

THE GRANULAR ORIGINS OF AGGREGATE FLUCTUATIONS

BY XAVIER GABAIX¹

This paper proposes that idiosyncratic firm-level shocks can explain an important part of aggregate movements and provide a microfoundation for aggregate shocks. Existing research has focused on using aggregate shocks to explain business cycles, arguing that individual firm shocks average out in the aggregate. I show that this argument breaks down if the distribution of firm sizes is fat-tailed, as documented empirically. The idiosyncratic movements of the largest 100 firms in the United States appear to explain about one-third of variations in output growth. This “granular” hypothesis suggests new directions for macroeconomic research, in particular that macroeconomic questions can be clarified by looking at the behavior of large firms. This paper’s ideas and analytical results may also be useful for thinking about the fluctuations of other economic aggregates, such as exports or the trade balance.

KEYWORDS: Business cycle, idiosyncratic shocks, productivity, Solow residual, granular residual.

1. INTRODUCTION

THIS PAPER PROPOSES a simple origin of aggregate shocks. It develops the view that a large part of aggregate fluctuations arises from idiosyncratic shocks to individual firms. This approach sheds light on a number of issues that are difficult to address in models that postulate aggregate shocks. Although economy-wide shocks (inflation, wars, policy shocks) are no doubt important, they have difficulty explaining most fluctuations (Cochrane (1994)). Often, the explanation for year-to-year jumps of aggregate quantities is elusive. On the other hand, there is a large amount of anecdotal evidence of the importance of idiosyncratic shocks. For instance, the Organization for Economic Cooperation and Development (OECD (2004)) analyzed that, in 2000, Nokia contributed 1.6 percentage points of Finland’s gross domestic product (GDP) growth.² Likewise, shocks to GDP may stem from a variety of events, such as successful

¹For excellent research assistance, I thank Francesco Franco, Jinsook Kim, Farzad Saidi, Hei-wai Tang, Ding Wu, and, particularly, Alex Chinco and Fernando Duarte. For helpful comments, I thank the co-editor, four referees, and seminar participants at Berkeley, Boston University, Brown, Columbia, ECARES, the Federal Reserve Bank of Minneapolis, Harvard, Michigan, MIT, New York University, NBER, Princeton, Toulouse, U.C. Santa Barbara, Yale, the Econometric Society, the Stanford Institute for Theoretical Economics, and Kenneth Arrow, Robert Barsky, Susanto Basu, Roland Bénabou, Olivier Blanchard, Ricardo Caballero, David Canning, Andrew Caplin, Thomas Chaney, V. V. Chari, Larry Christiano, Diego Comin, Don Davis, Bill Dupor, Steve Durlauf, Alex Edmans, Martin Eichenbaum, Eduardo Engel, John Fernald, Jesus Fernandez-Villaverde, Richard Frankel, Mark Gertler, Robert Hall, John Haltiwanger, Chad Jones, Boyan Jovanovic, Finn Kydland, David Laibson, Arnaud Manas, Ellen McGrattan, Todd Mitton, Thomas Philippon, Robert Solow, Peter Temin, Jose Tessada, and David Weinstein. I thank for NSF (Grant DMS-0938185) for support.

²The example of Nokia is extreme but may be useful. In 2003, worldwide sales of Nokia were \$37 billion, representing 26% of Finland’s GDP of \$142 billion. This is not sufficient for a proper




FIGURE 1.—Sum of the sales of the top 50 and 100 non-oil firms in Compustat, as a fraction of GDP. Hulten’s theorem (Appendix B) motivates the use of sales rather than value added.

innovations by Walmart, the difficulties of a Japanese bank, new exports by Boeing, and a strike at General Motors.³

Since modern economies are dominated by large firms, idiosyncratic shocks to these firms can lead to nontrivial aggregate shocks. For instance, in Korea, the top two firms (Samsung and Hyundai) together account for 35% of exports, and the sales of those two firms account for 22% of Korean GDP (di Giovanni and Levchenko (2009)). In Japan, the top 10 firms account for 35% of exports (Canals, Gabaix, Vilarrubia, and Weinstein (2007)). For the United States, Figure 1 reports the total sales of the top 50 and 100 firms as a fraction of GDP. On average, the sales of the top 50 firms are 24% of GDP, while the sales of the top 100 firms are 29% of GDP. The top 100 firms hence represent a large part of the macroeconomic activity, so understanding their actions offers good insight into the aggregate economy.

