Empirical Methods for Dynamic Power Law Distributions in the Social Sciences

Ricardo T. Fernholz*

Claremont McKenna College

February 10, 2017

Abstract

This paper introduces nonparametric econometric methods that characterize general power law distributions under basic stability conditions. These methods extend the literature on power laws in the social sciences in several directions. First, we show that any stationary distribution in a random growth setting is shaped entirely by two factors—the idiosyncratic volatilities and reversion rates (a measure of cross-sectional mean reversion) for different ranks in the distribution. This result is valid regardless of how growth rates and volatilities vary across different economic agents, and hence applies to Gibrat's law and its extensions. Second, we present techniques to estimate these two factors using panel data. Third, we describe how our results imply predictability as higher-ranked processes must on average grow more slowly than lower-ranked processes. We employ our empirical methods using data on commodity prices and show that our techniques accurately describe the empirical distribution of relative commodity prices. We also show that rank-based out-of-sample forecasts of future commodity prices outperform random-walk forecasts at a one-month horizon.

JEL Codes: C10, C14

Keywords: power laws, Pareto distribution, Gibrat's law, Zipf's law, nonparametric methods, commodity prices, forecasting

---

I would like to thank Christoffer Koch, Serkan Ozbeklik, Angela Vossmeyer, and seminar participants at the Federal Reserve Bank of Dallas and the Federal Reserve Bank of Philadelphia for their helpful comments. All remaining errors are my own.

*Robert Day School of Economics and Finance, Claremont McKenna College, 500 E. Ninth St., Claremont, CA 91711, rfernholz@cmc.edu.

1 Introduction

Power laws are ubiquitous in economics, finance, and the social sciences more broadly. They are found across many different phenomena, ranging from the distribution of income and wealth (Atkinson et al., 2011; Piketty, 2014) to the city size distribution (Gabaix, 1999) to the distribution of assets of financial intermediaries (Janicki and Prescott, 2006; Fernholz and Koch, 2016). Although a number of potential mechanisms explaining the appearance of power laws have been proposed (Newman, 2005; Gabaix, 2009), one of the most broad and influential involves random growth processes.

A large literature in economics, both theoretical and empirical, models different power laws and Pareto distributions as the result of random growth processes that are stabilized by the presence of some friction (Champernowne, 1953; Gabaix, 1999; Luttmer, 2007; Benhabib et al., 2011). In this paper, we present rank-based, nonparametric methods that allow for the characterization of general power law distributions in any continuous random growth setting. These techniques, which are well-established and the subject of active research in statistics and mathematical finance, are general and can be applied to Gibrat's law and many of its extensions in economics and finance.[1] According to our general characterization, any stationary distribution in a random growth setting is shaped entirely by two factors—the idiosyncratic volatilities and reversion rates (a measure of cross-sectional mean reversion) for different ranks in the distribution. An increase in idiosyncratic volatilities increases concentration, while an increase in reversion rates decreases concentration. We also present results that allow for the estimation of these two factors using panel data.

Our characterization of a stationary distribution in a general, nonparametric setting provides a framework in which we can understand the shaping forces for almost all power law distributions that emerge in random growth settings. After all, one implication of our results is that the distributional effect of any economic mechanism can be inferred by determining the effect of that mechanism on the idiosyncratic volatilities and reversion rates for different ranked processes. In addition to our general characterization, this is, to our knowledge, the first paper in economics to provide empirical methods to measure the econometric factors that shape power law distributions using panel data.

These empirical methods allow us to understand the causes, in an econometric sense, of

[1] There is a growing and extensive literature analyzing these rank-based methods. See, for example, Banner et al. (2005), Pal and Pitman (2008), Ichiba et al. (2011), and Shkolnikov (2011).

distributional changes that occur for power laws in economics and finance. This econometric analysis has many potential applications. For example, our methods can establish that increasing U.S. income and wealth inequality (Atkinson et al., 2011; Saez and Zucman, 2016) is the result of changes in either reversion rates or the magnitude of idiosyncratic shocks to household income and wealth. These econometric changes could then be linked to the evolution of policy, skill-biased technological change, or other changes in the economic environment. This is the approach of Fernholz and Koch (2016), who analyze the increasing concentration of U.S. bank assets using the econometric techniques presented in this paper. A similar analysis of increasing U.S. house price dispersion (Van Nieuwerburgh and Weill, 2010) should also yield new conclusions and useful insight.

One of this paper's main contributions is to present a new and more general approach to power laws in economics and finance. Indeed, the most common approach to modeling random growth processes and power laws involves solving a single stochastic differential equation that yields a parametric distribution representing a continuum of agents, an approach that started with Gabaix (1999). In contrast, our approach considers a discrete system of multiple stochastic differential equations in which no parametric assumptions are imposed. This generality is useful for many applications, since empirical distributions often do not conform to a single power law at all points (Axtell, 2001; Ioannides and Skouras, 2013).

Our general characterization of the stationary distribution in a random growth setting imposes a piecewise linear relationship between log size and log rank—a power law distribution in which the power law exponent can change at every rank. We show that the power law exponent varies across ranks in the same way as the reversion rates and idiosyncratic volatilities of the underlying random growth process vary across ranks. By varying these two shaping factors across ranks, then, our methods can replicate any empirical distribution, including those which do not conform to a single power law at all points. To our knowledge, this is the first paper in economics to achieve such generality.

In order to demonstrate the validity and accuracy of our empirical methods, we estimate reversion rates and idiosyncratic volatilities using monthly commodity prices data from 1980 - 2015 and compare the predicted distribution of relative commodity prices using our methods to the average distribution of relative commodity prices observed during this period. Although our methods apply most naturally to distributions such as wealth, firm size, and city size, they can also be applied to the distributions of relative asset prices. By testing our methods using normalized commodity prices data, we are able to examine the applicability of these methods to relative asset price distributions.

Commodity prices must be normalized so that they can be compared in an economically meaningful way, but as long as these appropriately-normalized prices satisfy the basic regularity conditions that our econometric theory relies on, then our methods should be applicable to the distribution of relative commodity prices. Furthermore, because the distribution of relative normalized commodity prices appears to be stationary during the 1980 - 2015 period, the rank-based reversion rates and idiosyncratic volatilities that we estimate should provide an accurate description of the observed relative commodity price distribution during this period. We confirm that this is in fact the case. One of the contributions of this paper, then, is to show that our empirical methods can validly be applied not only to standard size distributions but also to relative asset price distributions. This result highlights the potential for future applications of our econometric techniques using other data sets.

In addition to our characterization of a stationary distribution in a general random growth setting, we also show that a cross-sectional mean-reversion condition is necessary for the existence of such a stationary distribution. Specifically, a stationary distribution exists only if the growth rates of higher-ranked processes are on average lower than the growth rates of lower-ranked processes. This condition offers a testable prediction of our econometric theory. Specifically, commodity price rank should forecast future commodity prices. We confirm that this is in fact the case in two different ways. First, we show that the log growth rates of all subsets of higher-ranked commodities are lower than the log growth rates of lower-ranked commodities and that these differences are highly statistically significant over the 1980 - 2015 time period. Second, we use our econometric results to generate out-of-sample forecasts of one-month-ahead commodity prices from 1990 - 2015, and show that these rank-based forecasts outperform random-walk forecasts in all cases. In many cases, these rank-based forecasts outperform random-walk forecasts by more than forecasts based on factor models or commodity price fundamentals (Alquist and Coibion, 2014; West and Wong, 2014).

The rest of this paper is organized as follows. Section 2 presents our nonparametric framework and derives the main result that characterizes general stationary power law distributions. Section 3 presents results that show how to estimate the two shaping factors of a power law distribution using panel data. Section 4 presents estimates of rank-based reversion rates and idiosyncratic volatilities using commodity prices data, and also examines rank-based forecasts of future commodity prices derived from our econometric results. Section 5 concludes. Appendix A investigates the efficiency of local-time-based reversion-rate estimators, Appendix B discusses the regularity assumptions needed for our main results, and Appendix C contains all proofs.

2 A Nonparametric Approach to Dynamic Power Law Distributions

For consistency, we shall refer to agents holding units throughout this section. However, it is important to note that in this general setup agents can represent households, firms, cities, countries, and other entities, with the corresponding units representing income, wealth, total employees, population, and other quantities. Furthermore, we can also interpret agents' holdings of units as the prices of different assets, as we shall do for commodity prices in Section 4 below.

Consider a population that consists of $N > 1$ agents. Time is continuous and denoted by $t \in [0, \infty)$, and uncertainty in this population is represented by a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$. Let $\mathbf{B}(t) = (B_1(t), \ldots, B_M(t))$, $t \in [0, \infty)$, be an M-dimensional Brownian motion defined on the probability space, with $M \geq N$. We assume that all stochastic processes are adapted to $\{\mathcal{F}_t; t \in [0, \infty)\}$, the augmented filtration generated by $\mathbf{B}$.[2]

2.1 Dynamics

The total units held by each agent $i = 1, \ldots, N$ is given by the process x_i. Each of these unit processes evolves according to the stochastic differential equation

$$d \log x_i(t) = \mu_i(t)\, dt + \boldsymbol{\delta}_i(t) \cdot d\mathbf{B}(t), \tag{2.1}$$

[2]In order to simplify the exposition, we shall omit many of the less important regularity conditions and technical details involved with continuous-time stochastic processes.

where μ_i and $\boldsymbol{\delta}_i = (\delta_{i1}, \ldots, \delta_{iM})$ are measurable and adapted processes. The growth rates and volatilities, μ_i and $\boldsymbol{\delta}_i$, respectively, are general and practically unrestricted, having only to satisfy a few basic regularity conditions that are discussed in Appendix B. These conditions imply that the unit processes for the agents are Itô processes, which represent a broad class of stochastic processes that are common in the continuous-time finance literature (for a detailed discussion, see Karatzas and Shreve, 1991).

Indeed, the martingale representation theorem (Nielsen, 1999) implies that any plausible continuous process for agents' unit holdings can be written in the nonparametric form of equation (2.1).[3] The $M \geq N$ sources of volatility in this equation allow for a rich structure of time-varying idiosyncratic, correlated, and aggregate shocks to agents' unit holdings that need not conform to any particular distribution. Furthermore, equation (2.1) also allows for unit growth rates and volatilities that vary across individual agents based on any characteristics. The generality of this nonparametric framework implies that our econometric results are consistent with essentially any underlying economic model, including models based on Gibrat's law or specific extensions to Gibrat's law (Gabaix, 1999, 2009).

It is useful to describe the dynamics of the total units held by all agents, which we denote by $x(t) = x_1(t) + \cdots + x_N(t)$. In order to do so, we first characterize the covariance of unit holdings across different agents over time. For all $i, j = 1, \ldots, N$, let the covariance process ρ_{ij} be given by

$$\rho_{ij}(t) = \boldsymbol{\delta}_i(t) \cdot \boldsymbol{\delta}_j(t). \tag{2.2}$$

Applying Itô's Lemma to equation (2.1), we are now able to describe the dynamics of the total units process x.

Lemma 2.1. *The dynamics of the process for total units held by all agents x are given by*

$$d \log x(t) = \mu(t)\, dt + \sum_{i=1}^{N} \theta_i(t) \boldsymbol{\delta}_i(t) \cdot d\mathbf{B}(t), \quad \text{a.s.}, \tag{2.3}$$

where

$$\theta_i(t) = \frac{x_i(t)}{x(t)}, \tag{2.4}$$

[3] Many of this section's results can also apply to processes that are subject to sporadic, discontinuous jumps. This is an open area for research, with such extensions examined by Shkolnikov (2011) and Fernholz (2016a).

for $i = 1, \ldots, N$, *and*

$$\mu(t) = \sum_{i=1}^{N} \theta_i(t)\mu_i(t) + \frac{1}{2}\left(\sum_{i=1}^{N} \theta_i(t)\rho_{ii}(t) - \sum_{i,j=1}^{N} \theta_i(t)\theta_j(t)\rho_{ij}(t)\right). \tag{2.5}$$

2.2 Rank-Based Dynamics

In order to characterize the stationary distribution of units in this setup, it is necessary to consider the dynamics of agents' unit holdings by rank. One of the key insights of our approach and of this paper more generally is that rank-based unit dynamics are the essential determinants of the distribution of units. As we demonstrate below, there is a simple, direct, and robust relationship between rank-based unit growth rates and the distribution of units. This relationship is a purely statistical result and hence can be applied to essentially any economic environment, no matter how complex.

