_{k+1}. \tag{24}$$

Note that this general estimator for the VaR condenses to the specific VaR estimator that assumes heavy tails in Eq. 23 whenever $\hat{\gamma}_G = \hat{\gamma}_H$. In our application using the general quantile estimator in Eq. 24, we report VaR estimates based on the general estimator of Fraga Alves et al. (2009), but the estimators of Dekkers et al. (1989) and Beirlant et al. (1996) generate qualitatively similar results.

We compare the performance of the VaR estimator assuming heavy tails and the general VaR estimator in the back-testing procedure in Section 8.

3.4 Other estimation details

We obtain standard errors of the shape parameters and risk measures from a bootstrap procedure. The asymptotic normality of the Hill estimator also holds in the presence of serial dependence; see, e.g., Drees (2008). Following Hartmann et al. (2006), we refrain from assumptions on the specific dependence structure and apply a bootstrap procedure with fixed block length and 10,000 replications. Following Hall et al. (1995), we set the optimal block length equal to $n^{1/3}$. We aim to choose a threshold parameter k that balances the bias and variance of the estimator as usual (Drees et al. 2000).

4 Data

For our purposes, we rely on two different categories of price series. The first series concerns monthly prices of agricultural commodities in the spot market going back to 1928. The second series concerns weekly returns on contracts traded on the futures market from 1979 onward. Our estimates for weekly returns on futures contracts are in particular interesting from a risk management perspective. The lower frequency of the series with spot prices is better aligned with the macroeconomic model in which power law spikes observed in agricultural commodity prices may arise endogenously in the economy. For each price series, we calculate discrete returns from the observed prices as $R_t = (P_t - P_{t-1})/P_{t-1}$.

The monthly spot prices from 1928 to 2014 are obtained from Global Financial Data. These series are likely to be impacted by price controls, resulting in a lower level of volatility during the World War II period. Price controls were effective since May 1942, and, although most were removed in 1946, they may have had some impact on prices until 1948Q2; see Evans (1982). To avoid these price controls from impacting our estimates, we removed all observations from 1942M1 until 1948M6 from our dataset. To account for potential illiquidity in spot markets impacting our estimates, we removed from each spot price series the observations in months without price changes and those in months following a month without price change. This resulted in the removal of on average 47 observations in each series.

Futures contracts for delivery at a particular date are usually traded for a relatively short period, ranging from several months to several years. To obtain long-term futures returns series, or so-called *continuous* series, we combine consecutive data from several futures contracts, see, e.g., De Roon et al. (2000). We take considerable effort to construct high-quality continuous futures return series. Our procedure is as follows. First, we download daily open interest and price series of all available futures contracts from Thomson Reuters Datastream for each commodity. Those time series are available over a period of 34 years: from January 1979 until December 2012. Subsequently, daily returns are calculated for all futures price series. Finally, we construct the continuous futures returns series from the individual return series. In January 1979 we start with the futures contract that has the largest open interest. For each day we include its returns in the new continuous series until six weeks before the contract's last trading day. At this date we switch to the futures contract with the

largest open interest and a later last trading day. Again we include the returns until six weeks before the last trading day and repeat the last step. This procedure results for each commodity in a daily continuous futures returns series, which is then converted into a weekly returns series of 1,773 observations from, on average, 164 different futures contracts.

Our method has an important advantage compared to Datastream's procedure to construct continuous futures series. By calculating returns prior to constructing the continuous series, no returns are calculated over price observations from two different futures series. Therefore, our series represents the return that investors could achieve by rolling over futures contracts as opposed to the continuous Datastream series which includes price jumps due to changes in the underlying futures series. The extreme returns in our series thus represent genuine financial risks to market participants.[14]

From all agricultural commodities traded in the United States, the following commodities are investigated: Corn, cotton, oats, soybeans, wheat, live cattle, orange juice (futures market) or oranges (spot market), and sugar.[15]

5 Testing for heavy-tailedness

Table 1 reports the descriptive statistics of the monthly returns in the spot market and the weekly returns on futures contracts. A quick overview of the data confirms the non-normality of the returns. Six out of eight series with monthly returns in the spot market contain at least one observation with a distance of more than 6 standard deviations from the mean, an observation that would occur roughly once every 42 million years under the assumption of a normal distribution.[16] For the weekly returns, six out of eight series contain at least one observation with a distance of more than five standard deviations from the mean. The probability of such a return occurring under the assumption of the normal distribution is about once every 33,500 years. Figure 2 reports QQ-plots of three arbitrarily chosen weekly return series (corn, cotton and live cattle) as an illustration. The non-normality is strongly confirmed by QQ-plots of the return series against the normal distribution. Only the right tail of the return series for live cattle seems to be quite close to the normal distribution.

We proceed with a formal test of whether the upper tail distribution is heavy-tailed by estimating parameter γ in Eq. 19 for each of the returns series. The theory shows how heavy-tailedness can arise in the right tail of agricultural commodity returns as

---

[14]In addition, shifts in the roll-over date often occur in the Datastream continuous series. To give an extreme example: the second largest daily price fall during the last 30 years in the unadjusted Datastream series for cotton (NCTCS00) is caused by a delayed roll-over date. The return of -26.3% is caused by the difference between 113.6, which is the price for delivery in July 1995 listed on the 4th of July, 1995, and 83.75, which is the price for delivery in October 1995 listed on the 5th of July, 1995. Such extreme observations may distort the assessment of the actual tail of the risk distribution.

[15]See Appendix C for details on the selection process of the commodities. The continuous futures returns are available from the corresponding author on request.

[16]If $\tilde{x}$ follows a standard normal distribution, then $\Pr(|\tilde{x}| > 6) \approx 1.97 \times 10^{-9}$.

Table 1 Descriptive Statistics

Commodity	Mean	St.dev.	Min.	Min. Date	Max.	Max. Date	Obs.
Panel (a): Monthly returns in the spot market (percentages), 1928-2014							
Corn	0.36	8.1	–45.7	1937-10	48.1	1936-07	944
Cotton	0.33	5.8	–29.2	2010-12	44.5	1986-09	929
Oats	0.49	8.9	–41.5	1935-08	68.4	1988-06	940
Soybeans	0.42	8.6	–36.6	1931-10	58.1	1974-07	922
Wheat	0.59	8.2	–35.1	2014-09	39.7	1973-07	912
Live cattle	0.30	5.1	–16.6	1931-12	33.6	1935-01	953
Oranges	2.37	25.1	–85.7	2000-08	393.1	2000-10	941
Sugar	0.67	11.1	–39.3	1979-08	75.4	1983-05	914
Panel (b): Weekly returns in the futures market (percentages), 1979-2012							
Corn	–0.02	3.4	–16.4	2008-10-03	20.8	2008-12-12	1,773
Cotton	0.01	3.2	–13.9	2008-10-10	18.4	2010-12-03	1,773
Oats	–0.04	4.1	–16.5	1988-07-22	31.4	1988-06-10	1,773
Soybeans	0.04	3.2	–14.8	2008-10-03	15.4	1999-07-23	1,773
Wheat	–0.06	3.5	–15.3	2008-12-05	16.5	2012-05-18	1,773
Live cattle	0.05	2.1	–14.4	2004-01-02	8.8	1982-01-08	1,773
Orange juice	0.00	4.1	–16.6	2012-05-18	37.7	1981-01-16	1,773
Sugar	0.09	4.9	–18.5	1988-07-29	38.8	1985-07-26	1,773

Note: The numbers in the table are expressed in percentage price changes. The first two columns report the mean and the standard deviation of the monthly and weekly returns series. The other columns report the minima and maxima of the returns and the dates of these observations

a consequence of price spikes from adverse productivity shocks that are not heavy-tailed themselves. Values of $\hat{\gamma}$ larger than 0 suggest that the tail of a distribution is heavy-tailed. Table 2, panels (a) and (b) report the estimates for the upper tail of the monthly returns in the spot market and the weekly returns in the futures market, respectively. Heavy-tailedness is rejected at a 5 percent significance level for t-statistics smaller than -1.6. Rejection of the thin-tailed distribution at a five percent significance level requires a t-statistic larger than 1.6.