In this view, many economic fluctuations are not, primitively, due to small diffuse shocks that directly affect every firm. Instead, many economic fluctuations are attributable to the incompressible “grains” of economic activity, the

assessment of Nokia’s importance, but gives some order of magnitude, as the Finnish base of Nokia is an important residual claimant of the fluctuations of Nokia International.

³Other aggregates are affected as well. For instance, in December 2004, a \$24 billion one-time Microsoft dividend boosted growth in personal income from 0.6% to 3.7% (Bureau of Economic Analysis, January 31, 2005). A macroeconomist would find it difficult to explain this jump in personal income without examining individual firm behavior.

large firms. I call this view the “granular” hypothesis. In the granular view, idiosyncratic shocks to large firms have the potential to generate nontrivial aggregate shocks that affect GDP, and via general equilibrium, all firms.

The granular hypothesis offers a microfoundation for the aggregate shocks of real business cycle models (Kydland and Prescott (1982)). Hence, real business cycle shocks are not, at heart, mysterious “aggregate productivity shocks” or “a measure of our ignorance” (Abramovitz (1956)). Instead, they are well defined shocks to individual firms. The granular hypothesis sheds light on a number of other issues, such as the dependence of the amplitude of GDP fluctuations on GDP level, the microeconomic composition of GDP, and the distribution of GDP and firm-level fluctuations.

In most of this paper, the standard deviation of the percentage growth rate of a firm is assumed to be independent of its size.⁴ This explains why individual firms can matter in the aggregate. If Walmart doubles its number of supermarkets and thus its size, its variance is not divided by 2—as would be the case if Walmart were the amalgamation of many independent supermarkets. Instead, the newly acquired supermarkets inherit the Walmart shocks, and the total percentage variance of Walmart does not change. This paper conceptualizes these shocks as productivity growth, but the analysis holds for other shocks.⁵

The main argument is summarized as follows. First, it is critical to show that $1/\sqrt{N}$ diversification does not occur in an economy with a fat-tailed distribution of firms. A simple diversification argument shows that, in an economy with N firms with independent shocks, aggregate fluctuations should have a size proportional to $1/\sqrt{N}$. Given that modern economies can have millions of firms, this suggests that idiosyncratic fluctuations will have a negligible aggregate effect. This paper points out that when firm size is power-law distributed, the conditions under which one derives the central limit theorem break down and other mathematics apply (see Appendix A). In the central case of Zipf’s law, aggregate volatility decays according to $1/\ln N$, rather than $1/\sqrt{N}$. The strong $1/\sqrt{N}$ diversification is replaced by a much milder one that decays according to $1/\ln N$. In an economy with a fat-tailed distribution of firms, diversification effects due to country size are quite small.

Having established that idiosyncratic shocks do not die out in the aggregate, I show that they are of the correct order of magnitude to explain business cycles. We will see that if firm i has a productivity shock $d\pi_i$, these shocks

⁴The benchmark that the variance of the percentage growth rate is approximately independent of size (“Gibrat’s law” for variances) appears to hold to a good first degree; see Section 2.5.

⁵The productivity shocks can come from a decision of the firm’s research department, of the firm’s chief executive officer, of how to process shipments, inventories, or which new line of products to try. They can also stem from changes in capacity utilization, and, particularly, strikes. Suppose a firm, which uses only capital and labor, is on strike for half the year. For many purposes, its effective productivity that year is halved. This paper does not require the productivity shocks to arise from any particular source.

are independent and identically distributed (i.i.d.) and there is no amplification mechanism, then the standard deviation of total factor productivity (TFP) growth is $\sigma_{\text{TFP}} = \sigma_{\pi} h$, where σ_{π} is the standard deviation of the i.i.d. productivity shocks and h is the sales herfindahl of the economy. Using the estimate of annual productivity volatility of $\sigma_{\pi} = 12\%$ and the sales herfindahl of $h = 5.3\%$ for the United States in 2008, one predicts a TFP volatility equal to $\sigma_{\text{TFP}} = 12\% \cdot 5.3\% = 0.63\%$. Standard amplification mechanisms generate the order of magnitude of business cycle fluctuations, $\sigma_{\text{GDP}} = 1.7\%$. Non-U.S. data lead to even larger business cycle fluctuations. I conclude that idiosyncratic granular volatility seems quantitatively large enough to matter at the macroeconomic level.