The first step in achieving this characterization is to introduce notation for agent rank based on unit holdings. For $k = 1, \ldots, N$, let

$$x_{(k)}(t) = \max_{1 \le i_1 < \cdots < i_k \le N} \min\left(x_{i_1}(t), \ldots, x_{i_k}(t)\right), \tag{2.6}$$

so that $x_{(k)}(t)$ represents the units held by the agent with the k-th most units among all the agents in the population at time t. For brevity, we shall refer to this agent as the k-th largest agent throughout this paper. One consequence of this definition is that

$$\max(x_1(t), \ldots, x_N(t)) = x_{(1)}(t) \ge x_{(2)}(t) \ge \cdots \ge x_{(N)}(t) = \min(x_1, \ldots, x_N(t)). \tag{2.7}$$

Next, let $\theta_{(k)}(t)$ be the share of total units held by the k-th largest agent at time t, so that

$$\theta_{(k)}(t) = \frac{x_{(k)}(t)}{x(t)}, \tag{2.8}$$

for $k = 1, \ldots, N$.

One of this paper's central contributions is to provide a characterization of the asymptotic behavior of the unit shares held by every single ranked agent $\theta_{(k)}$. This characterization in fact describes the entire distribution of agents' unit holdings. To see this, note that the statement that the agent with the k-th largest unit holdings holds $\theta_{(k)}$ units is equivalent to the statement that k agents hold more than $\theta_{(k)}$ units. Of course, this latter statement yields the probability of observing unit holdings greater than $\theta_{(k)}$, which implies that the unit shares $\theta_{(k)}$ describe the cumulative distribution function (CDF) of the distribution of unit holdings. In other words, if we can describe the asymptotic behavior of each $\theta_{(k)}$, $k = 1, \ldots, N$, then we can also describe the asymptotic distribution of agents' unit holdings in full.

The next step in our approach is to describe the dynamics of the agent rank unit processes $x_{(k)}$ and rank unit share processes $\theta_{(k)}$, $k = 1, \ldots, N$. This task is complicated by the fact that the max and min functions from equation (2.6) are not differentiable, and hence we cannot simply apply Itô's Lemma in this case. Instead, we introduce the notion of a local time to solve this problem. For any continuous process z, the *local time* at 0 for z is the process Λ_z defined by

$$\Lambda_z(t) = \frac{1}{2}\left(|z(t)| - |z(0)| - \int_0^t \operatorname{sgn}(z(s))\, dz(s)\right). \tag{2.9}$$

As detailed by Karatzas and Shreve (1991), the local time for z measures the amount of time the process z spends near zero and can also be defined as

$$\Lambda_z(t) = \lim_{\epsilon \downarrow 0} \frac{1}{2\epsilon} \int_0^t 1_{\{|z(s)|<\epsilon\}}\, ds. \tag{2.10}$$

To be able to link agent rank to agent index, let p_t be the random permutation of $\{1, \ldots, N\}$ such that for $1 \leq i, k \leq N$,

$$p_t(k) = i \quad \text{if} \quad x_{(k)}(t) = x_i(t). \tag{2.11}$$

This definition implies that $p_t(k) = i$ whenever agent i is the k-th largest agent in the population at time t, with ties broken in some consistent manner.[4]

Lemma 2.2. *For all $k = 1, \ldots, N$, the dynamics of the agent rank unit processes $x_{(k)}$ and rank unit share processes $\theta_{(k)}$ are given by*[5]

$$d\log x_{(k)}(t) = d\log x_{p_t(k)}(t) + \frac{1}{2}d\Lambda_{\log x_{(k)} - \log x_{(k+1)}}(t) - \frac{1}{2}d\Lambda_{\log x_{(k-1)} - \log x_{(k)}}(t), \tag{2.12}$$

---

[4]For example, if $x_i(t) = x_j(t)$ and $i > j$, then we can set $p_t(k) = i$ and $p_t(k+1) = j$.

[5]For brevity, we write $dz_{p_t(k)}(t)$ to refer to the process $\sum_{i=1}^N 1_{\{i=p_t(k)\}} dz_i(t)$ throughout this paper.

a.s, and

$$d\log\theta_{(k)}(t) = d\log\theta_{p_t(k)}(t) + \frac{1}{2}d\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}(t) - \frac{1}{2}d\Lambda_{\log\theta_{(k-1)}-\log\theta_{(k)}}(t), \tag{2.13}$$

a.s., with the convention that $\Lambda_{\log x_{(0)}-\log x_{(1)}}(t) = \Lambda_{\log x_{(N)}-\log x_{(N+1)}}(t) = 0$.

According to equation (2.12) from the lemma, the dynamics of units for the k-th largest agent in the population are the same as those for the agent that is the k-th largest at time t (agent $i = p_t(k)$), plus two local time processes that capture changes in agent rank (one agent overtakes another in unit holdings) over time. Equation (2.13) describes the similar dynamics of the rank unit share processes $\theta_{(k)}$.

Using equations (2.1) and (2.3) and the definition of $\theta_i(t)$, we have that for all $i = 1,\ldots,N$,

$$\begin{aligned} d\log\theta_i(t) &= d\log x_i(t) - d\log x(t) \\ &= \mu_i(t)\,dt + \boldsymbol{\delta}_i(t)\cdot d\mathbf{B}(t) - \mu(t)\,dt - \sum_{i=1}^{N}\theta_i(t)\boldsymbol{\delta}_i(t)\cdot d\mathbf{B}(t). \end{aligned} \tag{2.14}$$

If we apply Lemma 2.2 to equation (2.14), then it follows that

$$\begin{aligned} d\log\theta_{(k)}(t) = \big(\mu_{p_t(k)}(t) - \mu(t)\big)\;dt + \boldsymbol{\delta}_{p_t(k)}(t)\cdot d\mathbf{B}(t) - \sum_{i=1}^{N}\theta_i(t)\boldsymbol{\delta}_i(t)\cdot d\mathbf{B}(t) \\ + \frac{1}{2}d\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}(t) - \frac{1}{2}d\Lambda_{\log\theta_{(k-1)}-\log\theta_{(k)}}(t), \end{aligned} \tag{2.15}$$

a.s, for all $k = 1,\ldots,N$. Equation (2.15), in turn, implies that the process $\log\theta_{(k)} - \log\theta_{(k+1)}$ satisfies, a.s., for all $k = 1,\ldots,N-1$,

$$\begin{aligned} d\left(\log\theta_{(k)}(t) - \log\theta_{(k+1)}(t)\right) = \big(\mu_{p_t(k)}(t) - \mu_{p_t(k+1)}(t)\big)\;dt + d\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}(t) \\ - \frac{1}{2}d\Lambda_{\log\theta_{(k-1)}-\log\theta_{(k)}}(t) - \frac{1}{2}d\Lambda_{\log\theta_{(k+1)}-\log\theta_{(k+2)}}(t) \\ + \left(\boldsymbol{\delta}_{p_t(k)}(t) - \boldsymbol{\delta}_{p_t(k+1)}(t)\right)\cdot d\mathbf{B}(t). \end{aligned} \tag{2.16}$$

The processes for relative unit holdings of adjacent agents in the distribution of units as given by equation (2.16) are key to describing the distribution of units in this setup.

2.3 Stationary Distribution

The results presented above allow us to analytically characterize the stationary distribution of units in this setup. Let α_k equal the time-averaged limit of the expected growth rate of units for the k-th largest agent relative to the expected growth rate of units for the entire population of agents, so that

$$\alpha_k = \lim_{T\to\infty} \frac{1}{T} \int_0^T \left(\mu_{p_t(k)}(t) - \mu(t)\right)\, dt, \tag{2.17}$$

for $k = 1, \ldots, N$. The relative growth rates α_k are a rough measure of the rate at which agents' unit holdings cross-sectionally revert to the mean. We shall refer to the $-\alpha_k$ as *reversion rates*, since lower values of α_k (and hence higher values of $-\alpha_k$) imply faster cross-sectional mean reversion.

In a similar manner, we wish to define the time-averaged limit of the volatility of the process $\log\theta_{(k)} - \log\theta_{(k+1)}$, which measures the relative unit holdings of adjacent agents in the distribution of units. For all $k = 1, \ldots, N-1$, let σ_k be given by

$$\sigma_k^2 = \lim_{T\to\infty} \frac{1}{T} \int_0^T \left\| \boldsymbol{\delta}_{p_t(k)}(t) - \boldsymbol{\delta}_{p_t(k+1)}(t) \right\|^2 dt. \tag{2.18}$$

The relative growth rates α_k together with the volatilities σ_k entirely determine the shape of the stationary distribution of units in this population, as we shall demonstrate below.

We shall refer to the volatility parameters σ_k, which measure the standard deviations of the processes $\log\theta_{(k)} - \log\theta_{(k+1)}$, as *idiosyncratic volatilities*. An idiosyncratic shock to the unit holdings of either the k-th or $(k+1)$-th ranked agent alters the value of $\log\theta_{(k)} - \log\theta_{(k+1)}$ and hence will be measured by σ_k. In addition, however, a shock that affects the unit holdings of multiple agents that do not occupy adjacent ranks in the distribution will also alter this value. Indeed, any shock that affects $\log\theta_{(k)}$ and $\log\theta_{(k+1)}$ differently, must necessarily alter the value of $\log\theta_{(k)} - \log\theta_{(k+1)}$ and hence will be measured by σ_k. In this sense, the volatility parameters σ_k are slightly more general than pure idiosyncratic volatilities that capture only shocks that affect one single agent at a time.

Finally, for all $k = 1, \ldots, N$, let

$$\kappa_k = \lim_{T\to\infty} \frac{1}{T} \Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(T). \tag{2.19}$$

Let $\kappa_0 = 0$, as well. Throughout this paper, we assume that the limits in equations (2.17)-(2.19) do in fact exist. In Appendix C, we show that the parameters α_k and κ_k are related by $\alpha_k - \alpha_{k+1} = \frac{1}{2}\kappa_{k-1} - \kappa_k + \frac{1}{2}\kappa_{k+1}$, for all $k = 1, \ldots, N-1$.

The *stable version* of the process $\log \theta_{(k)} - \log \theta_{(k+1)}$ is the process $\log \theta^*_{(k)} - \log \theta^*_{(k+1)}$ defined by

$$d\left(\log \theta^*_{(k)}(t) - \log \theta^*_{(k+1)}(t)\right) = -\kappa_k \, dt + d\Lambda_{\log \theta^*_{(k)} - \log \theta^*_{(k+1)}}(t) + \sigma_k \, dB(t), \tag{2.20}$$

for all $k = 1, \ldots, N-1$.[6] The stable version of $\log \theta_{(k)} - \log \theta_{(k+1)}$ replaces all of the processes from the right-hand side of equation (2.16) with their time-averaged limits, with the exception of the local time process $\Lambda_{\log \theta_{(k)} - \log \theta_{(k+1)}}$. By considering the stable version of these relative unit holdings processes, we are able to obtain a simple characterization of the distribution of units.

Theorem 2.3. *There is a stationary distribution for the stable version of unit holdings by agents in this population if and only if $\alpha_1 + \cdots + \alpha_k < 0$, for $k = 1, \ldots, N-1$. Furthermore, if there is a stationary distribution of units, then for $k = 1, \ldots, N-1$, this distribution satisfies*

$$E\left[\log \theta^*_{(k)}(t) - \log \theta^*_{(k+1)}(t)\right] = \frac{\sigma_k^2}{-4(\alpha_1 + \cdots + \alpha_k)}, \quad \text{a.s.} \tag{2.21}$$

Theorem 2.3 provides an analytic rank-by-rank characterization of the entire distribution of units. This is achieved despite minimal assumptions on the processes that describe the dynamics of agents' unit holdings over time. As long as the relative growth rates, volatilities, and local times that we take limits of in equations (2.17)-(2.19) do not change drastically and frequently over time, then the distribution of the stable versions of $\theta_{(k)}$ from Theorem 2.3 will accurately reflect the distribution of the true versions of these rank unit share processes.[7]

The theorem yields a system of $N-1$ equations which together with the identity $\theta_{(1)} + \cdots + \theta_{(N)} = 1$ can be solved for the unit shares held by every single ranked agent $\theta_{(k)}$. As discussed in Section 2.2 above, this description of each $\theta_{(k)}$ is equivalent to a description of the cumulative distribution function (CDF) of the distribution of unit holdings. One of

[6] For each $k = 1, \ldots, N-1$, equation (2.20) implicitly defines another Brownian motion $B(t)$, $t \in [0, \infty)$. These Brownian motions can covary in any way across different k.