Table 2, panel (a) shows that all estimates for γ, except one, are positive for the upper tail distribution of monthly returns in the spot market. This suggests that the right tails of these distributions are heavy-tailed. The exception to the rule is live cattle, which did also seem relatively close to the thin-tailed alternative in the QQ-plot. It is possible to reject the thin-tailed alternative for 5 out of 8 series based on the estimators of Beirlant et al. (1996) and Fraga Alves et al. (2009), with the strongest test results for corn, cotton, oats and sugar. The Dekkers et al. (1989) estimator, which exhibits somewhat larger standard errors, also rejects the thin-tailed alternative for 4 out of 8 series. None of the estimates suggest the rejection of heavy-tailedness for
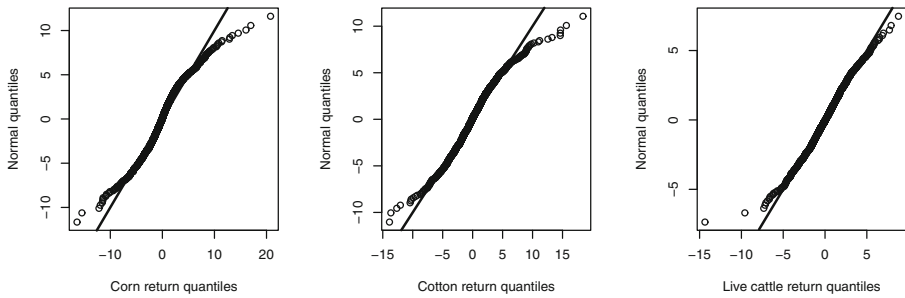
Fig. 2 QQ-plots of agricultural commodity returns. Note: QQ-plots of the weekly returns of futures contracts for corn, cotton and live cattle against a normal distribution with the same mean and standard deviation

the upper tail of monthly returns in the spot market. The pattern of heavy-tailedness is very similar when turning to the right tail distribution of weekly futures returns reported in Table 2, panel (b), although the results for weekly futures returns are somewhat less strong than for the monthly returns in the spot market. The thin-tailed alternative can be rejected at a 5 percent significance level for 2 out of 8 futures returns series based on the Dekkers et al. (1989) estimator, and for 3 out of 8 futures returns series based on the Beirlant et al. (1996) and Fraga Alves et al. (2009) estimators. The point estimates for live cattle and wheat are slightly negative. None of the estimates suggest the rejection of heavy-tailedness for the upper tail of weekly future returns.

6 Empirical tail estimates

Table 3 documents the estimated tail parameters under the assumption of heavy-tailed distributions. Unreported Hill plots show that the Hill estimates are relatively stable when a total of $k = 50$ tail observations are used for the monthly returns series and a total of $k = 75$ tail observations for the weekly returns series. This corresponds to 5.4% of all monthly observations and 4.2% of all weekly observations.

Table 3, panel (a) documents the estimated tail parameters using the Hill estimator for the monthly returns in the spot market. The value of most shape parameters is estimated to be in the range from 2 to 3. The most risky commodities with respect to the shape parameter describing price spikes are oranges, with an estimated tail index of around 1.95, and sugar, with an estimated tail index of around 2.36. The table also reports the Hill estimator for the left tail of the distribution. These estimates can be used to approximate the distribution using a Taylor approximation in the form of a power law function even though it is known that the returns in the left tail are bounded.

Table 3, panel (b) documents the estimated tail parameters for the weekly futures returns. Theoretically, the estimates of the shape parameters should be robust for changes in the data frequency in case of independent and identically distributed

Table 2 Testing for Heavy-tailedness

	$\hat{\gamma}_{DEH}$	(s.e.)	(t-stat)	$\hat{\gamma}_{BVT}$	(s.e.)	(t-stat)	$\hat{\gamma}_{FGHN}$	(s.e.)	(t-stat)
Panel (a): Right tail of monthly returns in the spot market, 1928-2014									
Corn	0.231	(0.13)	(1.8)	0.242	(0.10)	(2.3)	0.228	(0.12)	(1.8)
Cotton	0.268	(0.15)	(1.8)	0.268	(0.12)	(2.2)	0.262	(0.13)	(2.1)
Oats	0.438	(0.11)	(4.0)	0.338	(0.11)	(3.1)	0.597	(0.20)	(3.0)
Soybeans	0.260	(0.20)	(1.3)	0.231	(0.15)	(1.6)	0.225	(0.15)	(1.5)
Wheat	0.165	(0.14)	(1.2)	0.088	(0.12)	(0.7)	0.194	(0.09)	(2.1)
Live cattle	-0.016	(0.20)	(-0.1)	0.087	(0.17)	(0.5)	0.032	(0.12)	(0.3)
Oranges	0.320	(0.23)	(1.4)	0.350	(0.16)	(2.1)	0.248	(0.18)	(1.4)
Sugar	0.248	(0.14)	(1.7)	0.216	(0.11)	(1.9)	0.263	(0.11)	(2.3)
Panel (b): Right tail of weekly returns in the futures market, 1979-2012									
Corn	0.187	(0.12)	(1.5)	0.154	(0.09)	(1.7)	0.186	(0.10)	(1.9)
Cotton	0.125	(0.12)	(1.0)	0.116	(0.10)	(1.1)	0.143	(0.09)	(1.6)
Oats	0.280	(0.13)	(2.1)	0.278	(0.12)	(2.4)	0.269	(0.15)	(1.8)
Soybeans	0.032	(0.12)	(0.3)	0.047	(0.10)	(0.5)	0.077	(0.07)	(1.1)
Wheat	-0.015	(0.13)	(-0.1)	-0.019	(0.10)	(-0.2)	0.055	(0.07)	(0.7)
Livecattle	-0.108	(0.15)	(-0.7)	-0.064	(0.10)	(-0.7)	-0.018	(0.07)	(-0.2)
Orange juice	0.307	(0.11)	(2.8)	0.297	(0.12)	(2.6)	0.320	(0.14)	(2.2)
Sugar	0.149	(0.21)	(0.7)	0.147	(0.15)	(1.0)	0.122	(0.14)	(0.9)

Note: The table reports the point estimates and standard errors for parameter γ in the generalized extreme value distribution in Eq. 19. The columns correspond to the generalized Hill estimator of Beirlant et al. (1996) ($\hat{\gamma}_{BVT}$), the moment estimator of Dekkers et al. (1989) ($\hat{\gamma}_{DEH}$) and the mixed moment estimator of Fraga Alves et al. (2009) ($\hat{\gamma}_{FGHN}$). For the monthly returns in the spot market, each tail consists of 50 observations, or approximately 5.4% of the observations. For the weekly returns in the futures markets, each tail consists of 75 observations, or approximately 4.2% of the observations. Heavy-tailedness is not rejected for any of the series (rejection of heavy-tailedness at the 5% significance level requires a t-statistic smaller than -1.6)

returns.[17] Nevertheless, time series characteristics such as volatility clustering could result in differences in tail behavior at different data frequencies. Moreover, there are differences between the futures market and the spot market for commodities, where the latter is, in general, characterized by lower liquidity and higher trading costs in terms of organizing the logistics of storage and transportation. The tail parameter estimates for the weekly returns in the futures market are overall somewhat higher

---

[17]Mandelbrot (1963b) shows that power law distributions are invariant with respect to the shape parameter under several basic transformations. The shape parameter is invariant with regard to summation of random variables, mixing random variables with different scale parameters and selection of maxima. It follows that the power law distribution is independent of data frequency choices, distribution mixture assumptions and missing data. As a consequence, sample-specific data problems are unlikely to affect the observed shape parameter. But at lower frequencies, estimates are less efficient.