Section 3 then investigates accordingly the proportion of aggregate shocks that can be accounted for by idiosyncratic fluctuations. I construct the “granular residual” Γ_t , which is a parsimonious measure of the shocks to the top 100 firms:

$$\Gamma_t := \sum_{i=1}^K \frac{\text{sales}_{i,t-1}}{\text{GDP}_{t-1}} (g_{it} - \bar{g}_t),$$

where $g_{it} - \bar{g}_t$ is a simple measure of the idiosyncratic shock to firm i . Regressing the growth rate of GDP on the granular residual yields an R^2 of roughly one-third. Prima facie, this means that idiosyncratic shocks to the top 100 firms in the United States can explain one-third of the fluctuations of GDP. More sophisticated controls for common shocks confirm this finding. In addition, the granular residual turns out to be a useful novel predictor of GDP growth which complements existing predictors. This supports the view that thinking about firm-level shocks can improve our understanding of GDP movements.

Previous economists have proposed mechanisms that generate macroeconomic shocks from purely microeconomic causes. A pioneering paper is by Jovanovic (1987), whose models generate nonvanishing aggregate fluctuations owing to a multiplier proportional to \sqrt{N} , the square root of the number of firms. However, Jovanovic’s theoretical multiplier of $\sqrt{N} \simeq 1000$ is much larger than is empirically plausible.⁶ Nonetheless, Jovanovic’s model spawned a lively intellectual quest. Durlauf (1993) generated macroeconomic uncertainty with idiosyncratic shocks and local interactions between firms. The drivers of his results are the nonlinear interactions between firms, while in this paper it is the skewed distribution of firms. Bak, Chen, Scheinkman, and Woodford (1993) applied the physical theory of self-organizing criticality. While there is much to learn from their approach, it generates fluctuations more fat-tailed than in reality, with infinite means. Nirei (2006) proposed a model where aggregate fluctuations arise from (s, S) rules at the firm level, in the spirit of Bak

⁶If the actual multiplier were so large, the impact of trade shocks, for instance, would be much higher than we observe.

et al. (1993). These models are conceptually innovative, but they are hard to work with theoretically and empirically. The mechanism proposed in this paper is tractable and relies on readily observable quantities.

Long and Plosser (1983) suggested that sectoral (rather than firm) shocks might account for GDP fluctuations. As their model has a small number of sectors, those shocks can be viewed as miniaggregate shocks. Horvath (2000), as well as Conley and Dupor (2003), explored this hypothesis further. They found that sector-specific shocks are an important source of aggregate volatility. Finally, Horvath (1998) and Dupor (1999) debated whether N sectors can have a volatility that does not decay according to $1/\sqrt{N}$. I found an alternative solution to their debate, which is formalized in Proposition 2. My approach relies on those earlier contributions and clarifies that the fat-tailed nature of the sectoral shocks is important theoretically, as it determines whether the central limit theorem applies.

Studies disagree somewhat on the relative importance of sector-specific shocks, aggregate shocks, and complementarities. Caballero, Engel, and Haltiwanger (1997) found that aggregate shocks are important, while Horvath (1998) concluded that sector-specific shocks go a long way toward explaining aggregate disturbances. Many of these effects in this paper could be expressed in terms of sectors.

Granular effects are likely to be even stronger outside the United States, as the United States is more diversified than most other countries. One number reported in the literature is the value of the assets controlled by the richest 10 families, divided by GDP. Claessens, Djankov, and Lang (2000) found a number equal to 38% in Asia, including 84% of GDP in Hong Kong, 76% in Malaysia, and 39% in Thailand. Faccio and Lang (2002) also found that the top 10 families control 21% of listed assets in their sample of European firms. It would be interesting to transpose the present analysis to those countries and to entities other than firms—for instance, business groups or sectors.

This paper is organized as follows. Section 2 develops a simple model. It also provides a calibration that indicates that the effects are of the right order of magnitude to account for macroeconomic fluctuations. Section 3 shows directly that the idiosyncratic movements of firms appear to explain, year by year, about one-third of actual fluctuations in GDP, and also contains a narrative of the granular residual and GDP. Section 4 concludes.