[7] Fernholz (2002) and Fernholz and Koch (2016) demonstrate the accuracy of Theorem 2.3 in matching, respectively, the distribution of total market capitalizations of U.S. stocks and the distribution of assets of U.S. financial intermediaries.

the most important implications of equation 2.21 from Theorem 2.3 is that the only two econometric factors that affect the distribution of unit holdings are the rank-based reversion rates, $-\alpha_k$, and the rank-based volatilities, σ_k. This implies that it is not necessary to directly model and estimate agents' unit holdings dynamics by name, denoted by index i, as is the usual approach in economics.[8] Instead, an understanding of unit holdings dynamics by rank, denoted by k, is sufficient to describe the entire distribution of units.

2.4 Discussion: Power Laws, Gibrat's Law, and Theorem 2.3

Our characterization of the distribution of agents' unit holdings in Theorem 2.3 is different from what is standard for power laws in economics and finance. Usually, a single stochastic differential equation is solved and this solution yields a parametric distribution that represents a continuum of agents. This single-equation approach was introduced to economics by Gabaix (1999), and has since been used in many applications (Luttmer, 2007; Benhabib et al., 2011; Gabaix et al., 2016).

In contrast to this literature, our approach involves solving a discrete system of multiple stochastic differential equations. This granularity is essential for real-world applications, since there is never a continuum of agents in the data. Indeed, the practical importance of this novel feature of our approach is apparent in Sections 3 and 4, where we first introduce methods to estimate the rank-based reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k for each discrete rank $k = 1, \ldots, N-1$ using panel data, and then apply this estimation procedure to commodity price data from 1980 - 2015. This procedure is more flexible than the standard single-equation approach, and this flexibility allows for more precision when describing the distribution of relative commodity prices from 1980 - 2015.

A second advantage of our approach is that it imposes no parametric assumptions about the stationary distribution of agents' unit holdings. Indeed, equation (2.21) from Theorem 2.3 is flexible enough to replicate any empirical distribution given appropriate values for the reversion rates and volatilities. To see this, note that the distribution of agents' unit holdings follows a power law if the relationship between log-unit shares and log rank is linear, at least for the highest ranks (Newman, 2005; Gabaix, 2009). According to equation (2.21) from Theorem 2.3, for all $k = 1, \ldots, N-1$, the slope of a log-unit shares versus log rank plot is

[8]In the literatures on income and wealth inequality, for example, the standard approach is to model the income or wealth dynamics of individual households rather than ranked households (Guvenen, 2009; Benhabib et al., 2011; Altonji et al., 2013).

given by

$$\frac{E\left[\log\theta^*_{(k)}(t) - \log\theta^*_{(k+1)}(t)\right]}{\log(k) - \log(k+1)} \approx \frac{-(k+0.5)\sigma_k^2}{-4(\alpha_1 + \cdots + \alpha_k)}, \tag{2.22}$$

where we use the asymptotic approximation $\log(k) - \log(k+1) \approx -(k+0.5)^{-1}$. Equation (2.22) characterizes a power law distribution for unit shares in that it imposes piecewise linearity—a different linear relationship between ranks 1 and 2, 2 and 3, and so on up to ranks $N-1$ and N. As the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k vary across different ranks in the distribution, equation (2.22) shows that the slope of the log-unit shares versus log rank plot varies correspondingly.

According to equation (2.22), then, our general methods allow for a power law relationship that can vary across every single rank in the distribution of agents' unit holdings. To our knowledge, our approach is the first in economics or finance to achieve such generality. As with granularity, this generality is essential for many applications, since many empirical distributions do not uniformly conform to one single power law. There is ample evidence, for example, that income and wealth distributions follow power laws at the top while appearing more lognormal at lower levels of income and wealth (Guvenen, 2009; Atkinson et al., 2011; Benhabib et al., 2011). The same basic pattern has been documented for city size distributions as well (Eeckhout, 2004; Ioannides and Skouras, 2013). Similarly, firm size distributions usually follow power laws that vary across different parts of the distribution, regardless of whether firm size is measured by total employees (Axtell, 2001; Luttmer, 2007), total market capitalization (Fernholz, 2002), or total assets (Fernholz and Koch, 2016). The fact that these real-world applications do not conform to a single distribution fits easily into our empirical framework because of the flexibility and lack of parametric assumptions of Theorem 2.3. Indeed, equations (2.21) and (2.22) imply that these changing or partial power laws are the result of reversion rates and idiosyncratic volatilities that vary correspondingly across different ranks in the distribution.

Theorem 2.3 generalizes far beyond the equal growth rates and volatilities imposed by Gibrat's law (Gabaix, 2009). The strongest form of Gibrat's law for unit holdings requires growth rates and volatilities that do not vary across the distribution of unit holdings. In terms of the reversion rates $-\alpha_k$ (which measure relative unit growth rates for different ranked agents) and idiosyncratic volatilities σ_k, this requirement is equivalent to there existing some common $\alpha < 0$ and $\sigma > 0$ such that

$$\alpha = \alpha_1 = \cdots = \alpha_{N-1}, \tag{2.23}$$

and

$$\sigma = \sigma_1 = \cdots = \sigma_{N-1}. \tag{2.24}$$

In terms of equation (2.21) from Theorem 2.3, then, Gibrat's law yields unit shares that satisfy

$$\frac{E\left[\log\theta^*_{(k)}(t) - \log\theta^*_{(k+1)}(t)\right]}{\log(k) - \log(k+1)} \approx \frac{-(k+0.5)\sigma_k^2}{-4(\alpha_1 + \cdots + \alpha_k)} = \frac{-(k+0.5)\sigma^2}{-4k\alpha} = \frac{\sigma^2}{4\alpha} + \frac{\sigma^2}{8k\alpha} \tag{2.25}$$

for all $k = 1, \ldots, N-1$.

The distribution of agents' unit holdings follows a Pareto distribution if a log-unit shares versus log rank plot appears as a straight line (Newman, 2005; Gabaix, 2009). Furthermore, if the slope of such a straight line plot is -1, then agents' unit shares obey Zipf's law (Gabaix, 1999). Equation (2.25) shows that Gibrat's law yields a Pareto distribution in which, for large k, the log-log plot of unit shares versus rank has slope $\sigma^2/4\alpha < 0$, which is equivalent to the Pareto distribution having parameter $-4\alpha/\sigma^2 > 0$. Furthermore, we see that agents' unit shares obey Zipf's law only if $\sigma^2 = -4\alpha$, in which case the log-log plot has slope -1, for large k.

One interesting implication of Theorem 2.3 and equation (2.22) relates to a result of Gabaix (1999) regarding deviations from Zipf's law for cities. Gabaix (1999) uses a single-equation approach to show that a process with higher volatility at lower city sizes produces a stationary distribution in which the power law varies with size. In particular, he shows that in this case a plot of log size versus log rank is steeper for smaller cities. We can see that equation (2.22) confirms this result, since higher values of the volatility parameters σ_k for large k (lower ranks) mean a steeper curve at the lower ranks. In other words, Theorem 2.3 implies that a negative relationship between unit shares and volatility yields a concave log-log plot of unit shares versus rank.

Theorem 2.3 also extends this result of Gabaix (1999) in several directions. First, we see that it is not only idiosyncratic volatilities σ_k that are increasing in k that generates a concave log-log plot of unit shares versus rank, but also reversion rates $-\alpha_k$ that are decreasing in k. Indeed, according to equation (2.22), the slope of each linear piece of such a log-log plot is increasing in σ_k and decreasing in $-(\alpha_1 + \cdots + \alpha_k)/k$. Second, because we adopt a nonparametric approach using general Itô processes in equation (2.1), our results in Theorem 2.3 and equation (2.22) provide both necessary and sufficient conditions for any power law exponent to obtain in any part of the distribution curve. In other words, not only do our results confirm the result of Gabaix (1999) regarding volatility and city size, but they also show that the *only* two factors that can alter the power law exponent of the stationary distribution of continuous random growth processes are the rank-based reversion rates $-\alpha_k$ and volatility parameters σ_k.

Perhaps most importantly, the precise definitions of the parameters α_k and σ_k in equations (2.17) and (2.18) econometrically define the growth rates and volatilities that shape a dynamic power law distribution. Equation (2.18) shows that the time-averaged volatilities of the processes $\log \theta_{(k)} - \log \theta_{(k+1)}$, $k = 1, \ldots, N-1$, are the volatility parameters that shape the distribution of agents' unit holdings. As mentioned above, these volatility parameters σ_k do measure idiosyncratic shocks to agents' unit holdings as in Gabaix (1999), but they also measure correlated shocks that alter the relative unit holdings of adjacent ranked agents $\log \theta_{(k)} - \log \theta_{(k+1)}$. Similarly, equation (2.17) provides a precise econometric definition of the growth rates that shape the distribution of agents' unit holdings. According to this equation, these growth rate parameters α_k measure the log growth rates of unit holdings for each rank k relative to the log growth rate of all units in the economy—a measure of cross-sectional mean reversion. To our knowledge, these are new results in economics.

Theorem 2.3 demonstrates that Gibrat's law and the Pareto distributions it generates are special cases of more general processes in which growth rates and volatilities potentially vary across different ranks in the distribution of unit holdings. This flexibility is a novel feature of our empirical methodology and is necessary to accurately match many empirical distributions. For example, Fernholz and Koch (2016) find that asset growth rates and volatilities vary substantially across different size-ranked U.S. financial intermediaries. Similarly, Fernholz (2002) finds that growth rates and volatilities of total market capitalization vary substantially across different size-ranked U.S. stocks, while Neumark et al. (2011) find that employment growth rates vary across different size-ranked U.S. firms. In Section 4, we confirm this general pattern and show that the growth rates of commodity prices also differ across ranks in a statistically significant and economically meaningful way.

According to Theorem 2.3, stationarity of the distribution of agents' unit holdings requires that the reversion rates $-\alpha_k$ must sum to positive quantities, for all $k = 1, \ldots, N-1$. Stability, then, requires a cross-sectional mean-reversion condition in the sense that the growth rates of units for the agents with the most units in the population must be strictly below the growth rates of units for agents with smaller unit holdings. This condition is different from standard mean-reversion conditions that are imposed on a single stochastic process for stationarity. Indeed, our results in this paper do not require that the individual unit holdings processes for different agents x_i mean-revert or are stationary. Instead, our results require only that agents' unit holdings grow at different rates depending on their ranks in the distribution, regardless of the underlying stationarity properties of these individual processes. As we shall demonstrate in Section 4, this condition has significant implications for the dynamics of different ranked commodity prices.

The unstable case in which the cross-sectional mean-reversion condition of Theorem 2.3 does not hold is examined in detail by Fernholz and Fernholz (2014) and Fernholz (2016b). These authors show that this case generates unbounded concentration of units and divergence in the sense that asymptotically over time one agent comes to hold practically all the units in the economy (the time-averaged limit of $\theta_{(1)}$ is equal to one). This is an important insight that can only be observed using a discrete system of multiple stochastic differential equations. In the single-equation approach, this result is usually stated as the non-stationarity and infinite variance of an integrated process with no "frictions" present (Gibrat, 1931; Champernowne, 1953; Gabaix, 2009).[9] Of course, the statement that a single process is non-stationary and has an infinite variance is difficult to interpret meaningfully in empirical applications. In contrast, using our multiple-equation approach, we can derive the more general and empirically meaningful result that in the absence of cross-sectional mean-reversion as defined in Theorem 2.3, units will grow increasingly concentrated in the hands of one single agent over time.