Table 3 Tail parameter estimates

Commodity	Left tail			Right tail		
	Shape	(s.e.)	Scale	Shape	(s.e.)	Scale
Panel (a): Monthly returns in the spot market, 1928-2014						
Corn	3.10	(0.45)	130.42	2.88	(0.41)	85.87
Cotton	3.52	(0.52)	96.84	2.57	(0.33)	17.12
Oats	2.54	(0.41)	28.17	2.94	(0.46)	130.58
Soybeans	2.72	(0.46)	53.78	2.59	(0.42)	41.03
Wheat	2.84	(0.48)	55.04	3.14	(0.39)	227.09
Live cattle	2.94	(0.42)	18.28	2.93	(0.42)	24.56
Oranges	2.46	(0.30)	173.72	1.95	(0.32)	56.90
Sugar	3.49	(0.41)	554.36	2.36	(0.41)	56.04
Panel (b): Weekly returns in the futures market, 1979-2012						
Corn	3.21	(0.36)	10.55	3.79	(0.39)	49.73
Cotton	3.50	(0.38)	13.47	3.06	(0.35)	8.29
Oats	3.81	(0.40)	72.21	3.40	(0.52)	30.67
Soybeans	3.04	(0.33)	6.70	3.30	(0.36)	11.46
Wheat	3.40	(0.46)	12.92	3.20	(0.27)	14.63
Live cattle	4.04	(0.54)	8.63	3.93	(0.40)	7.49
Orange juice	3.10	(0.30)	13.38	2.85	(0.35)	11.77
Sugar	4.20	(0.51)	348.86	3.63	(0.45)	108.49

Note: The columns report the shape and scale parameters estimated using the Hill estimator in Eqs. 21 and Eq. 22. The standard errors are obtained from the block bootstrap procedure. The third columns report the estimated scale parameter from Eq. 22. For the monthly returns in the spot market, each tail consists of $k = 50$ observations, or approximately 5.4% of the observations. For the weekly returns in the futures markets, each tail consists of $k = 75$ observations, or approximately 4.2% of the observations

than those of the monthly returns in the spot market with estimates that are more often in the range from 3 to 4. Hence, our estimates suggest the presence of a higher risk of extreme price spikes in the spot market than is documented for the futures market. Finally, live cattle is the commodity with the highest estimates for the shape parameters, implying less heavy tails, which is consistent with our previous results in Table 2.

While the shape parameter α is relatively similar across commodities within each market, there is a larger variation in the estimates of the scale parameter C. The variation in the scale at least partly explains the variation in the VaR estimates within each market; see Tables 4 and 5 discussed in the next section. Suppose that the shape parameter α is identical across commodities within a particular market. Then, assuming independence of returns across commodities, an optimal portfolio could be constructed on the basis of the mean and the scale (Hyung and De Vries 2005). Here,

Table 4 Risk estimates for monthly returns in the spot market, 1928-2014

Commodity	Probability level: 1.00%		Probability level: 0.25%	
	VaR	(95% conf. int.)	VaR	(95% conf. int.)
Panel (a): Left tail				
Corn	19.68	(16.87, 23.31)	26.79	(21.39, 36.57)
Cotton	13.04	(11.03, 14.79)	17.97	(14.00, 22.91)
Oats	20.41	(17.52, 24.44)	29.12	(22.92, 39.88)
Soybeans	20.47	(17.73, 23.25)	26.57	(23.17, 31.27)
Wheat	18.30	(16.02, 21.93)	24.21	(20.28, 32.10)
Live cattle	11.50	(10.08, 12.97)	14.66	(12.78, 16.40)
Oranges	45.20	(40.20, 51.53)	61.06	(53.88, 71.08)
Sugar	21.15	(18.78, 23.56)	27.49	(23.74, 32.58)
Panel (b): Right tail				
Corn	23.23	(19.40, 27.75)	37.59	(28.47, 50.01)
Cotton	18.14	(14.48, 22.70)	31.11	(22.70, 44.16)
Oats	25.03	(20.69, 31.86)	40.09	(30.34, 61.07)
Soybeans	24.77	(19.93, 31.12)	42.28	(31.01, 60.20)
Wheat	24.33	(20.68, 29.37)	37.82	(30.54, 51.91)
Live cattle	14.41	(12.16, 16.27)	23.15	(17.50, 27.49)
Oranges	84.61	(65.27, 108.60)	172.37	(111.73, 259.72)
Sugar	38.76	(32.33, 47.48)	69.74	(52.32, 105.24)

Note: The Value-at-Risk (VaR) estimates are expressed in percentage price changes. VaR estimates for the left tail in panel (a) are calculated from Eq. 24 using the general estimator of Fraga Alves et al. (2009) for $\hat{\gamma}_G$. VaR estimates for the right tail in panel (b) are calculated from Eq. 23 using the Hill estimator. Estimates use $k = 50$ observations, or approximately 5.4% of the observations. The confidence intervals (conf. int.) at a 95% level in parentheses are based on a block bootstraps procedure

the scale substitutes for the variance as in more traditional mean-variance portfolio optimization.[18]

7 Risk estimates

This section discusses the empirical results of the risk estimates. Table 4 reports risk estimates for monthly returns in the spot market. The risk estimates for the right tail are based on the estimator in Eq. 23, which assumes heavy tails in line with the theory and results shown above. The risk estimates for the left tails are based on the generalized estimator in Eq. 24. As an illustration take the VaR levels for increases

---

[18]Without independence of returns across commodities, the tail dependence among commodity returns would also play a role; see, e.g., Van Oordt and Zhou (2016).

Table 5 Risk estimates for weekly returns in the futures market, 1979-2012

Commodity	Probability level: 0.25%		Probability level: 0.05%	
	VaR	(95% conf. int.)	VaR	(95% conf. int.)
Panel (a): Left tail				
Corn	11.19	(9.52, 13.11)	15.09	(12.13, 19.81)
Cotton	9.81	(8.44, 11.74)	12.84	(10.18, 17.17)
Oats	12.31	(11.26, 13.59)	15.31	(13.50, 17.95)
Soybeans	10.81	(9.35, 12.30)	14.56	(11.77, 18.12)
Wheat	10.55	(8.99, 12.40)	14.21	(11.24, 18.94)
Live cattle	7.09	(5.74, 9.56)	9.88	(6.95, 18.35)
Orange juice	12.40	(11.17, 13.99)	15.98	(13.84, 19.66)
Sugar	14.33	(13.11, 16.17)	17.63	(15.29, 21.52)
Panel (b): Right tail				
Corn	13.62	(12.08, 16.39)	20.82	(17.49, 27.85)
Cotton	14.14	(11.21, 18.12)	23.92	(17.18, 34.28)
Oats	15.93	(12.64, 20.99)	25.57	(18.04, 38.59)
Soybeans	12.88	(10.59, 15.39)	20.97	(15.79, 27.43)
Wheat	14.98	(12.63, 17.27)	24.75	(19.64, 30.52)
Live cattle	7.65	(6.81, 8.74)	11.52	(9.73, 14.36)
Orange juice	19.50	(15.82, 24.49)	34.32	(25.03, 49.07)
Sugar	18.90	(16.39, 22.92)	29.43	(23.72, 40.19)