2. THE CORE IDEA

2.1. *A Simple “Islands” Economy*

This section uses a concise model to illustrate the idea. I consider an islands economy with N firms. Production is exogenous, like in an endowment econ-

omy, and there are no linkages between firms (those will be added later). Firm i produces a quantity S_{it} of the consumption good. It experiences a growth rate

$$(1) \quad \frac{\Delta S_{i,t+1}}{S_{it}} = \frac{S_{i,t+1} - S_{it}}{S_{it}} = \sigma_i \varepsilon_{i,t+1},$$

where σ_i is firm i 's volatility and $\varepsilon_{i,t+1}$ are uncorrelated random variables with mean 0 and variance 1. Firm i produces a homogeneous good without any factor input. Total GDP is

$$(2) \quad Y_t = \sum_{i=1}^N S_{it}$$

and GDP growth is

$$\frac{\Delta Y_{t+1}}{Y_t} = \frac{1}{Y_t} \sum_{i=1}^N \Delta S_{i,t+1} = \sum_{i=1}^N \sigma_i \frac{S_{it}}{Y_t} \varepsilon_{i,t+1}.$$

As the shocks $\varepsilon_{i,t+1}$ are uncorrelated, the standard deviation of GDP growth is $\sigma_{\text{GDP}} = (\text{var } \frac{\Delta Y_{t+1}}{Y_t})^{1/2}$:

$$(3) \quad \sigma_{\text{GDP}} = \left(\sum_{i=1}^N \sigma_i^2 \cdot \left(\frac{S_{it}}{Y_t} \right)^2 \right)^{1/2}.$$

Hence, the variance of GDP, σ_{GDP}^2 , is the weighted sum of the variance σ_i^2 of idiosyncratic shocks with weights equal to $(\frac{S_{it}}{Y_t})^2$, the squared share of output that firm i accounts for. If the firms all have the same volatility $\sigma_i = \sigma$, we obtain

$$(4) \quad \sigma_{\text{GDP}} = \sigma h,$$

where h is the square root of the sales herfindahl of the economy:

$$(5) \quad h = \left[\sum_{i=1}^N \left(\frac{S_{it}}{Y_t} \right)^2 \right]^{1/2}.$$

For simplicity, h will be referred to as the herfindahl of the economy.

This paper works first with the basic model (1)–(2). The arguments apply if general equilibrium mechanisms are added.

2.2. *The 1/√N Argument for the Irrelevance of Idiosyncratic Shocks*

Macroeconomists often appeal to aggregate (or at least sectorwide) shocks, since idiosyncratic fluctuations disappear in the aggregate if there is a large number of firms N . Consider firms of initially identical size equal to $1/N$ of GDP and identical standard deviation $\sigma_i = \sigma$. Then (4)–(5) gives:

$$\sigma_{\text{GDP}} = \frac{\sigma}{\sqrt{N}}.$$

To estimate the order of magnitude of the cumulative effect of idiosyncratic shocks, take an estimate of firm volatility $\sigma = 12\%$ from Section 2.4 and consider an economy with $N = 10^6$ firms.⁷ Then

$$\sigma_{\text{GDP}} = \frac{\sigma}{\sqrt{N}} = \frac{12\%}{10^3} = 0.012\% \text{ per year.}$$

Such a GDP volatility of 0.012% is much too small to account for the empirically measured size of macroeconomic fluctuations of around 1%. This is why economists typically appeal to aggregate shocks. More general modelling assumptions predict a $1/\sqrt{N}$ scaling, as shown by the next proposition.

PROPOSITION 1: *Consider an islands economy with N firms whose sizes are drawn from a distribution with finite variance. Suppose that they all have the same volatility σ . Then the economy’s GDP volatility follows, as $N \rightarrow \infty$*

$$(6) \quad \sigma_{\text{GDP}} \sim \frac{\mathbb{E}[S^2]^{1/2}}{\mathbb{E}[S]} \frac{\sigma}{\sqrt{N}}.$$

PROOF: Since $\sigma_{\text{GDP}} = \sigma h$, I examine h : $N^{1/2}h = \frac{(N^{-1} \sum_{i=1}^N S_i^2)^{1/2}}{N^{-1} \sum_{i=1}^N S_i}$. The law of large numbers ensures that $N^{-1} \sum_{i=1}^N S_i^2 \xrightarrow{\text{a.s.}} \mathbb{E}[S^2]$ and $N^{-1} \sum_{i=1}^N S_i \xrightarrow{\text{a.s.}} \mathbb{E}[S]$. This yields $N^{1/2}h \xrightarrow{\text{a.s.}} \mathbb{E}[S^2]^{1/2}/\mathbb{E}[S]$. Q.E.D.