[9] These "frictions" can take several forms including birth/death of processes, a reflecting barrier at some lower bound for unit values, or some positive shock that is regularly added to each agent's unit holdings over time. See Gabaix (2009) for an extensive discussion.

3 Estimation

One of the contributions of this paper is to provide techniques to estimate the shaping econometric factors α_k and σ_k using panel data. In this section, we describe how to accomplish this efficiently using the local time processes defined and discussed in Section 2.

In order to estimate the reversion rates and volatilities from equation (2.21) from Theorem 2.3, we use discrete-time approximations of the continuous processes that yield the theorem. For the estimation of the volatility parameters σ_k^2, we use the discrete-time approximation of equation (2.18) above. In particular, these estimators are given by

$$\hat{\sigma}_k^2 = \frac{1}{T-1}\sum_{t=1}^{T-1}\left[\big(\log\theta_{p_t(k)}(t+1) - \log\theta_{p_t(k+1)}(t+1)\big) - \big(\log\theta_{p_t(k)}(t) - \log\theta_{p_t(k+1)}(t)\big)\right]^2, \tag{3.1}$$

for all $k = 1, \ldots, N-1$. Note that T is the total number of periods covered in the data.

The estimation of the rank-based relative growth rates α_k is more difficult. In order to estimate these parameters, we first estimate the local time parameters κ_k and then exploit the relationship that exists between these local times and the rank-based relative growth rates.

Lemma 3.1. *The relative growth rate parameters α_k and the local time parameters κ_k satisfy*

$$\alpha_k = \frac{1}{2}\kappa_{k-1} - \frac{1}{2}\kappa_k, \tag{3.2}$$

for all $k = 1, \ldots, N-1$, and $\alpha_N = -(\alpha_1 + \cdots + \alpha_{N-1})$.

Lemma 3.2. *The ranked agent unit share processes $\theta_{(k)}$ satisfy the stochastic differential equation*

$$d\log\big(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\big) = d\log\big(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t)\big) - \frac{\theta_{(k)}(t)}{2(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t))}d\Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(t), \quad \text{a.s.}, \tag{3.3}$$

for all $k = 1, \ldots, N$.

These lemmas together allow us to generate estimates of the rank-based relative growth rates α_k. In order to accomplish this, we first estimate the local time processes $\Lambda_{\log\theta^*_{(k)} - \log\theta^*_{(k+1)}}$

using the discrete-time approximation of equation (3.3). This discrete-time approximation implies that for all $k = 1, \ldots, N$,

$$\log\left(\theta_{p_t(1)}(t+1) + \cdots + \theta_{p_t(k)}(t+1)\right) - \log\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right) \approx \log\left(\theta_{p_{t+1}(1)}(t+1) + \cdots + \theta_{p_{t+1}(k)}(t+1)\right) - \log\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right) - \frac{\theta_{p_t(k)}(t)}{2\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right)}\left(\Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(t+1) - \Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(t)\right), \tag{3.4}$$

which, after simplification and rearrangement, yields the local time estimators

$$\hat{\Lambda}_{\log\theta_{(k)} - \log\theta_{(k+1)}}(t+1) - \hat{\Lambda}_{\log\theta_{(k)} - \log\theta_{(k+1)}}(t) = \Big[\log\left(\theta_{p_{t+1}(1)}(t+1) + \cdots + \theta_{p_{t+1}(k)}(t+1)\right) - \log\left(\theta_{p_t(1)}(t+1) + \cdots + \theta_{p_t(k)}(t+1)\right)\Big]\frac{2\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right)}{\theta_{p_t(k)}(t)}. \tag{3.5}$$

As with our estimates of the volatility parameters σ_k^2, we estimate the values of the local times in equation (3.5) for $t = 1, \ldots, T-1$, where T is the total number of periods covered in the data. We also set $\hat{\Lambda}_{\log\theta_{(k)} - \log\theta_{(k+1)}}(0) = 0$, for all $k = 1, \ldots, N$.

Equation (3.5) states that the change in the local time estimator $\hat{\Lambda}_{\log\theta_{(k)} - \log\theta_{(k+1)}}$ is increasing in the difference between the time $t+1$ unit holdings of the largest k agents *at time* $t+1$ and the time $t+1$ unit holdings of the largest k agents *at time* t, a nonnegative number. Of course, this difference measures the intensity of cross-sectional mean reversion, since a large difference implies that the k agents with the largest unit holdings at time t have seen their units grow substantially slower than some other subset of agents that had smaller unit holdings at time t and are now themselves the agents with the largest unit holdings in the economy.

After estimating the local times in equation (3.5), we then use equation (2.19) to generate estimates of κ_k according to

$$\hat{\kappa}_k = \frac{1}{T}\hat{\Lambda}_{\log\theta_{(k)} - \log\theta_{(k+1)}}(T), \tag{3.6}$$

for all $k = 1, \ldots, N$. Finally, we can use the relationship between the parameters α_k and κ_k

established by Lemma 3.1 to define the estimator

$$\hat{\alpha}_k = \frac{1}{2}\hat{\kappa}_{k-1} - \frac{1}{2}\hat{\kappa}_k, \tag{3.7}$$

for all $k = 1, \ldots, N-1$. In Appendix A, we investigate the efficiency of this estimator.

While the methods described in this section explain how to generate point estimates of the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k, it is important to also understand how much variation there is in these estimates. It is not possible to generate confidence intervals using classical techniques in this setting because the empirical distribution of the parameters α_k and σ_k is unknown. However, it is possible to use bootstrap resampling to generate confidence intervals for these estimated factors.

Equations (3.1) and (3.5) show that the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k are measured as changes from one period, t, to the next, $t+1$. As a consequence, the bootstrap resamples we construct consist of $T-1$ pairs of observations of agents' unit holdings from adjacent time periods (periods t and $t+1$). Such resamples, of course, are equivalent to the full sample which has observations over T periods and hence consists of $T-1$ pairs of observations from adjacent periods. The confidence intervals are then generated by determining the range of values that obtain for the parameters α_k and σ_k over all of the bootstrap resamples. In Section 4, we apply our techniques to the distribution of relative commodity prices and generate confidence intervals for our estimates of the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k following this procedure.

4 Application: The Distribution of Commodity Prices

We wish to confirm the validity and accuracy of the empirical methods we presented in Section 2. We do this using a publicly available data set on the global monthly spot prices of 22 common commodities for 1980 - 2015 obtained from the Federal Reserve Bank of St. Louis (FRED).[10]

[10]These commodities are aluminum, bananas, barley, beef, Brent crude oil, cocoa, copper, corn, cotton, iron, lamb, lead, nickel, orange, poultry, rubber, soybeans, sugar, tin, wheat, wool (fine), and zinc.

In order to accomplish this, we shall use the results and procedure described in Section 3 to estimate rank-based reversion rates, $-\alpha_k$, and idiosyncratic volatilities, σ_k, for the distribution of relative commodity prices over our sample period 1980 - 2015. In this section, then, we shall interpret agents' holdings of units $x_i(t)$ from equation (2.1) as the prices of different commodities. Because commodities are sold in different units and hence their prices cannot be compared in an economically meaningful way, it is important to normalize these prices by equalizing them in the initial period.

As long as the distribution of relative commodity prices is approximately stationary, then Theorem 2.3 should accurately describe this distribution. This is true regardless of whether commodity price formation is modeled using a simple CES-utility approach or a richer framework that includes production and commodity-specific technology shocks as in Alquist and Coibion (2014). Indeed, any plausible model of commodity price formation is consistent with our nonparametric framework based on the general Itô processes of equation (2.1).

The econometric results of Sections 2 and 3 apply to the distribution of the parameters $\theta_{(k)}$, $k = 1, \ldots, N$, which in those sections represented the shares of total units held by different ranked agents. If we interpret the x_i as commodity prices, then the parameters $\theta_{(k)}$ represent commodity "price shares," a quantity that is well defined but difficult to interpret economically. It is easy to show, however, that the distribution of these commodity "price shares" $\theta_{(k)}$ is the same as the distribution of commodity prices relative to the average of all commodity prices. This latter quantity has a clear economic interpretation. In this section, we estimate reversion rates and idiosyncratic volatilities that describe the stationary distribution of relative normalized commodity prices according to equation (2.21). As Figure 1 demonstrates, the distribution of these relative normalized prices appears to be roughly stationary over time. Consistent with this observation, we confirm below that the methods presented in Sections 2 and 3 do in fact accurately describe this stationary distribution.

If we let $\bar{x}(t)$ equal the average price of all N commodities at time t, then for all $i = 1, \ldots, N$, the relative price of commodity i at time t is defined as

$$\tilde{x}_i(t) = \frac{x_i(t)}{\bar{x}(t)} = \frac{N x_i(t)}{x_1(t) + \cdots + x_N(t)}. \tag{4.1}$$

The relative price $\tilde{x}_i(t)$ is equal to the price of commodity i at time t relative to the average price of all N commodities at time t. If we let $\tilde{x}_{(k)}(t)$ denote the relative price of the k-th ranked commodity at time t, then equations (2.4) and (2.8) imply that, for all $i, k = 1, \dots, N$,

$$\tilde{x}_i(t) = N\theta_i(t) \qquad \text{and} \qquad \tilde{x}_{(k)}(t) = N\theta_{(k)}(t). \tag{4.2}$$

Note that the k-th ranked commodity at time t refers to the commodity with the k-th highest price at time t. It follows from equation (4.2) that for all $k = 1, \dots, N-1$ and all t,

$$\log \tilde{x}_{(k)}(t) - \log \tilde{x}_{(k+1)}(t) = \log \theta_{(k)}(t) - \log \theta_{(k+1)}(t), \tag{4.3}$$

and hence equation (2.21) from Theorem 2.3 describes both the distribution of commodity "price shares," $\theta_{(k)}$, and relative commodity prices, $\tilde{x}_{(k)}$. In other words, all of our previous results apply to the distribution of relative commodity prices as well.

4.1 Prediction and Data

The econometric results of Section 2 suggest that any stationary size distribution can be accurately characterized by the reversion rates, $-\alpha_k$, and idiosyncratic volatilities, σ_k, according to equation (2.21). Fernholz (2002) and Fernholz and Koch (2016) show, respectively, that this is in fact true for the size distributions of total market capitalizations of U.S. stocks and total assets of U.S. financial intermediaries. One of this paper's contributions is to further demonstrate the validity of our econometric techniques using a new data set.

The first step is to estimate the reversion rates $-\alpha_k$ for each rank $k = 1, \dots, N$. As described in Section 2, these reversion rates measure the growth rates of different ranked commodity prices relative to the growth rate of all commodity prices together. In Figure 2, we plot annualized values of minus the reversion rates α_k for each rank in the distribution of relative normalized commodity prices together with 95% confidence intervals based on the results of 10,000 bootstrap resample estimates.

The figure shows a substantial and statistically significant deviation from Gibrat's law for commodity prices. According to Figure 2, the lowest-ranked, lowest-priced spot prices on average grow 3-4% faster than all commodity prices together while the highest-ranked, highest-priced spot prices on average grow between 5 and 10% slower than the aggregate. This means that there is nearly a 15% difference between the average annual growth rates of low- and high-ranked commodity prices. These differential growth rates give rise to the cross-sectional mean reversion discussed and analyzed in Section 2, and are a shaping force of the stationary distribution of relative commodity prices according to equation (2.21) from Theorem 2.3. These differential growth rates also imply that there is predictability of future commodity prices using rank, a point we shall explore in more detail below.

The reversion rates $-\alpha_k$ are estimated using the procedure described in Section 3. In particular, these estimated reversion rates are generated by first estimating the local time parameters κ_k according to equations (3.5) and (3.6), and then generating estimates of the parameters α_k according to equation (3.7). Figure 3 plots the evolution of the local time processes $\Lambda_{\log \tilde{x}_{(k)} - \log \tilde{x}_{(k+1)}}$, $k = 1, \ldots, N-1$, which we use to construct our estimates of the reversion rates $-\alpha_k$.

The confidence intervals in Figure 2 are generated using this same procedure, only with bootstrap resamples instead of the original full sample. These confidence intervals show that the deviations from Gibrat's law for commodity prices during the 1980 - 2015 period are highly statistically significant. This observation confirms the usefulness of our rank-based methods, since these methods allow for growth rates that vary across the distribution of relative commodity prices in the realistic manner shown in Figure 2.