Note: The Value-at-Risk (VaR) estimates are expressed in percentage price changes. VaR estimates for the left tail in panel (a) are calculated from Eq. 24 using the general estimator of Fraga Alves et al. (2009) for $\hat{\gamma}_G$. VaR estimates for the right tail in panel (b) are calculated from Eq. 23 using the Hill estimator. Estimates use $k = 75$ observations, or approximately 4.2% of the observations. The confidence intervals (conf. int.) at a 95% level in parentheses are based on a block bootstraps procedure

in the price of oranges of 84.6% and 172.6% at a 1.00% and 0.25% probability level, respectively.[19] The first estimate means that once every 100 months (around 8 years), one may expect the spot price of oranges to increase by more than 80% within a single month. The price is likely to more than double once every 400 months, or about once every 33 years. The risk estimates also reflect the pattern that the largest price spikes tend to exceed the largest price declines in the data (Table 1), with larger risk estimates for the right tail than the left tail.

The highly volatile behaviour of the agricultural commodity prices is again evidenced in the table with the risk estimates for weekly futures returns (Table 5). In the same vein, the table reports the 0.25% and 0.05% VaR for the weekly futures

---

[19] We refer to Boudoukh et al. (2007) for a study of the underlying causes for highly volatile orange juice prices.

returns.[20] The 0.25% VaR is expected to be exceeded about once every eight years, and the 0.05% VaR about once every 40 years, or about once during the career of an investment professional.

The relevant risk management information for an investor with a long position is contained in the left tail of the distribution, shown in Table 5, panel (a). We find that sugar has the highest risk estimates for price declines in the futures market of all commodities studied. Once every forty years the price of sugar is expected to fall by more than 17.6% within a week. The safest commodity, in terms of price development, appears to be live cattle. Once every forty years the price of live cattle is expected to fall by more than 9.9% within a week. Orange juice tends to have the highest upside price risk in the futures market, with a price spike that exceeds 34.3% within a week once every forty year.

In the introduction, a question was posed regarding the likelihood of a price change occurring in excess of an market participant's tolerance level for price increases of 25% (e.g., a food processor or importing country). We apply Eq. 1 using the estimated parameters $\hat{C}$ and $\hat{\alpha}$ for the right tail of the weekly return distribution for wheat in Table 3. This gives the probability of exceeding the threshold as $\Pr(\tilde{x} > u) \approx 14.63 \times 25^{-3.20} \approx 4.84 \times 10^{-4}$. The inverse of this number yields the number of weeks in which at least one price spike of more than 25% is expected to occur. The outcome is around 2,065 weeks. Hence, with 52 weeks per year, we expect to see such a large price jump for wheat once every 40 years.

8 Back-testing

To examine the accuracy of the $VaR(p)$ estimates, we employ an out-of-sample back-testing procedure. In this method, the $VaR(p)$ estimates based on historical price changes are compared to the realized price changes. Thus, first $VaR_{t-n+m+1}(p)$ is estimated using a horizon of m past returns: $\{R_{t-n}, ..., R_{t-n+m}\}$. If the realized return $R_{t-n+m+1}$ exceeds the estimated $VaR_{t-n+m+1}(p)$, then a "VaR-exceedance" is registered. The above procedure is repeated with a rolling window at time $t - n + m + 2$ et cetera. For a series containing n returns, the procedure yields a total of $n - m$ observations of whether a VaR-exceedance occurred. According to the VaR definition, if the methodology provides accurate predictions of $VaR(p)$, then the proportion of VaR-exceedances should have a value close to p. Thus, $\sum_{j=m+1}^{n} \mathbf{1}(R_{t-n+j} > VaR_{t-n+j}(p))/(n - m) = p$ should match approximately for accurate VaR estimates, where $\mathbf{1}(\cdot)$ denotes the indicator function. An unconditional coverage test is used to test whether the number of observed VaR-exceedances deviates statistically significantly from the expected number (Kupiec 1995).

The procedure is implemented as follows. The back-testing procedure is executed under the assumption of the normal distribution, the generalized EVT distribution

---

[20]Note that the magnitude of the VaR estimates in Tables 5 and 4 cannot be compared directly by the "square-root-of-time" as would be possible under normality. With heavy tails, one would have to scale up the weekly VaRs by approximately $4^{1/\alpha}$ (Dacorogna et al. 2001, Proposition 2). For $\alpha \approx 3$, the proper scaling factor is approximately 1.6, not $\sqrt{4} = 2$.

and the power law. For the monthly returns in the spot market, the length of the rolling window for the estimation is set at $m = 360$ observations (30 years) and each tail is assumed to contain the most extreme 8.3% (1/12th) of the observations in the estimation horizon. For the weekly returns in the futures market, the length of the rolling window is set at $m = 520$ weekly observations (10 years) and each tail is assumed to contain the most extreme 5.0% of the observations in the estimation horizon. A visual representation of the procedure's results is given in Fig. 3. The spikes show weekly wheat returns. The bold lines show the 0.25% VaR estimates from the normal distribution (in gray), the generalized EVT distribution (in blue) and the power law (in red). The first period without the VaR bands reflects the window needed for estimation.

The results of the back-testing procedure are reported in Tables 6 and 7 for monthly and weekly data, respectively. The columns report the number of VaR-exceedances at three different VaR probabilities for each commodity and each distribution. The line labelled "average" reports the average number of observed VaR-exceedances in each column. The bottom line indicates the expected number of VaR-exceedances for accurate VaR estimates. A risk assessment methodology that underestimates the risk would lead to an average of observed VaR-exceedances that is higher than the expected number, while risk estimates that are too conservative would lead to a number of VaR-exceedances that is too low. Finally, the bottom line reports the number of unconditional coverage tests that are rejected at a 5 percent significance level for each column.

Back-testing risk estimates for the right tail of the monthly returns in the spot market for agricultural commodities exemplify the benefits of assuming a power law
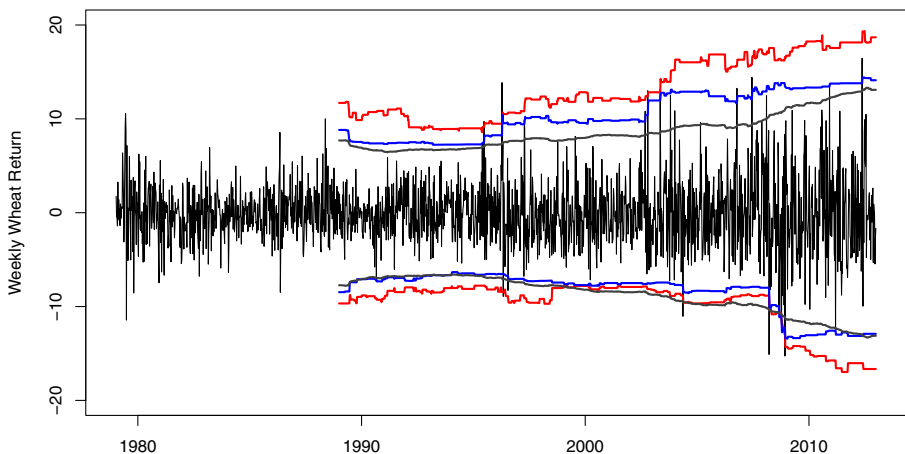