Proposition 1 will be contrasted with Proposition 2 below, which shows that different models of the size distribution of firms lead to dramatically different results.

2.3. *The Failure of the 1/√N Argument When the Firm Size Distribution Is Power Law*

The firm size distribution, however, is not thin-tailed, as assumed in Proposition 1. Indeed, Axtell (2001), using Census data, found a power law with exponent $\zeta = 1.059 \pm 0.054$. Hence, the size distribution of U.S. firms is well

⁷Axtell (2001) reported that in 1997 there were 5.5 million firms in the United States.

approximated by the power law with exponent $\zeta = 1$, the “Zipf” distribution (Zipf (1949)). This finding holds internationally, and the origins of this distribution are becoming better understood (see Gabaix (2009)). The next proposition examines behavior under a “fat-tailed” distribution of firms.

PROPOSITION 2: *Consider a series of island economies indexed by $N \geq 1$. Economy N has N firms whose growth rate volatility is σ and whose sizes S_1, \dots, S_N are drawn from a power law distribution*

$$(7) \quad \mathbb{P}(S > x) = ax^{-\zeta}$$

for $x > a^{1/\zeta}$, with exponent $\zeta \geq 1$. Then, as $N \rightarrow \infty$, GDP volatility follows

$$(8) \quad \sigma_{\text{GDP}} \sim \frac{v_\zeta}{\ln N} \sigma \quad \text{for } \zeta = 1,$$

$$(9) \quad \sigma_{\text{GDP}} \sim \frac{v_\zeta}{N^{1-1/\zeta}} \sigma \quad \text{for } 1 < \zeta < 2,$$

$$(10) \quad \sigma_{\text{GDP}} \sim \frac{v_\zeta}{N^{1/2}} \sigma \quad \text{for } \zeta \geq 2,$$

where v_ζ is a random variable. The distribution of v_ζ does not depend on N and σ . When $\zeta \leq 2$, v_ζ is the square root of a stable Lévy distribution with exponent $\zeta/2$. When $\zeta > 2$, v_ζ is simply a constant. In other terms, when $\zeta = 1$ (Zipf’s law), GDP volatility decays like $1/\ln N$ rather than $1/\sqrt{N}$.

In the above proposition, an expression like $\sigma_{\text{GDP}} \sim \frac{v_\zeta}{N^{1-1/\zeta}} \sigma$ means $\sigma_{\text{GDP}} \times N^{1-1/\zeta}$ converges to $v_\zeta \sigma$ in distribution. More formally, for a series of random variables X_N and of positive numbers a_N , $X_N \sim a_N Y$ means that $X_N/a_N \xrightarrow{d} Y$ as $N \rightarrow \infty$, where \xrightarrow{d} is the convergence in distribution.

I comment on the economics of Proposition 2 before proving it. The firm size distribution has thin tails, that is, finite variance, if and only if $\zeta > 2$. Proposition 1 states that if the firm size distribution has thin tails, then σ_{GDP} decays according to $1/\sqrt{N}$. In contrast, Proposition 2 states that if the firm size distribution has fat tails ($\zeta < 2$), then σ_{GDP} decays much more slowly than $1/\sqrt{N}$: it decays as $1/N^{1-1/\zeta}$.

To get the intuition for the scaling, take the case $a = 1$ and observe that (7) implies that “typical” size S_1 of the largest firm is such that $S_1^{-\zeta} = 1/N$, hence $S_1 = N^{1/\zeta}$ (see Sornette (2006) for that type of intuition). In contrast, GDP is $Y \simeq N\mathbb{E}[S]$ when $\zeta > 1$ by the law of large numbers. Hence, the share of the largest firm is $S_1/Y = N^{-(1-1/\zeta)}/\mathbb{E}[S] \propto N^{-(1-1/\zeta)}$.⁸ this is a small decay when

⁸Here $f(Y) \propto g(Y)$ for some functions f, g means that the ratio $f(Y)/g(Y)$ tends, for large Y , to be a positive real number. So f and g have the same scaling “up to a constant factor.”