The next step is to estimate the idiosyncratic volatilities σ_k, which is accomplished using the discrete-time approximation given by equation (3.1). Figure 4 plots annualized estimates of these parameter values for each rank in the distribution of relative normalized commodity prices together with 95% confidence intervals based on the results of 10,000 bootstrap resample estimates. The estimates and confidence intervals for the parameters α_k and σ_k in Figures 2 and 4 are smoothed across different ranks using a Gaussian kernel smoother. Following Fernholz and Koch (2016), we smooth these parameters between 1 and 100 times and then choose the number of smoothings within this range that minimizes the squared deviation between the predicted relative commodity prices according to equation (2.21) and the average observed relative commodity prices for the period 1981-2015.[11]

How well do the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k reported in Figures 2 and 4 replicate the true distribution of relative commodity prices? Figure 5 shows that these estimated parameters generate predicted relative commodity prices according to equation (2.21) that do in fact match the average relative commodity prices observed during the

[11] The commodity prices are normalized to all equal each other at the start of our sample period in 1980. Since it takes a number of months for these initially equal relative prices to converge to a stationary distribution, we remove the first year of data when generating observed average relative prices to compare to predicted relative prices for the purposes of smoothing the parameters α_k and σ_k.

1980 - 2015 sample period. The squared deviation between predicted and observed average relative commodity prices over this sample period is 0.143. Thus, we further confirm the validity of our econometric methods using commodity prices data.

The straight line in the left part of Figure 5 implies that the distribution of relative normalized commodity prices approximately follows a power law at the highest commodity price ranks during the 1980 - 2015 time period. Specifically, this figure indicates that at the highest ranks, rank and price are approximately related by the linear relationship

$$\log(\text{price relative to average}) = -0.386 \log(\text{rank}) + 0.897. \tag{4.4}$$

It is important to note, however, that even though relative commodity prices are well-approximated by a power law at the highest ranks, the pattern in Figure 5 is also similar to a lognormal distribution.

4.2 Rank-Based Forecasts of Future Commodity Prices

One of the implications of Theorem 2.3 is that there is a stationary distribution of relative commodity prices if and only if $\alpha_1 + \cdots + \alpha_k < 0$. This necessary condition is a cross-sectional mean-reversion condition which states that the growth rates of the prices of the higher-priced, higher-ranked commodities must on average be lower than the growth rates of the prices of the lower-priced, lower-ranked commodities. This condition is different from standard notions of mean reversion necessary for the stationarity of a single process. Indeed, we make no assumptions about mean reversion or stationarity for the individual processes x_i, since such single-process stationarity is not necessary for any of the results in this paper. Instead, we require only the *cross-sectional* mean-reversion condition specified across ranks in Theorem 2.3.

This condition of Theorem 2.3 offers a testable prediction for our commodity prices data—there should be predictability of future commodity prices relative to each other based on rank. In particular, the cross-sectional mean-reversion condition requires that the monthly log growth rates of the top k ranked commodities be smaller on average than the monthly log growth rates of the bottom $N-k$ ranked commodities, for all rank cutoffs $k = 1, \ldots, N-1$. There are 21 different rank cutoffs to test with our data set consisting of 22 commodity prices from 1980 - 2015. Table 1 reports the average difference in monthly log growth rates between the top k and bottom $N-k$ commodities, for all 21 rank cutoffs k. The table also reports the standard deviation of this difference in log growth rates and the t-statistic for the null hypothesis that the difference in log growth rates is equal to zero.

Table 1 presents empirical evidence of statistically significant predictability of future commodity prices based on commodity rank. This is exactly what is predicted by the cross-sectional mean-reversion condition of Theorem 2.3. The table shows that for all 21 rank cutoffs the null hypothesis can be rejected at 5% significance, and for most rank cutoffs the null hypothesis can be rejected at 0.1% significance. In all cases, the log growth rate of the top k ranked commodity prices is below the log growth rate of the bottom $N-k$ ranked commodity prices, just as our econometric theory in Section 2 requires.

We also wish to investigate the power of our rank-based econometric theory for forecasting future commodity prices out of sample. The cross-sectional mean-reversion condition of Theorem 2.3 and the estimated reversion rates $-\alpha_k$ shown in Figure 2 suggest that rank may have some power in forecasting future commodity prices, although any such forecasting power is limited by the high estimates for the volatility parameters σ_k reported in Figure 4. We generate forecasts using equation (3.3) from Lemma 3.2 in Section 3. In particular, if we write equation (3.3) in terms of relative prices $\tilde{x}$ rather than unit shares θ, then we have

$$E_t\left[d\log\left(\tilde{x}_{p_t(1)}(t)+\cdots+\tilde{x}_{p_t(k)}(t)\right)\right] = -\frac{\tilde{x}_{(k)}(t)}{2(\tilde{x}_{(1)}(t)+\cdots+\tilde{x}_{(k)}(t))}E_t\left[d\Lambda_{\log\tilde{x}_{(k)}-\log\tilde{x}_{(k+1)}}(t)\right], \tag{4.5}$$

for all $k=1,\ldots,N$. Note that the expected value of $d\log\left(\tilde{x}_{(1)}(t)+\cdots+\tilde{x}_{(k)}(t)\right)$ in equation (3.3) is zero since a stationary distribution of relative commodity prices implies a zero expected change in the ranked log relative prices $\tilde{x}_{(k)}$.

According to equation (4.5), we can forecast the change in relative log prices of each subset of top k ranked commodities using our estimates of the change in the local time processes $\Lambda_{\log\tilde{x}_{(k)}-\log\tilde{x}_{(k+1)}}$. Our estimates of the change in the local time processes are denoted by $\hat{\kappa}_k$, and are defined in equation (3.6) in Section 3. Thus, for each $k=1,\ldots,N-1$, we can use our estimates $\hat{\kappa}_k$ to forecast relative log commodity prices one month ahead for the top k ranked commodities in that month. In order for these forecasts to be meaningful, they must be outside of the sample period in which the local time parameters κ_k are estimated.

We adopt the following two out-of-sample forecast procedures. First, we use the first ten years of our 1980 - 2015 sample period to generate estimates $\hat{\kappa}_k$ as described in Section 3, and then we use those fixed estimates together with equation (4.5) to forecast relative log commodity prices one month ahead over the remaining 1990 - 2015 time period. Second, we use rolling estimates $\hat{\kappa}_k$ that are updated each month to forecast relative log commodity prices one month ahead over the 1990 - 2015 time period. For this procedure, each month's forecast uses estimates $\hat{\kappa}_k$ that are generated using commodity price data from all months prior to that forecast month. In other words, the first forecast in January 1990 uses the same ten-year sample period as in the first procedure described above to generate the estimates $\hat{\kappa}_k$, but each subsequent month's forecast is generated using updated estimates $\hat{\kappa}_k$.

Table 2 reports the root mean squared error (RMSE) of one-month out-of-sample forecasts using the two procedures described in the previous paragraph relative to the RMSE of a random-walk forecast that uses the current month's relative commodity prices to forecast the next month's relative commodity prices. Values of this RMSE ratio below one indicate that our procedures generated more accurate predictions than a random-walk forecast. These ratios are reported for all rank cutoffs $k = 1, \ldots, N-1$. For all 21 rank cutoffs, the RMSE ratio is below one and hence the out-of-sample forecasts generated using equation (4.5) and estimated local times $\hat{\kappa}_k$ in all cases outperform a random-walk forecast using either procedure. In some cases, especially at the highest rank cutoffs, our local-time-based forecasts substantially outperform the random-walk forecasts.

The results reported in Table 2 are notable because these forecasts do not use any information about supply, demand, economic conditions, or any commodity-specific fundamentals. Instead, the only information the forecasts rely on is commodity price rank, which is used in conjunction with the econometric theory developed in this paper. Nonetheless, these simple forecasts are able to improve on random-walk forecasts for every subset of top k ranked commodities over the 1990 - 2015 forecast period. Our forecast exercise is different from other commodity price forecast exercises such as those of Chinn and Coibion (2014) and West and Wong (2014) since we generate forecasts of sums of multiple top-ranked commodity prices relative to all commodity prices rather than forecasts of individual real commodity prices. Nonetheless, the success of this simple local-time-based forecast methodology is notable. Note, for example, that most of the RMSE ratios reported in Table 2 are lower than the RMSE ratios reported by West and Wong (2014), who use a factor model to forecast individual one-month-ahead commodity prices.

It is beyond the scope of this paper to compare our local-time-based forecasts in detail to other more common methods of forecasting future commodity prices. Nonetheless, there is no reason why the local-time-based forecasting methodology presented in this section cannot be combined with more common commodity forecasting methods to generate even better forecasts. In fact, because commodity price rank probably does not correlate strongly with fundamentals or forecasting factors, it is likely that forecasts using both the rank-based methods presented in this paper and more traditional commodity forecasting methods such as those of Alquist and Coibion (2014) or West and Wong (2014) can generate some of the best forecasts in the literature so far.

5 Conclusion

This paper presents rank-based, nonparametric methods that allow for the characterization of general power law distributions in random growth settings. We show that any stationary distribution in a random growth setting is shaped entirely by two factors—the idiosyncratic volatilities and reversion rates (a measure of cross-sectional mean reversion) for different ranks in the distribution. An increase in idiosyncratic volatilities increases concentration, while an increase in reversion rates decreases concentration. We also provide methods for estimating these two shaping factors using panel data.

Using data on a set of 22 global commodity prices from 1980 - 2015, we show that our rank-based, nonparametric methods accurately describe the distribution of relative normalized commodity prices. According to our econometric results, a necessary condition for the existence of a stationary distribution is that higher-ranked commodity prices must grow more slowly than lower-ranked commodity prices. In other words, our results predict that commodity price rank will forecast future commodity prices. We confirm this prediction and show that out-of-sample rank-based forecasts derived from our econometric results are more accurate than random-walk forecasts of one-month-ahead commodity prices during the 1990 - 2015 period.

A The Efficiency of Local-Time-Based Estimation

We wish to investigate the efficiency of the local-time-based estimator $\hat{\alpha}_k$ defined in equation (3.7) in Section 3. In order to accomplish this, we consider a simple rank-based econometric model that yields Zipf's law for the distribution of agents' unit shares $\theta_{(k)}$.

Let $r_t(i)$ denote the rank of agent i in terms of unit holdings at time t, so that $r_t(i) = k$ if and only if $x_i(t) = x_{(k)}(t)$. We assume that Gibrat's law holds in the two econometric models we consider in this appendix. In particular, we assume that the units held by each agent i follow the stochastic differential equation

$$d\log x_i = (g_{r_t(i)} + c(t))\,dt + \gamma\,dB_i(t), \tag{A.1}$$

where $g_k = g < 0$ for all $k > N$, $g_N = -(N-1)g$, $c(t)$ is a common growth rate across all agents, and $\gamma > 0$ is the common standard deviation of idiosyncratic shocks to agents' unit holdings. For this econometric model, the reversion rates $-\alpha_k$ defined in equation (2.17) are equal to $-g > 0$, for all $k = 1, \ldots, N-1$. Similarly, the idiosyncratic volatilities σ_k defined in equation (2.18) are equal to $2\gamma^2$, for all $k = 1, \ldots, N-1$. Thus, the econometric model (A.1) follows Gibrat's law, with the high growth rate of the agent at rank N corresponding to the "friction" that is needed for a stationary distribution to exist (Champernowne, 1953; Gabaix, 2009).

If the parameters from equation (A.1) satisfy $\gamma^2 = -2g$, then this econometric model generates a stationary power law distribution of unit shares that follows Zipf's law. This follows because the parameters α and σ^2 from equation (2.25) are equal to g and $2\gamma^2$ from equation (A.1), respectively.[12] In Section 2.4, we showed that Zipf's law follows from Gibrat's law if $\sigma^2 = -4\alpha$. It follows, then, that Zipf's law obtains in the econometric model (A.1) if $2\gamma^2 = -4g$.

We consider two parameterizations of this econometric model. In the first, we set $g = -0.02$ and $\gamma = 0.2$. In the second, we set $g = -0.2$ and $\gamma = 0.632 = \sqrt{0.4}$. Because $\gamma^2 = -2g$ in both cases, these two parameterizations generate stationary distributions that follow Zipf's law. We also set the total number of agents N equal to 100 for both parameterizations.