Fig. 3 Backtesting agricultural prices Value-at-Risk Note: The spikes show the weekly wheat returns. The lines report the VaR estimated using a rolling window of 520 weeks. The estimates are based on the normal distribution (in gray), generalized EVT distribution using the estimator of Fraga Alves et al. (2009) (in blue) and a power law tail distribution using the estimator of Hill (1975) (in red). The VaR estimate from the normal distribution and generalized EVT distribution are exceeded at a frequency higher than once every eight years

Table 6 Back-testing Monthly Returns in the Spot Market, 1928-2014

Distribution:	Normal			Generalized EVT			Heavy-tailed		
	1.00%	0.50%	0.25%	1.00%	0.50%	0.25%	1.00%	0.50%	0.25%
Panel (a): Number of VaR-exceedances in left tail									
Corn	14	10	6	19	11	6	8	2	0
Cotton	15	12	11	16	11	5	8	2	1
Oats	12	9	5	17	10	3	6	1	0
Soybeans	13	9	7	16	8	1	7	2	1
Wheat	17	12	7	17	9	6	10	5	0
Live cattle	6	4	3	9	4	3	4	1	0
Orange juice	11	9	7	19	13	9	7	0	0
Sugar	6	4	3	9	4	1	3	1	1
Average	11.8	8.6	6.1	15.2	8.8	4.2	6.6	1.8	0.4
(Expected)	(6.7)	(3.3)	(1.7)	(6.7)	(3.3)	(1.7)	(6.7)	(3.3)	(1.7)
Rejected tests	4	6	5	6	6	3	0	0	0
Panel (b): Number of VaR-exceedances in right tail									
Corn	12	9	8	11	6	4	7	5	4
Cotton	24	19	16	16	9	5	8	2	2
Oats	18	14	10	15	7	4	10	3	2
Soybeans	13	7	5	8	5	4	5	4	4
Wheat	22	17	14	22	11	7	13	7	2
Live cattle	15	9	7	12	8	5	6	2	0
Orange juice	19	17	14	14	10	4	11	3	2
Sugar	25	19	17	16	7	4	5	0	0
Average	18.5	13.9	11.4	14.2	7.9	4.6	8.1	3.2	2.0
(Expected)	(6.7)	(3.3)	(1.7)	(6.7)	(3.3)	(1.7)	(6.7)	(3.3)	(1.7)
Rejected tests	7	7	7	5	4	1	1	0	0

Note: The table reports the number of VaR-exceedances in the back-testing procedure. We estimate the VaR with a rolling-window of 360 months based on the normal distribution, the generalized EVT distribution using the estimator of Fraga Alves et al. (2009) and the power law tail distribution. For each commodity, we count how often the return in any particular month exceeds the VaR estimate based on the past 360 months (a "VaR-exceedance"). The lines labelled "average" and "expected" report the average number and expected number of VaR-exceedances for accurate VaR estimates for each column. The bottom line reports the number of unconditional coverage tests that were rejected at a 5 percent significance level for each column

tail. Table 6, panel (b) shows that assuming a normal distribution leads to far too many VaR-exceedances in the right tail. The predictions from the normal distribution increasingly underestimate the risk when one moves further into the tail. The third column in the table shows that assuming a normal distribution leads to a number of 0.25% VaR-exceedances that is more than 5 times too high. Moving to the generalized EVT distribution partly alleviates this problem, since the actual number

Table 7 Back-testing Weekly Returns in the Futures Market, 1979-2012

Distribution:	Normal			Generalized EVT			Heavy-tailed		
	1.00%	0.50%	0.25%	1.00%	0.50%	0.25%	1.00%	0.50%	0.25%
Panel (a): Number of VaR-exceedances in left tail									
Corn	27	16	11	29	20	11	25	12	5
Cotton	23	13	10	31	18	9	23	9	5
Oats	25	17	11	26	15	7	21	5	2
Soybeans	26	17	13	23	13	8	16	8	4
Wheat	17	10	8	30	14	9	24	10	5
Live cattle	17	12	8	14	8	4	13	6	3
Orange juice	28	19	15	28	16	8	24	12	3
Sugar	21	14	10	18	11	6	13	5	1
Average	23.0	14.8	10.8	24.9	14.4	7.8	19.9	8.4	3.5
(Expected)	(12.5)	(6.3)	(3.1)	(12.5)	(6.3)	(3.1)	(12.5)	(6.3)	(3.1)
Rejected tests	6	6	8	6	6	5	5	0	0
Panel (b): Number of VaR-exceedances in right tail									
Corn	37	27	19	23	12	4	15	6	2
Cotton	30	22	17	24	15	7	20	8	3
Oats	21	13	7	12	7	4	10	5	5
Soybeans	23	13	10	19	10	3	17	3	1
Wheat	35	23	18	24	15	10	18	11	4
Live cattle	18	11	8	20	9	4	17	6	3
Orange juice	32	24	20	22	13	7	18	9	2
Sugar	18	12	6	18	9	4	13	4	1
Average	26.8	18.1	13.1	20.2	11.2	5.4	16.0	6.5	2.6
(Expected)	(12.5)	(6.3)	(3.1)	(12.5)	(6.3)	(3.1)	(12.5)	(6.3)	(3.1)
Rejected tests	6	6	6	4	3	1	0	0	0

Note: The table reports the number of VaR-exceedances in the back-testing procedure. We estimate the VaR with a rolling-window of 520 weeks based on the normal distribution, the generalized EVT distribution using the estimator of Fraga Alves et al. (2009) and the power law tail distribution. For each commodity, we count how often the return in any particular week exceeds the VaR estimate based on the past 520 weeks (a "VaR-exceedance"). The lines labelled "average" and "expected" report the average number and expected number of VaR-exceedances for accurate VaR estimates for each column. The bottom line reports the number of unconditional coverage tests that were rejected at a 5 percent significance level for each column

of VaR-exceedances is closer to the expected number of VaR-exceedances, but still exceeds the expected number of VaR-exceedances by more than a factor two for the probability levels 0.50% and 0.25%. In contrast, assuming heavy-tailedness leads to a number of VaR-exceedances that is close to the expected number. For the probability levels of 0.50% and 0.25%, the average number of exceedances is, respectively, 3.2 and 2.0, which is close to their expected number of 3.3 and 1.7, respectively. These

estimates concern the most extreme events, or about one observation in 16.7 and 33.3 year, respectively. The generally better performance of assuming a heavy-tailed distribution is also reflected in the lower number of rejected unconditional coverage tests in the bottom line of the table, in particular for the smaller probability levels.

As discussed, one may expect for the left tail that the heavy-tailed approximation leads to risk estimates that are somewhat too conservative. This is indeed confirmed in Table 6, panel (a), where the heavy-tailed approximation leads to fewer VaR-exceedances than expected. For the probability levels of 0.50% and 0.25%, the average number of exceedances is, respectively, 1.8 and 0.4, which is less than the expected number of 3.3 and 1.7, respectively. However, depending on the context, it may still be better to use the power law approximation than the generalized EVT distribution, which turns out to lead to more VaR-exceedances than expected, and to a rejection of the unconditional coverage test for some of the series.