ζ is close to 1. Likewise, the size of the top k firms satisfies $S_k^{-\zeta} = k/N$, so $S_k = (N/k)^{1/\zeta}$. Hence, the share of the largest K firms (for a fixed K) is proportional to $N^{-(1-1/\zeta)}$. Plugging this into (5), we see that the herfindahl, and GDP volatility, is proportional to $N^{-(1-1/\zeta)}$.

In the case $\zeta = 1$, $\mathbb{E}[S] = \infty$, so GDP cannot be $Y \simeq N\mathbb{E}[S]$. The following heuristic reasoning gives the correct value. As firm size density is x^{-2} and we saw that the largest firm has typical size N , the typical average firm size is $\bar{S}_N = \int_1^N x^{-2}x dx = \ln N$, and then $Y \simeq N\bar{S}_N = N \ln N$. Hence, the share of the top firm is $S_1/Y = 1/\ln N$. By the above reasoning, GDP volatility is proportional to $1/\ln N$.

The perspective of Proposition 2 is that of an economist who knows the GDP of various countries, but not the size of their respective firms, except that, for instance, they follow Zipf’s law. Then he would conclude that the volatility of a country of size N should be proportional to $1/\ln N$. This explains the v_ζ terms in the distribution of σ_{GDP} : when $\zeta < 2$, GDP volatility (and the herfindahl h) depends on the specific realization of the size distribution of top firms. Because of the fat-tailedness of the distribution of firms, σ_{GDP} does not have a degenerate distribution even as $N \rightarrow \infty$. For the same reason, when $\zeta > 2$, the law of large numbers applies and the distribution of volatility does become degenerate. Of course, if the economist knows the actual size of the firms, then she could calculate the standard deviation of GDP directly by calculating the herfindahl index. Note also that as GDP is made of some large firms, GDP fluctuations are typically not Gaussian (mathematically, the Lindeberg–Feller theorem does not apply, because there are some large firms). The ex ante distribution is developed further in Proposition 3.

Having made these remarks about the meaning of Proposition 2, let me present its proof.

PROOF OF PROPOSITION 2: Since $\sigma_{\text{GDP}} = \sigma h$, I examine

$$(11) \quad h = \frac{N^{-1} \left(\sum_{i=1}^N S_i^2 \right)^{1/2}}{N^{-1} \sum_{i=1}^N S_i}.$$

I observe that when $\zeta > 1$, the law of large numbers gives

$$(12) \quad N^{-1} \sum_{i=1}^N S_i \rightarrow \mathbb{E}[S]$$

almost surely, so

$$h \sim \frac{N^{-1} \left(\sum_{i=1}^N S_i^2 \right)^{1/2}}{\mathbb{E}[S]}.$$

I will first complete the above heuristic proof for the scaling as a function N , which will be useful to ground the intuition, and then present a formal proof which relies on the heavier machinery of Lévy's theorem.

Heuristic Proof. For simplicity, I normalize $a = 1$. I observe that the size of the i th largest firm is approximately

$$(13) \quad S_{i,N} = \left(\frac{i}{N} \right)^{-1/\xi}.$$

The reason for (13) is the following. As the counter-cumulative distribution function (CDF) of the distribution is $x^{-\xi}$, the random variable $S^{-\xi}$ follows a uniform distribution. Hence, the size of firm number i out of N follows $\mathbb{E}[S_{i,N}^{-\xi}] = i/(N + 1)$. So in a heuristic sense, we have $S_{i,N}^{-\xi} \simeq i/(N + 1)$ or, more simply, (13).

From representation (13), the herfindahl can be calculated as

$$h_N \sim \frac{N^{-1+1/\xi} \left(\sum_{i=1}^N i^{-2/\xi} \right)^{1/2}}{\mathbb{E}[S]}.$$

In the fat-tailed case, $\xi < 2$, the series $\sum_{i=1}^{\infty} i^{-2/\xi}$ converges, hence

$$h_N \sim \frac{N^{-1+1/\xi} \left(\sum_{i=1}^{\infty} i^{-2/\xi} \right)^{1/2}}{\mathbb{E}[S]} = CN^{-1+1/\xi}$$

for a constant C . Volatility scales as $N^{-1+1/\xi}$, as in (9).