In addition to the local-time-based estimation of the reversion rates $-\alpha_k$ described in Section 3, it is also possible to directly estimate these parameters using equation (2.17). In particular, we define the direct estimator of the parameters α_k, for all $k = 1, \ldots, N$, using a discrete-time approximation of equation (2.17):

$$\hat{\alpha}_k^D = \frac{1}{T-1}\sum_{t=1}^{T-1}\left[\log\theta_{p_t(k)}(t+1) - \log\theta_{p_t(k)}(t)\right]. \tag{A.2}$$

[12]The econometric model (A.1) satisfies the stability condition of Theorem 2.3 because $\alpha_1 + \cdots + \alpha_k = kg < 0$, for all $k = 1, \ldots, N-1$.

We investigate the efficiency of local-time-based estimation of the reversion rates by comparing the asymptotic properties of the estimator $\hat{\alpha}_k$ defined in equation (3.7) relative to the direct estimator $\hat{\alpha}_k^D$ defined in equation (A.2).

We ran 1000 simulations of the two parameterizations of the econometric model (A.1) over a period of 50 years. The data are assumed to be observable at a monthly frequency, as in the case of the commodity price data examined in Section 4. For both parameterizations, we produce local-time-based estimates and direct estimates of the parameters α_k using samples that range from 1 to 50 years of monthly data. The results are shown in Table 3, which reports the total root mean squared error (RMSE) of the estimated parameters relative to the true parameters summed over all 100 ranks.[13] These total RMSEs are averages over all 1000 simulations, and are reported for different sample lengths.

According to Table 3, the local-time-based estimator $\hat{\alpha}_k$ generates smaller total RMSEs than the direct estimator $\hat{\alpha}_k^D$ at all sample lengths for both parameterizations of the econometric model (A.1). The efficiency advantage of the local-time-based estimator is substantial, and is larger in the first parameterization than in the second. This first result highlights the desirable small-sample properties of the local-time-based estimator $\hat{\alpha}_k$ presented in this paper. This second result suggests that more volatile systems with higher values of the parameters α_k and σ_k may yield smaller efficiency gains for the local-time-based estimator relative to the direct estimator $\hat{\alpha}_k^D$. The performance of the local-time-based estimator $\hat{\alpha}_k$ in the first, less-volatile parameterization is worth highlighting, as the estimated parameters converge rapidly to the true parameters after just a few years of observations.

B Assumptions and Regularity Conditions

In this appendix, we present the assumptions and regularity conditions that are necessary for the stable distribution characterization in Theorem 2.3. As discussed in Section 2, these assumptions admit a large class of continuous unit processes for the agents in our setup. The first assumption establishes basic integrability conditions that are common for both continuous semimartingales and Itô processes.

Assumption B.1. *For all $i = 1, \ldots, N$, the growth rate processes μ_i satisfy*

$$\int_0^T |\mu_i(t)|\,dt < \infty, \quad T > 0, \quad \text{a.s.,} \tag{B.1}$$

[13]Because there are $N = 100$ agents in the econometric model, there are 100 parameters α_k to estimate.

and the volatility processes $\boldsymbol{\delta}_i$ *satisfy*

$$\int_0^T \left\|\boldsymbol{\delta}_i(t)\right\|^2 dt < \infty, \quad T > 0, \quad \text{a.s.}, \tag{B.2}$$

$$\left\|\boldsymbol{\delta}_i(t)\right\|^2 > 0, \quad t > 0, \quad \text{a.s.} \tag{B.3}$$

$$\lim_{t\to\infty} \frac{1}{t}\left\|\boldsymbol{\delta}_i(t)\right\|^2 \log\log t = 0, \quad \text{a.s.}, \tag{B.4}$$

Conditions (B.1) and (B.2) are standard in the definition of an Itô process, while condition (B.3) ensures that agents' holdings of units contain a nonzero random component at all times. Condition (B.4) is similar to a boundedness condition in that it ensures that the variance of agents' unit holdings does not diverge to infinity too rapidly.

The second assumption underlying our results establishes that no two agents' unit holdings be perfectly correlated over time. In other words, there must always be some idiosyncratic component to each agent's unit dynamics. Finally, we also assume that no agent's unit holdings relative to the total units for all agents shall disappear too rapidly.

Assumption B.2. *The symmetric matrix* $\rho(t)$*, given by* $\rho(t) = (\rho_{ij}(t))$*, where* $1 \leq i, j \leq N$*, is nonsingular for all* $t > 0$*, a.s.*

Assumption B.3. *For all* $i = 1, \dots, N$*, the unit share processes* θ_i *satisfy*

$$\lim_{t\to\infty} \frac{1}{t} \log \theta_i(t) = 0, \quad \text{a.s.} \tag{B.5}$$

C Proofs

This appendix presents the proofs of Lemmas 2.1, 2.2 3.1, and 3.2, and Theorem 2.3.

Proof of Lemma 2.1. By definition, $x(t) = x_1(t) + \cdots + x_N(t)$ and for all $i = 1, \dots, N$, $\theta_i(t) = x_i(t)/x(t)$. This implies that

$$dx(t) = \sum_{i=1}^N dx_i(t) = \sum_{i=1}^N \theta_i(t)x(t)\frac{dx_i(t)}{x_i(t)},$$

from which it follows that

$$\frac{dx(t)}{x(t)} = \sum_{i=1}^N \theta_i(t)\frac{dx_i(t)}{x_i(t)}. \tag{C.1}$$

We wish to show that the process satisfying equation (2.3) also satisfies equation (C.1).

If we apply Itô's Lemma to the exponential function, then equation (2.3) yields

$$\begin{aligned} dx(t) = x(t)\mu(t)\,dt &+ \frac{1}{2}x(t)\sum_{i,j=1}^{N}\theta_i(t)\theta_j(t)\boldsymbol{\delta}_i(t)\cdot\boldsymbol{\delta}_j(t)\,dt \\ &+ x(t)\sum_{i=1}^{N}\theta_i(t)\boldsymbol{\delta}_i(t)\cdot d\mathbf{B}(t), \end{aligned} \tag{C.2}$$

a.s., where $\mu(t)$ is given by equation (2.5). Using the definition of $\rho_{ij}(t)$ from equation (2.2), we can simplify equation (C.1) and write

$$\frac{dx(t)}{x(t)} = \left(\mu(t) + \frac{1}{2}\sum_{i,j=1}^{N}\theta_i(t)\theta_j(t)\rho_{ij}(t)\right)dt + \sum_{i=1}^{N}\theta_i(t)\boldsymbol{\delta}_i(t)\cdot d\mathbf{B}(t). \tag{C.3}$$

Similarly, the definition of $\mu(t)$ from equation (2.5) allows us to further simplify equation (C.3) and write

$$\begin{aligned} \frac{dx(t)}{x(t)} &= \left(\sum_{i=1}^{N}\theta_i(t)\mu_i(t) + \frac{1}{2}\sum_{i=1}^{N}\theta_i(t)\rho_{ii}(t)\right)dt + \sum_{i=1}^{N}\theta_i(t)\boldsymbol{\delta}_i(t)\cdot d\mathbf{B}(t) \\ &= \sum_{i=1}^{N}\theta_i(t)\left(\mu_i(t) + \frac{1}{2}\rho_{ii}(t)\right)dt + \sum_{i=1}^{N}\theta_i(t)\boldsymbol{\delta}_i(t)\cdot d\mathbf{B}(t). \end{aligned} \tag{C.4}$$

If we again apply Itô's Lemma to the exponential function, then equation (2.1) yields, a.s., for all $i = 1, \ldots, N$,

$$\begin{aligned} dx_i(t) &= x_i(t)\left(\mu_i(t) + \frac{1}{2}\big\|\boldsymbol{\delta}_i(t)\big\|^2\right)dt + x_i(t)\boldsymbol{\delta}_i(t)\cdot d\mathbf{B}(t) \\ &= x_i(t)\left(\mu_i(t) + \frac{1}{2}\rho_{ii}(t)\right)dt + x_i(t)\boldsymbol{\delta}_i(t)\cdot d\mathbf{B}(t). \end{aligned} \tag{C.5}$$

Substituting equation (C.5) into equation (C.4) then yields

$$\frac{dx(t)}{x(t)} = \sum_{i=1}^{N}\theta_i(t)\frac{dx_i(t)}{x_i(t)},$$

which completes the proof. □

Proof of Lemma 2.2. Agents' unit holding processes x_i are absolutely continuous in the sense that the random signed measures $\mu_i(t)\,dt$ and $\rho_{ii}(t)\,dt$ are absolutely continuous with respect to Lebesgue measure. As a consequence, we can apply Lemma 4.1.7 and Proposition 4.1.11 from Fernholz (2002), which yields equations (2.12) and (2.13). □

Proof of Lemma 3.1. This relationship between the rank-based relative growth rate parameters α_k and the local time parameters κ_k is established in the proof of Theorem 2.3 below (see equation (C.9) below). That proof also establishes the fact that $\alpha_N = -(\alpha_1 + \cdots + \alpha_{N-1})$ (see equation (C.11) below). □

Proof of Lemma 3.2. Consider the function $f_k(\theta_1, \ldots, \theta_N) = \theta_{(1)} + \cdots + \theta_{(k)}$, where $1 \leq k \leq N$. This function satisfies

$$\frac{\partial f_k}{\partial \theta_l} = 1,$$

for all $l = 1, \ldots, k$, and

$$\frac{\partial f_k}{\partial \theta_l} = 0,$$

for all $l = k+1, \ldots, N$. Furthermore, the support of the local time processes $\Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}$ is the set $\{t \ : \ \theta_{(k)}(t) = \theta_{(k+1)}(t)\}$, for all $k = 1, \ldots, N-1$. According to Theorem 4.2.1 and equations (3.1.1)-(3.1.2) of Fernholz (2002), then, the function $f_k(\theta_1, \ldots, \theta_N) = \theta_{(1)} + \cdots + \theta_{(k)}$ satisfies the stochastic differential equation

$$\begin{aligned} d\log(x_{p_t(1)}(t) + \cdots + x_{p_t(k)}(t)) - d\log x(t) &= d\log f_k(\theta_1(t), \ldots, \theta_N(t)) \\ &\quad - \frac{\theta_{(k)}(t)}{2(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t))}\, d\Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}, \quad \text{a.s.,} \end{aligned} \tag{C.6}$$

for all $k = 1, \ldots, N$.[14] Equation (C.6) is equivalent to

$$\begin{aligned} d\log\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right) &= d\log\left(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t)\right) \\ &\quad - \frac{\theta_{(k)}(t)}{2(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t))} d\Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(t), \end{aligned}$$

which confirms equation (3.3) from Lemma 3.2. □

Proof of Theorem 2.3. This proof follows arguments from Chapter 5 of Fernholz (2002).

[14]Equation (C.6) relies on the fact that $\log(x_{p_t(1)}(t) + \cdots + x_{p_t(k)}(t))$ is the value over time of a "portfolio" of unit holdings with weights of $\frac{\theta_{(l)}(t)}{\theta_{(1)} + \cdots + \theta_{(k)}}$ placed on each ranked unit holding $l = 1, \ldots, k$ and weights of zero placed on each ranked unit holding $l = k+1, \ldots, N$.