Assuming heavy-tailedness also works well when modeling the risk for weekly returns in the spot market for observations both far in the right and far in the left tail (Table 7). The average number of observed exceedances in the left and right tail at the 0.25% probability is, respectively, 3.5 and 2.6, which is close to the expected number of 3.1. These estimates concern events that are expected to occur only once every 8 years. The accuracy of risk estimates for weekly returns based on the power law is also substantially better when compared to the normal and the generalized EVT distribution, which both tend to underestimate the risk and lead to too many VaR-exceedances. The results for the left tail provide support for using a Taylor approximation in the form of a power law function to estimate extreme price risks at medium to higher loss levels in the context of weekly futures returns, even in situations where the theoretical return distribution is bounded.

9 Conclusion

A proper understanding of extreme commodity returns is instrumental to any market participant in agricultural commodity markets. Especially knowledge regarding the largest price swings is important for risk management purposes. We construct a two-sector general equilibrium model which describes how productivity shocks affect agricultural commodity prices. In our model, extreme price spikes arise as a result of productivity shocks in the agricultural sector, which are translated into upward price changes in a nonlinear manner, resulting in a heavy-tailed equilibrium return distribution in the commodity market. Our study shows how this mechanism may result in heavy-tailed price distributions, even if the productivity shocks are not heavy-tailed themselves.

We provide further empirical evidence on the heavy-tailedness of agricultural commodity price returns in both the futures market and the spot market. Extreme Value Theory is used to estimate the parameters of the power law in the tail of their distribution. These estimates are used to measure an agricultural commodity investor's exposure to extreme price risks. We calculate Value-at-Risk measures to provide estimates of the likelihood and size of the largest losses an investor may encounter. Back-testing shows that this methodology provides relatively accurate estimates of

risk measures for agricultural commodities when compared to estimates based on the normal distribution or the generalized EVT distribution, in particular for the right tail.

Our results have two implications. First, we document that market participants in this asset class are likely to encounter large future price swings. Our model supports the prevalence of fat-tailed commodity price distributions and thus a highly volatile return behaviour. Secondly, the risk measurement framework in this paper informs market participants on a well-suited method to quantitatively assess the largest price risks, with our empirical results providing essential knowledge on when the methodology performs particularly well.

An interesting observation is that the price spikes of live cattle tend to be less heavy-tailed (if at all) than the price spikes of the crop commodities. One potential explanation that is consistent with our theoretical framework is that the production of crop commodities in our dataset are strongly affected by adverse productivity shocks arising from short-term weather conditions, while the supply of live cattle is less impacted by such adverse productivity shocks. Arguably, a single commodity is insufficient to validate our framework, but exploring the factors that can explain the differences in heavy-tailedness among a broader set of commodities is an interesting direction for future research.

Appendix A: Derivation of equilibrium prices

A.1 Demand

The first order conditions for optimality entail, for each $j = 1, (...), n$,

$$\theta \left(\frac{Z}{\left[\sum_{i=1}^{n} Q_i^\rho\right]^{1/\rho}} \right)^{1-\theta} n^{-\theta/\rho} \left[\sum_{i=1}^{n} Q_i^\rho\right]^{\frac{1}{\rho}-1} Q_j^{\rho-1} - \lambda \frac{1}{n} p_j = 0,$$

$$(1-\theta) Z^{-\theta} n^{-\theta/\rho} \left[\sum_{i=1}^{n} Q_i^\rho\right]^{\theta/\rho} - \lambda q = 0,$$

$$-L^\delta + \lambda w = 0,$$

and

$$wL + \Pi(Q) = qZ + \frac{1}{n} \sum_{i=1}^{n} p_i Q_i.$$

The first order conditions imply the familiar price and wage ratios. For each $i, j = 1, (...), n$, we have

$$\frac{p_i}{p_j} = \frac{Q_i^{\rho-1}}{Q_j^{\rho-1}},$$

for each $j = 1, (...), n$, we have

$$\frac{p_j}{q} = \frac{\theta}{1-\theta} \frac{Z}{Q_j} \frac{p_j^{\rho/(\rho-1)}}{\frac{1}{n} \sum_{i=1}^{n} p_i^{\rho/(\rho-1)}},$$

and

$$\frac{w}{q} = \left(\frac{P}{q}\right)^{\theta} \frac{L^{\delta}}{(1-\theta)^{1-\theta} \theta^{\theta}},$$

where the price index for differentiated goods P is defined as in equation Eq. 8.

Then the labor supply can be written as

$$L = \left((1-\theta)^{1-\theta} \theta^{\theta} \frac{w}{q^{1-\theta} P^{\theta}}\right)^{1/\delta}. \tag{25}$$

The competitive good demand is

$$Z = (1-\theta) \frac{wL + \Pi(Q)}{q}. \tag{26}$$

The differentiated goods demanded can be expressed for each $i = 1, (...), n$ as

$$Q_i = \theta \frac{wL + \Pi(Q)}{p_i} \left(\frac{p_i}{P}\right)^{\rho/(\rho-1)}. \tag{27}$$

A.2 Supply

From the perfectly competitive agricultural market we have that equilibrium profits in the agricultural sector equal zero, that is,

$$\Pi_\theta(Z) = qZ - wN = \left(q - \frac{w}{B}\right) Z = 0,$$

so that

$$q = w/B. \tag{28}$$

The differentiated goods profit function for each $i = 1, (...), n$ reads

$$\Pi_i(Q_i) = p_i Q_i - wN_i = \left(p_i - \frac{w}{A}\right) Q_i$$

$$= \left(p_i - \frac{w}{A}\right) \theta \frac{wL + \Pi(Q)}{p_i} \left(\frac{p_i}{P}\right)^{\rho/(\rho-1)}.$$

The producer exploits his pricing power, but ignores his pricing effect on the price index P of the differentiated goods and the consumer income $wL + \Pi(Q)$.[21] Differentiation gives for each $i = 1, (...), n$

$$\frac{\partial \Pi_i(Q_i)}{\partial p_i} = \frac{1}{\rho - 1} Q_i \left(\rho - \frac{1}{A} \frac{w}{p_i} \right).$$

Exploiting the pricing power therefore implies setting prices for each $i = 1, (...), n$

$$p_i = \frac{w}{\rho A}. \tag{29}$$

Hence, $P = w/\rho A$ as all prices are identical. Total profits in the differentiated goods sector equal

$$\Pi(Q) = \frac{1}{n} \sum_{i=1}^{n} \Pi_i(Q_i) = \sum_{i=1}^{n} \left(1 - \frac{w/p_i}{A} \right) \theta \left[wL + \Pi(Q) \right] \left(\frac{p_i}{P} \right)^{\rho/(\rho-1)}$$
$$= (1 - \rho) \theta \left[wL + \Pi(Q) \right].$$

Solve for the total sectorial profits as

$$\Pi(Q) = \frac{(1 - \rho)\theta}{1 - (1 - \rho)\theta} wL. \tag{30}$$

A.3 Equilibrium

It follows in equilibrium, after substituting the price levels into the labor supply Eq. 25, that

$$L = \left(\theta^\theta (1 - \theta)^{1-\theta} A^\theta B^{1-\theta} \right)^{1/\delta} \rho^{\theta/\delta} = \varphi \rho^{\theta/\delta}, \tag{31}$$

say, and where

$$\varphi = \left(\theta^\theta (1 - \theta)^{1-\theta} A^\theta B^{1-\theta} \right)^{1/\delta}.$$

Furthermore, from Eqs. 26, 30 and 31, we obtain

$$Z = (1 - \theta) \frac{B}{1 - (1 - \rho)\theta} \varphi \rho^{\theta/\delta}. \tag{32}$$

Similarly, using Eqs. 27, 30 and 31, we obtain for each $j = 1, (...), n$

$$Q_j = \theta \frac{A}{1 - (1 - \rho)\theta} \rho \varphi \rho^{\theta/\delta}.$$

Hence,

$$\frac{1}{n} \sum_{j=1}^{n} Q_j = \theta \frac{A}{1 - (1 - \rho)\theta} \varphi \rho^{\theta/\delta+1}. \tag{33}$$

---

[21] One can easily incorporate this effect as well, if desired. For two reasons we do not follow this route. One may doubt that producers take this macro effect of their pricing behavior into account. Moreover, it adds little to the insights derived from specifying the differentiated goods sector. The macro literature usually ignores this effect as well.