According to equation (2.15), for all $k = 1, \dots, N$,

$$\log \theta_{(k)}(T) = \int_0^T \left(\mu_{p_t(k)}(t) - \mu(t)\right)\, dt + \frac{1}{2}\Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(T) - \frac{1}{2}\Lambda_{\log\theta_{(k-1)} - \log\theta_{(k)}}(T) + \int_0^T \boldsymbol{\delta}_{p_t(k)}(t) \cdot d\mathbf{B}(t) - \sum_{i=1}^N \int_0^T \theta_i(t)\boldsymbol{\delta}_i(t) \cdot d\mathbf{B}(t). \tag{C.7}$$

Consider the asymptotic behavior of the process $\log \theta_{(k)}$. Assuming that the limits from equation (2.19) exist, then according to the definition of α_k from equation (2.17), the asymptotic behavior of $\log \theta_{(k)}$ satisfies

$$\lim_{T\to\infty} \frac{1}{T} \log \theta_{(k)}(T) = \alpha_k + \frac{1}{2}\kappa_k - \frac{1}{2}\kappa_{k-1} + \lim_{T\to\infty} \frac{1}{T} \int_0^T \boldsymbol{\delta}_{p_t(k)}(t) \cdot d\mathbf{B}(t) - \lim_{T\to\infty} \frac{1}{T} \sum_{i=1}^N \int_0^T \theta_i(t)\boldsymbol{\delta}_i(t) \cdot d\mathbf{B}(t), \quad \text{a.s.} \tag{C.8}$$

Assumption B.3 ensures that the term on the left-hand side of equation (C.8) is equal to zero, while Assumption B.1 ensures that the last two terms of the right-hand side of this equation are equal to zero as well (see Lemma 1.3.2 from Fernholz, 2002). If we simplify equation (C.8), then, we have that

$$\alpha_k = \frac{1}{2}\kappa_{k-1} - \frac{1}{2}\kappa_k, \tag{C.9}$$

which implies that

$$\alpha_k - \alpha_{k+1} = \frac{1}{2}\kappa_{k-1} - \kappa_k + \frac{1}{2}\kappa_{k+1}, \tag{C.10}$$

for all $k = 1, \dots, N-1$. Since equation (C.9) is valid for all $k = 1, \dots, N$, this establishes a system of equations that we can solve for κ_k. Doing this yields the equality

$$\kappa_k = -2(\alpha_1 + \cdots + \alpha_k), \tag{C.11}$$

for all $k = 1, \dots, N$. Note that asymptotic stability ensures that $\alpha_1 + \cdots + \alpha_k < 0$ for all $k = 1, \dots, N$, while the fact that $\alpha_N = \frac{1}{2}\kappa_{N-1} = -(\alpha_1 + \cdots + \alpha_{N-1})$ ensures that $\alpha_1 + \cdots + \alpha_N = 0$. Furthermore, if $\alpha_1 + \cdots + \alpha_k > 0$ for some $1 \le k < N$, then equation (C.11) generates a contradiction since $\kappa_k \ge 0$ by definition. In this case, it must be that Assumption B.3 is violated and $\lim_{T\to\infty} \frac{1}{T} \log \theta_{(k)}(T) \neq 0$ for some $1 \le k \le N$.

The last term on the right-hand side of equation (2.16) is an absolutely continuous martingale, and hence can be represented as a stochastic integral with respect to Brownian motion

$B(t)$.[15] This fact, together with equation (3.2) and the definitions of α_k and σ_k from equations (2.17)-(2.18), motivates our use of the stable version of the process $\log\theta_{(k)} - \log\theta_{(k+1)}$. Recall that, by equation (2.20), this stable version is given by

$$d\left(\log\theta^*_{(k)}(t) - \log\theta^*_{(k+1)}(t)\right) = -\kappa_k\,dt + d\Lambda_{\log\theta^*_{(k)} - \log\theta^*_{(k+1)}}(t) + \sigma_k\,dB(t), \tag{C.12}$$

for all $k = 1, \ldots, N-1$. According to Fernholz (2002), Lemma 5.2.1, for all $k = 1, \ldots, N-1$, the time-averaged limit of this stable version satisfies

$$\lim_{T\to\infty}\frac{1}{T}\int_0^T \left(\log\theta^*_{(k)}(t) - \log\theta^*_{(k+1)}(t)\right)\,dt = \frac{\sigma_k^2}{2\kappa_k} = \frac{\sigma_k^2}{-4(\alpha_1 + \cdots + \alpha_k)}, \tag{C.13}$$

a.s., where the last equality follows from equation (C.11).

As shown by Banner et al. (2005), the processes $\log\theta^*_{(k)} - \log\theta^*_{(k+1)}$ are stationary if the condition $\alpha_1 + \cdots + \alpha_k < 0$ holds, for all $k = 1, \ldots, N$. Thus, by ergodicity, equation (2.21) follows from equation (C.13). To the extent that the stable version of $\log\theta_{(k)} - \log\theta_{(k+1)}$ from equation (C.12) approximates the true version of this process from equation (2.16), then, the expected value of the true process $\log\theta_{(k)} - \log\theta_{(k+1)}$ will be approximated by $-\sigma_k^2/4(\alpha_1 + \cdots + \alpha_k)$, for all $k = 1, \ldots, N-1$. $\square$

[15]This is a standard result for continuous-time stochastic processes (Karatzas and Shreve, 1991; Nielsen, 1999).

References

Alquist, R. and O. Coibion (2014, March). Commodity-price comovement and global economic activity. *NBER Working Paper 20003*.

Altonji, J. G., A. A. Smith Jr., and I. Vidangos (2013, July). Modeling earnings dynamics. *Econometrica 81*(4), 1395–1454.

Atkinson, A. B., T. Piketty, and E. Saez (2011, March). Top incomes in the long run of history. *Journal of Economic Literature 49*(1), 3–71.

Axtell, R. (2001, September). Zipf distribution of u.s. firm sizes. *Science 293*(5536), 1818–1820.

Banner, A., R. Fernholz, and I. Karatzas (2005). Atlas models of equity markets. *Annals of Applied Probability 15*(4), 2296–2330.

Benhabib, J., A. Bisin, and S. Zhu (2011, January). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica 79*(1), 123–157.

Champernowne, D. G. (1953, June). A model of income distribution. *Economic Journal 63*(250), 318–351.

Chinn, M. D. and O. Coibion (2014, July). The predictive content of commodity futures. *Journal of Futures Markets 34*(7), 607–636.

Eeckhout, J. (2004). Gibrat's law for (all) cities. *American Economic Review 94*(5), 1429–1451.

Fernholz, E. R. (2002). *Stochastic Portfolio Theory*. New York, NY: Springer-Verlag.

Fernholz, R. T. (2016a, December). A model of economic mobility and the distribution of wealth. *Journal of Macroeconomics 50*, 168–192.

Fernholz, R. T. (2016b, January). A statistical model of inequality. *arXiv:1601.04093v1 [q-fin.EC]*.

Fernholz, R. T. and R. Fernholz (2014, July). Instability and concentration in the distribution of wealth. *Journal of Economic Dynamics and Control 44*, 251–269.

Fernholz, R. T. and C. Koch (2016, February). Why are big banks getting bigger? *Federal Reserve Bank of Dallas Working Paper 1604*.

Gabaix, X. (1999, August). Zipf's law for cities: An explanation. *Quarterly Journal of Economics 114*(3), 739–767.

Gabaix, X. (2009, 05). Power laws in economics and finance. *Annual Review of Economics 1*(1), 255–294.

Gabaix, X., J.-M. Lasry, P.-L. Lions, and B. Moll (2016, November). The dynamics of inequality. *Econometrica 84*(6), 2071–2111.

Gibrat, R. (1931). *Les inégalités économiques*. Librairie du Recueil Sirey, Paris.

Guvenen, F. (2009, January). An empirical investigation of labor income processes. *Review of Economic Dynamics 12*(1), 58–79.

Ichiba, T., V. Papathanakos, A. Banner, I. Karatzas, and R. Fernholz (2011). Hybrid atlas models. *Annals of Applied Probability 21*(2), 609–644.

Ioannides, Y. and S. Skouras (2013, January). Us city size distribution: Robustly pareto, but only in the tail. *Journal of Urban Economics 73*(1), 18–29.

Janicki, H. and E. S. Prescott (2006). Changes in the size distribution of us banks: 1960-2005. *FRB Richmond Economic Quarterly 92*(4), 291–316.

Karatzas, I. and S. E. Shreve (1991). *Brownian Motion and Stochastic Calculus*. New York, NY: Springer-Verlag.

Luttmer, E. G. J. (2007, August). Selection, growth, and the size distribution of firms. *Quarterly Journal of Economics 122*(3), 1103–1144.

Neumark, D., B. Wall, and J. Zhang (2011, February). Do small businesses create more jobs? new evidence for the united states from the national establishment time series. *Review of Economics and Statistics 93*(1), 16–29.

Newman, M. E. J. (2005, September-October). Power laws, pareto distributions, and zipf's law. *Contemporary Physics 46*(5), 323–351.

Nielsen, L. T. (1999). *Pricing and Hedging of Derivative Securities*. New York, NY: Oxford University Press.

Pal, S. and J. Pitman (2008). One-dimensional brownian particle systems with rank-dependent drifts. *Annals of Applied Probability 18*(6), 2179–2207.

Piketty, T. (2014). *Capital in the Twenty-First Century.* Cambridge, MA: Harvard University Press.

Saez, E. and G. Zucman (2016, May). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *Quarterly Journal of Economics 131*(2), 519–578.

Shkolnikov, M. (2011). Competing particle systems evolving by interacting lévy processes. *Annals of Applied Probability 21*(5), 1911–1932.

Van Nieuwerburgh, S. and P.-O. Weill (2010, October). Why has house price dispersion gone up? *Review of Economic Studies 77*(4), 1567–1606.

West, K. D. and K.-F. Wong (2014, April). A factor model for co-movements of commodity prices. *Journal of International Money and Finance 42*, 289–309.

Rank Cutoff k	Average Difference	Standard Deviation	t-Statistic
1	-1.46%	9.34%	-3.24
2	-1.69%	7.06%	-4.96
3	-1.42%	5.89%	-4.98
4	-1.49%	4.98%	-6.19
5	-1.28%	4.57%	-5.80
6	-1.14%	4.20%	-5.62
7	-1.05%	3.93%	-5.51
8	-0.87%	3.72%	-4.83
9	-0.75%	3.45%	-4.47
10	-0.68%	3.45%	-4.07
11	-0.64%	3.40%	-3.89
12	-0.61%	3.37%	-3.73
13	-0.62%	3.29%	-3.90
14	-0.59%	3.37%	-3.59
15	-0.52%	3.41%	-3.17
16	-0.47%	3.53%	-2.77
17	-0.45%	3.82%	-2.41
18	-0.36%	3.80%	-1.98
19	-0.47%	4.22%	-2.32
20	-0.53%	4.87%	-2.23
21	-0.88%	7.16%	-2.55

Table 1: Averages and standard deviations of the difference between monthly log growth rates for top k ranked commodities minus bottom $N - k$ ranked commodities from 1980 - 2015, for different values of the rank cutoff k. The t-statistics are for the null hypothesis that the difference in log growth rates between high-ranked and low-ranked commodities is equal to zero.

Rank Cutoff k	Fixed Estimation of κ_k	Rolling Estimation of κ_k
1	0.988	0.988
2	0.972	0.973
3	0.971	0.972
4	0.958	0.959
5	0.963	0.963
6	0.964	0.964
7	0.966	0.965
8	0.982	0.978
9	0.993	0.989
10	0.990	0.986
11	0.983	0.980
12	0.986	0.983
13	0.981	0.978
14	0.989	0.987
15	0.990	0.988
16	0.987	0.985
17	0.990	0.988
18	0.996	0.992
19	0.994	0.990
20	0.994	0.991
21	0.990	0.987

Table 2: Root mean squared error (RMSE) ratios of one-month-ahead out-of-sample forecasts of log price of top k ranked commodities relative to the price of all N commodities from 1990 - 2015, for different values of the rank cutoff k. These error ratios report RMSEs for forecasts using equation (4.5) relative to RMSEs for random-walk forecasts. The fixed estimation column corresponds to one-month-ahead forecasts from 1990-2015 using estimates of the local time parameters κ_k generated from the price data before 1990 only, while the rolling estimation column corresponds to one-month-ahead forecasts from 1990-2015 using updated rolling estimates of the local time parameters κ_k from all months of price data before the forecast month.