With the above preparations, we now derive the implications for the equilibrium prices. From Eq. 28, combined with Eqs. 9 and 31, we obtain

$$q = \frac{w}{B} = \frac{M}{B}\frac{1}{L} = M\frac{1/\rho^{\theta/\delta}}{B\left(\theta^\theta(1-\theta)^{1-\theta}A^\theta B^{1-\theta}\right)^{1/\delta}}.$$

Similarly, using Eq. 29 combined with Eqs. 9 and 31 yields for each $i = 1, (...), n$

$$p_i = p = \frac{w}{\rho A} = \frac{M}{\rho A}\frac{1}{L} = M\frac{1/\rho^{\theta/\delta+1}}{A\left(\theta^\theta(1-\theta)^{1-\theta}A^\theta B^{1-\theta}\right)^{1/\delta}}.$$

Appendix B: Proof of Proposition 2

Given Eq. 12, we want to find the condition on the density for B such that probability distribution of the price $\tilde{q}$ follows a heavy-tailed distribution. We have that $\Pr(\tilde{q} > u) \sim \mathcal{L}(u)u^{-\alpha}$ as $u \to \infty$ if $\tilde{q}$ is regularly α-varying at infinity with $0 < \alpha < \infty$, i.e., if

$$\lim_{t\to\infty}\frac{1 - F_q(tu)}{1 - F_q(t)} = u^{-\alpha} \text{ with } \alpha \in \mathbb{R}^+, \qquad (34)$$

where F_q denotes the cumulative distribution function of $\tilde{q}$, see also De Haan (1970). We need to find the condition such that F_q it is regularly varying at infinity. Rewriting Eq. 34 with L'Hôpital's Rule gives the condition

$$\lim_{t\to\infty}\frac{uf_q(tu)}{f_q(t)} = u^{-\alpha} \text{ with } \alpha \in \mathbb{R}^+, \qquad (35)$$

where f_q denotes the density of $\tilde{q}$. Given Eq. 12, we have that the equilibrium price $q(B)$ is a strictly decreasing function of B for $\theta \in (0, 1)$. Therefore, by a transformation of variable we have that

$$f_q(\tilde{q}) = \left|\frac{dB(\tilde{q})}{d\tilde{q}}\right| f_B(B(\tilde{q})), \qquad (36)$$

where $B(\tilde{q})$ denotes the inverse of $\tilde{q}(B)$. With the inverse of Eq. 12 and the derivative of the inverse of Eq. 12 this gives

$$f_q(\tilde{q}) = \frac{1}{\eta}\Theta^{1/\eta}\tilde{q}^{-(1/\eta+1)}f_B(\Theta^{1/\eta}\tilde{q}^{-1/\eta}), \qquad (37)$$

where

$$\eta = \frac{1 + \delta - \theta}{\delta}.$$

Substituting Eq. 37 into Eq. 35 gives

$$\lim_{t\to\infty}\frac{u\frac{1}{\eta}\Theta^{1/\eta}(tu)^{-(1/\eta+1)}f_B(\Theta^{1/\eta}(tu)^{-1/\eta})}{\frac{1}{\eta}\Theta^{1/\eta}(t)^{-(1/\eta+1)}f_B(\Theta^{1/\eta}(t)^{-1/\eta})} = u^{-\alpha} \text{ with } \alpha \in \mathbb{R}^+.$$

Simplifying gives

$$\lim_{t \to \infty} \frac{u^{-1/\eta} f_B(\Theta^{1/\eta} t^{-1/\eta} u^{-1/\eta})}{f_B(\Theta^{1/\eta} t^{-1/\eta})} = u^{-\alpha} \text{ with } \alpha \in \mathbb{R}^+;$$

$$\lim_{s \downarrow 0} \frac{w f_B(sw)}{f_B(s)} = w^{\eta \alpha} \text{ with } \alpha \in \mathbb{R}^+. \tag{38}$$

where we use the notation $w = u^{-1/\eta}$ and $s = \Theta^{1/\eta} t^{-1/\eta}$. Hence, given Eq. 12, if Eq. 38 holds true for the density function of B, we have that $\Pr(\tilde{q} > u) \sim \mathcal{L}(u) u^{-\alpha}$ as $u \to \infty$. Proposition 2 is then obtained by writing $\xi = \eta \alpha$ in the condition in Eq. 38 and using $\eta > 0$.

Appendix C: Commodity selection

In this appendix we explain how the commodities for this research are selected. The employed futures series need to satisfy two conditions: Availability and relevance for the US agricultural industry. Initially, our sample contains all traded commodity futures within the US. A list of 19 commodities remains after removing non-agricultural and identical commodities. Four commodities from this list are removed because of data availability (butter, milk, dry whey and rice are only available from 1996 onwards or even later). Next, six of the remaining fifteen series are removed because of low relevance: Because soybeans is included, soy meal and soybean oil are removed. Because live cattle is included, cattle feeder is removed. Cocoa and coffee are removed because of relatively low relevance for the US agricultural sector. Finally, lean hogs and frozen pork bellies were dropped because of lack of reliable long time series with spot prices. With both the Thomson Reuters Datastream codes and Global Financial Data tickers provided in parenthesis, this leaves us with corn (CC.; __C_US2D), cotton (NCT; COT_AFRD), oats (CO.; OATS_RAD), soybeans (CS.; __SYB_TD), wheat (CW.; __W_USSD), live cattle (CLC; _ICXD), orange juice (NJO; CMORANGM), sugar (NSB; _SU1599D).

References

Aiyagari, S.R., Eckstein, Z., Eichenbaum, M.: Inventories and price fluctuations under perfect competition and monopoly. In: Kollintzas, T. (ed.) The Rational Expectations Equilibrium Inventory Model, pp. 34–68. Springer (1989)

Ardeni, P.G., Freebairn, J.: The macroeconomics of agriculture. In: Rausser, G., Gardner, B. (eds.) Handbook of Agricultural Economics, vol. 2a, pp. 1455–1485. Elsevier (2002)

Balagtas, J.V., Holt, M.T.: The commodity terms of trade, unit roots, and nonlinear alternatives: A smooth transition approach. Am. J. Agric. Econ. **91**(1), 87–105 (2009)

Balkema, A.A., De Haan, L.: Residual life time at great age. Ann. Probab. **91**(1), 792–804 (1974)

Beirlant, J., Vynckier, P., Teugels, J.L.: Excess functions and estimation of the extreme-value index. Bernoulli **2**(4), 293–318 (1996)

Bidarkota, P., Crucini, M.J.: Commodity prices and the terms of trade. Rev. Int. Econ. **8**(4), 647–666 (2000)

Boudoukh, J., Richardson, M., YuQing, S., Whitelaw, R.F.: Do asset prices reflect fundamentals? Freshly squeezed evidence from the OJ market. J. Financ. Econ. **83**(2), 397–412 (2007)

Dacorogna, M.M., Müller, U.A., Pictet, O.V., De Vries, C.G.: Extremal forex returns in extremely large data sets. Extremes **4**(2), 105–127 (2001)

De Haan, L.: On Regular Variation and its Application to the Weak Convergence of Sample Extremes. PhD thesis, Amsterdam (1970)

De Roon, F.A., Nijman, T.E., Veld, C.: Hedging pressure effects in futures markets. J. Financ. **55**(3), 1437–1456 (2000)

Deaton, A.: Commodity prices and growth in Africa. J. Econ. Perspect. **13**(3), 23–40 (1999)

Dekkers, A.L.M., Einmahl, J.H.J., De Haan, L.: A moment estimator for the index of an extreme-value distribution. Ann. Statist. **17**(4), 1833–1855 (1989)

Dixit, A.K., Stiglitz, J.E.: Monopolistic competition and optimum product diversity. Amer. Econ. Rev. **67**(3), 297–308 (1977)

Drees, H.: Some aspects of extreme value statistics under serial dependence. Extremes **11**(1), 35–53 (2008)

Drees, H., De Haan, L., Resnick, S.: How to make a Hill plot. Ann. Statist. **28**(1), 254–274 (2000)

ECB: Assessing banking system risk with extreme value analysis. Financial Stability Review (June), 155–162 (2006)

Erb, C.B., Harvey, C.R.: The strategic and tactical value of commodity futures. Financ. Anal. J. **62**(2), 69–97 (2006)

Evans, P.: The effects of general price controls in the United States during world war ii. J. Polit. Econ. **90**(5), 944–966 (1982)

Fisher, I.: The Purchasing Power of Money: Its Determination and Relation to Credit, Interest and Crises. Macmillan, New York (1911)

Fraga Alves, M.I., Gomes, M.I., de Haan, L., Neves, C.: Mixed moment estimator and location invariant alternatives. Extremes **12**(2), 149–185 (2009)

Frisch, R.: A complete scheme for computing all direct and cross demand elasticities in a model with many sectors. Econometrica **27**(2), 177–196 (1959)

Gabaix, X.: Power laws in economics and finance. Ann. Rev. Econ. **1**, 255–293 (2009)

Gabaix, X., Gopikrishnan, P., Plerou, V., Stanley, H.E.: A theory of power-law distributions in financial market fluctuations. Nature **423**(6937), 267–270 (2003)

Gabaix, X., Gopikrishnan, P., Plerou, V., Stanley, H.E.: Institutional investors and stock market volatility. Q. J. Econ. **121**(2), 461–504 (2006)

Gali, J.: Introduction to Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework Princeton University Press (2008)

Gorton, G.B., Hayashi, F., Rouwenhorst, K.G.: The fundamentals of commodity futures returns. Eur. Finan. Rev. **17**(1), 35–105 (2012)

Hall, P., Horowitz, J.L., Jing, B.Y.: On blocking rules for the bootstrap with dependent data. Biometrika **82**(3), 561–574 (1995)

Harri, A., Erdem, C., Coble, K.H., Knight, T.O.: Crop yield distributions: A reconciliation of previous research and statistical tests for normality. Rev. Agric. Econ. **31**(1), 163–182 (2009)

Hartmann, P., Straetmans, S., de Vries, C.G.: Banking system stability: A cross-atlantic perspective. In: Carey, M., Stulz, R.M. (eds.) The Risks of Financial Institutions, pp. 133–192. University of Chicago Press (2006)

Hill, B.M.: A simple general approach to inference about the tail of a distribution. Ann. Stat. **3**(5), 1163–1174 (1975)

Hyung, N., De Vries, C.G.: Portfolio diversification effects of downside risk. J. Financ. Econ. **3**(1), 107–125 (2005)

Kat, H.M., Oomen, R.C.A.: What every investor should know about commodities I: A univariate return analysis. J. Invest. Manag. **5**, 16–40 (2006)

Knittel, C.R., Pindyck, R.S.: The simple economics of commodity price speculation. Amer. Econ. J.: Macroecon. **8**(2), 85–110 (2016)

Kofman, P., De Vries, C.G.: Potato futures returns: A tail investigation. Rev. Fut. Markets **8**(2), 244–258 (1990)

Koundouri, P., Kourogenis, N.: On the distribution of crop yields: Does the central limit theorem apply? Am. J. Agric. Econ. **93**(5), 1341–1357 (2011)

Krehbiel, T., Adkins, L.C.: Price risk in the nymex energy complex: An extreme value approach. J. Futur. Mark. **25**(4), 309–337 (2005)

Kupiec, P.H.: Techniques for verifying the accuracy of risk measurement models. J. Deriv. **3**(2), 73–84 (1995)

Mandelbrot, B.: The variation of certain speculative prices. J. Business **36**(4), 394–419 (1963a)

Mandelbrot, B.: New methods in statistical economics. J. Polit. Econ. **71**(5), 421–440 (1963b)

Matia, K., Amaral, L.A.N., Goodwin, S.P., Stanley, H.E.: Different scaling behaviors of commodity spot and future prices. Phys. Rev. E. **66**(045103), 1–4 (2002)

Mills, F.C.: The behavior of prices. National Bureau of Economic Research (1927)

Morgan, W., Cotter, J., Dowd, K.: Extreme measures of agricultural financial risk. J. Agric. Econ. **63**(1), 65–82 (2012)

Newman, M.E.J.: Power laws, Pareto distributions and Zipf's law. Contemp. Phys. **46**(5), 323–351 (2005)

Ramírez, O.A., Misra, S., Field, J.: Crop-yield distributions revisited. Am. J. Agric. Econ. **85**(1), 108–120 (2003)

Reed, W.J.: The Pareto, Zipf and other power laws. Economics letters **74**(1), 15–19 (2001)

Reichling, F., Wahlen, C.: Review of estimates of the Frisch elasticity of labor supply. Congressional Budget Office Working Paper 2012–12 (2012)

Roy, A.D.: Safety first and the holding of assets. Econometrica **20**(3), 431–449 (1952)

Salvadori, G., De Michele, C., Kottegoda, N.T., Rosso, R.: Extremes in nature: An approach using copulas. Springer (2007)

Telser, L.G.: Safety first and hedging. Rev. Econ. Stud. **23**(1), 1–16 (1955)

Toda, A.A.: Incomplete market dynamics and cross-sectional distributions. J. Econ. Theory **154**, 310–348 (2014)

Toda, A.A., Walsh, K.: The double power law in consumption and implications for testing euler equations. J. Polit. Econ. **123**(5), 1177–1200 (2015)

Turnovsky, S.J.: Technological and price uncertainty in a Ricardian model of international trade. Rev. Econ. Stud. **41**(2), 201–217 (1974)

Van Oordt, M.R.C., Zhou, C.: Systematic tail risk. J. Financ. Quant. Anal. **51**(2), 685–705 (2016)

Walsh, C.E.: Monetary Theory and Policy, 4th edn. M.I.T. Press (2017)

Wang, D., Tomek, W.G.: Commodity prices and unit root tests. Am. J. Agric. Econ. **89**(4), 873–889 (2007)

Williams, J.C., Wright, B.D.: Storage and Commodity Markets. Cambridge University Press (2005)

Wright, B.D., Williams, J.C.: The economic role of commodity storage. Econ. J. **92**(367), 596–614 (1982)