Parameterization: $g = -0.02$, $\gamma = 0.2$		
Sample Length (Years)	Local-Time-Based Estimator	Direct Estimator
1	0.19%	4.75%
2	0.12%	3.22%
3	0.10%	2.55%
4	0.09%	2.27%
5	0.09%	2.01%
10	0.08%	1.42%
20	0.08%	0.97%
30	0.08%	0.80%
40	0.08%	0.69%
50	0.08%	0.62%

Parameterization: $g = -0.2$, $\gamma = 0.632$		
Sample Length (Years)	Local-Time-Based Estimator	Direct Estimator
1	1.77%	14.67%
2	1.19%	10.66%
3	1.04%	8.02%
4	0.97%	6.98%
5	0.92%	6.16%
10	0.79%	3.93%
20	0.66%	2.51%
30	0.57%	1.79%
40	0.51%	1.33%
50	0.48%	1.03%

Table 3: Total root mean squared errors (RMSEs) of direct and local-time-based estimates of parameters α_k relative to true values for two parameterizations of the econometric model (A.1) and different sample lengths. The reported total RMSEs are sums over all 100 estimated parameters α_k, and are averaged across 1000 simulations. The local-time-based estimator is defined in equation (3.7) and the direct estimator is defined in equation (A.2). The top panel presents total RMSEs for a parameterization with $g = -0.02$ and $\gamma = 0.2$ and the bottom panel presents total RMSEs for a parameterization with $g = -0.2$ and $\gamma = 0.632$.

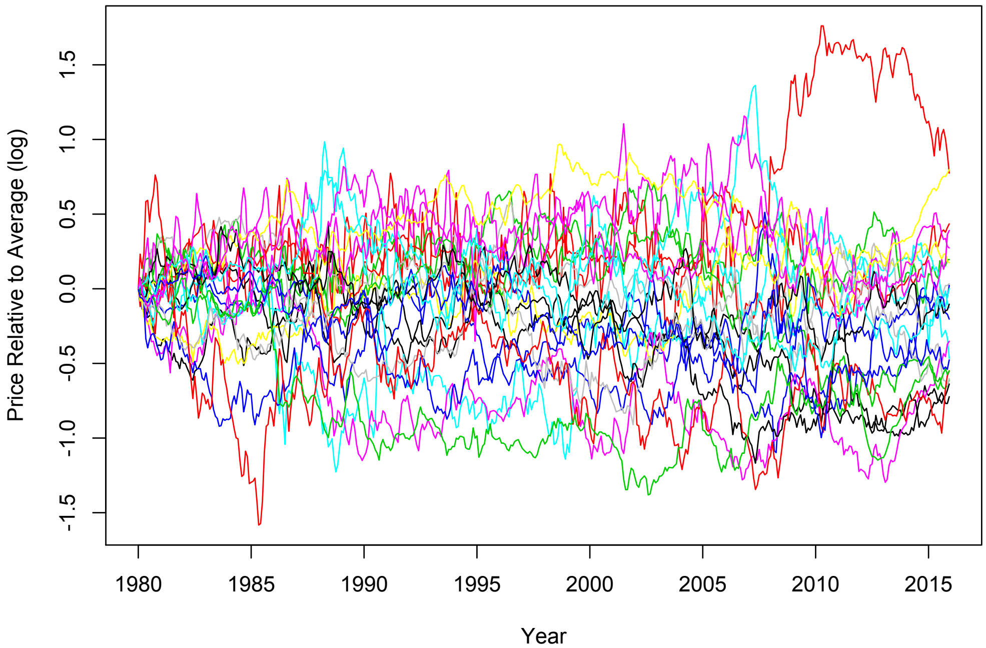

Figure 1: Log prices of commodities relative to the average price of all commodities, 1980 - 2015.

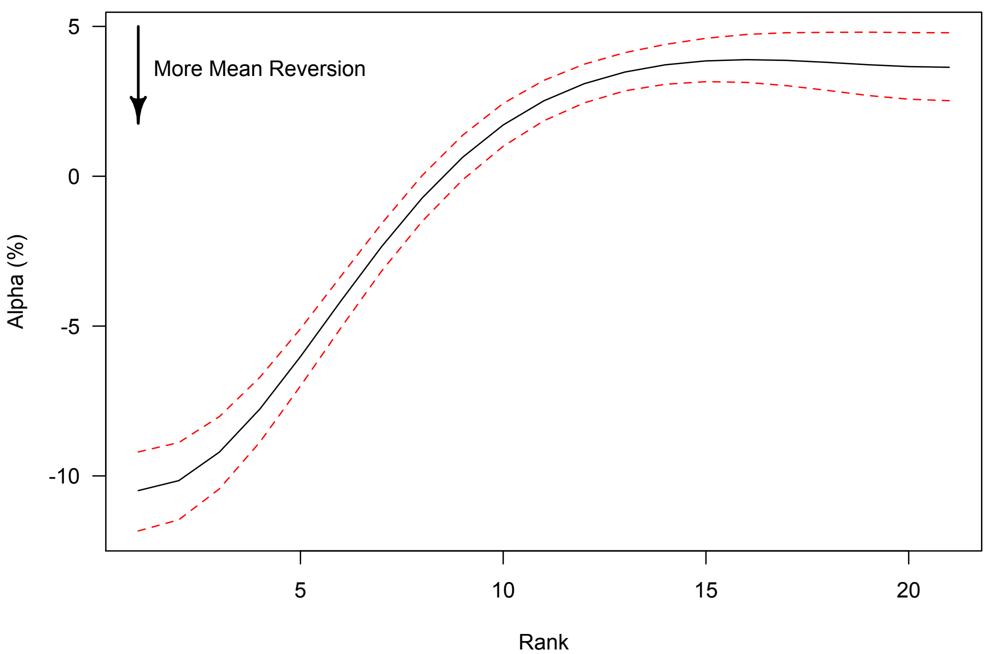

Figure 2: Point estimates and 95% confidence intervals of minus the reversion rates (α_k) for different ranked commodities, 1980 - 2015.

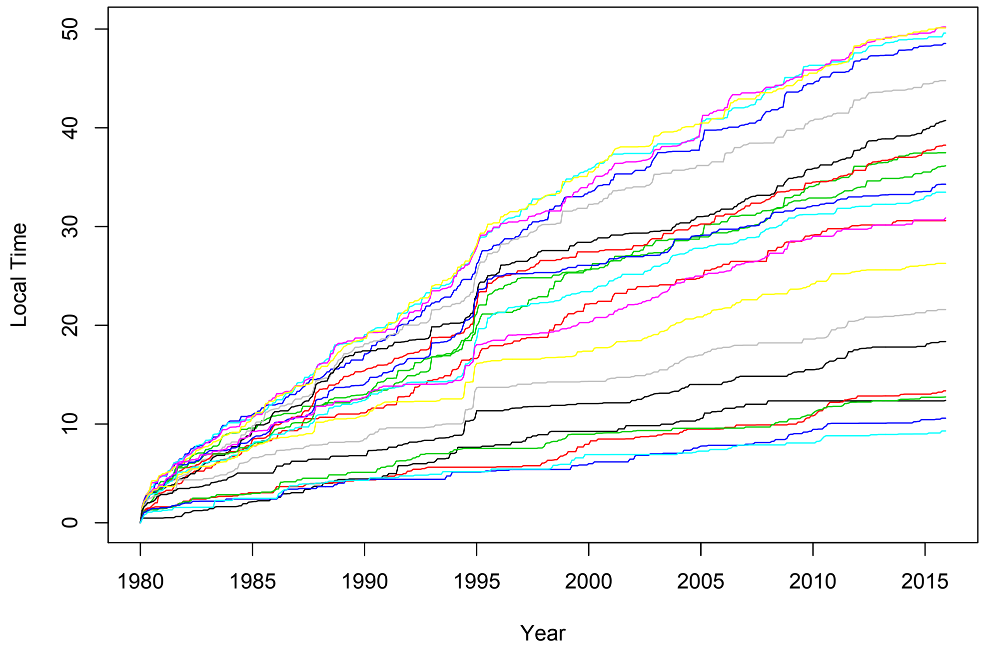

Figure 3: Local time processes ($\Lambda_{\log \tilde{x}_{(k)} - \log \tilde{x}_{(k+1)}}$) for different ranked commodities, 1980 - 2015.

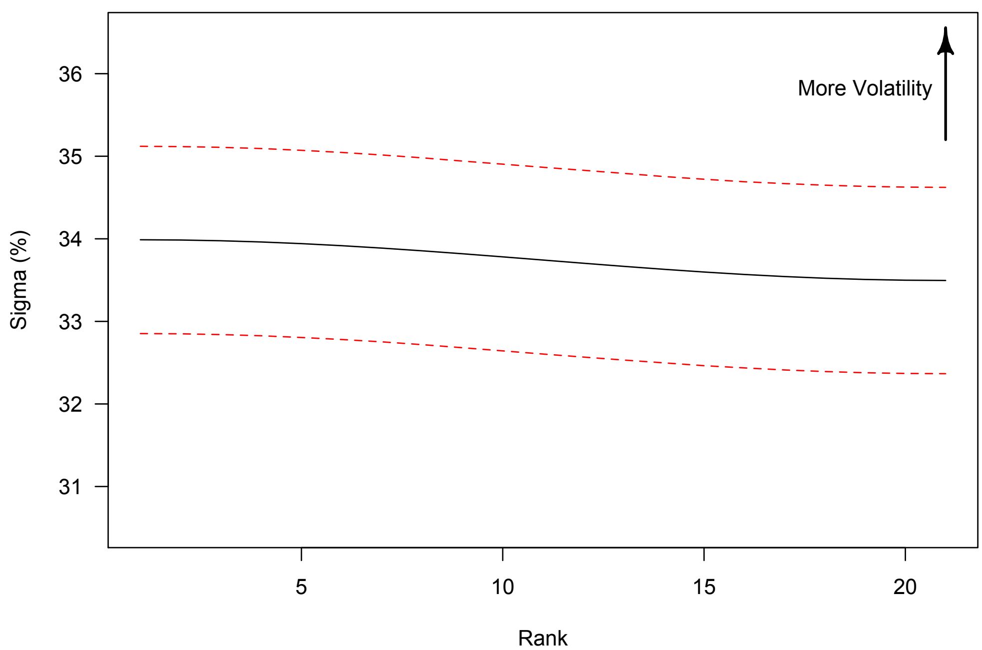

Figure 4: Point estimates and 95% confidence intervals of standard deviations of idiosyncratic commodity price volatilities (σ_k) for different ranked commodities, 1980 - 2015.

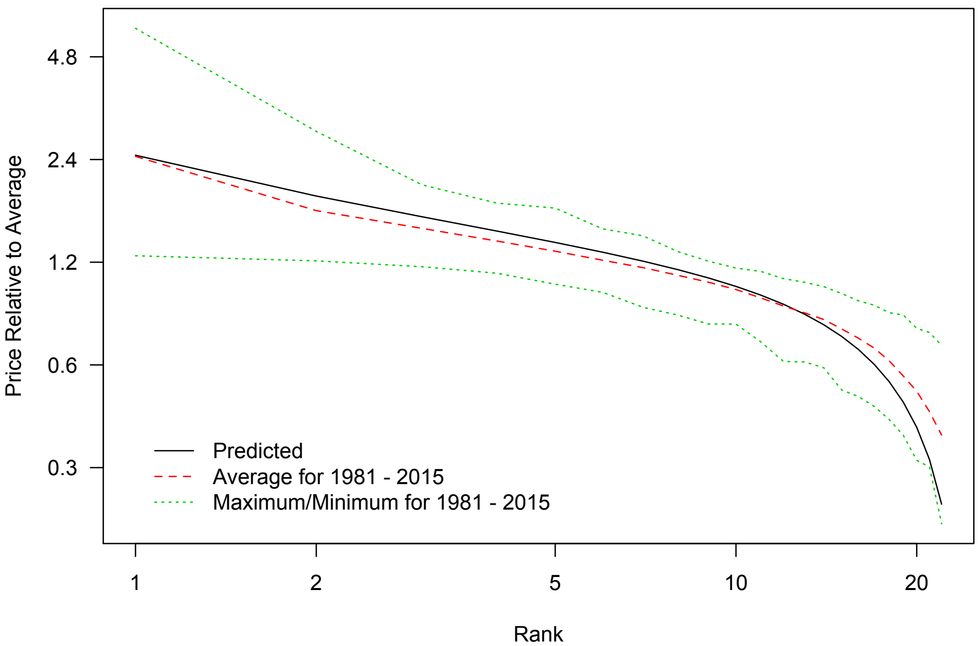

Figure 5: Relative commodity prices for different ranked commodities for 1981 - 2015 as compared to the predicted relative prices.