In contrast, in the finite-variance case, the series $\sum_{i=1}^{\infty} i^{-2/\xi}$ diverges and we have $\sum_{i=1}^N i^{-2/\xi} \sim \int_1^N i^{-2/\xi} di \sim N^{1-2/\xi}/(1 - 2/\xi)$, so that

$$h_N \sim \frac{N^{-1+1/\xi} (N^{1-2/\xi}/(1 - 2/\xi))^{1/2}}{\mathbb{E}[S]} = C'N^{-1/2},$$

and as expected volatility scales as $N^{-1/2}$.

Rigorous Proof. When $\zeta > 2$, the variance of firm sizes is finite and I use Proposition 1. When $\zeta \leq 2$, I observe that S_i^2 has power-law exponent $\zeta/2 \leq 1$, as shown by

$$\mathbb{P}(S^2 > x) = \mathbb{P}(S > x^{1/2}) = a(x^{1/2})^{-\zeta} = ax^{-\zeta/2}.$$

So to handle the numerator of (11), I use Lévy’s theorem from Appendix A. This implies

$$N^{-2/\zeta} \sum_{i=1}^N S_i^2 \xrightarrow{d} u,$$

where u is a Lévy-distributed random variable with exponent $\zeta/2$. So when $\zeta \in (1, 2]$, I can use the fact (12) to conclude

$$N^{1-1/\zeta} h = \frac{\left(N^{-2/\zeta} \sum_{i=1}^N S_i^2 \right)^{1/2}}{N^{-1} \sum_{i=1}^N S_i} \xrightarrow{d} \frac{u^{1/2}}{\mathbb{E}[S]}.$$

When $\zeta = 1$, additional care is required, because $\mathbb{E}[S] = \infty$. Lévy’s theorem applied to $X_i = S_i$ gives $a_N = N$ and $b_N = N \ln N$, hence

$$\frac{1}{N} \left(\sum_{i=1}^N S_i - N \ln N \right) \xrightarrow{d} g,$$

where g follows a Lévy distribution with exponent 1, which implies

$$(14) \quad Y = \sum_{i=1}^N S_i \sim N \ln N.$$

I conclude $h \sim u^{1/2} / \ln N$.

Q.E.D.

I conclude with a few remarks. Proposition 2 offers a resolution to the debate between Horvath (1998, 2000) and Dupor (1999). Horvath submitted evidence that sectoral shocks may be enough to generate aggregate fluctuations. Dupor (1999) debated this on theoretical grounds and claimed that Horvath was able to generate large aggregate fluctuations only because he used a moderate number of sectors ($N = 36$). If he had many more finely disaggregated sectors (e.g., 100 times as many), then aggregate volatility would decrease in $1/\sqrt{N}$ (e.g., 10 times smaller). Proposition 2 illustrates that both viewpoints are correct, but apply in different settings. Dupor’s reasoning holds only in a world

of small firms, when the central limit theorem can apply. Horvath's empirical world is one where the size distribution of firms is sufficiently fat-tailed that the central limit theorem does not apply. Instead, Proposition 2 applies and GDP volatility remains substantial even if the number N of subunits is large.

Though the benchmark case of Zipf's law is empirically relevant, and theoretically clean and appealing, many arguments in this paper do not depend on it. The results only require that the herfindahl of actual economies is sufficiently large. For instance, if the distribution of firm sizes were lognormal with a sufficiently high variance, then quantitatively very little would change.

The herfindahls generated by a Zipf distribution are reasonably high. For $N = 10^6$ firms, with an equal distribution of sizes, $h = 1/\sqrt{N} = 0.1\%$, but in a Zipf world with $\zeta = 1$, Monte Carlo simulations show that the median $h = 12\%$. With a firm volatility of $\sigma = 12\%$, this corresponds to a GDP volatility σh of 0.012% for identically sized firms and a more respectable 1.4% for a Zipf distribution of firm sizes. This is the theory under the Zipf benchmark, which has a claim to hold across countries and clarifies what we can expect independently of the imperfections of data sets and data collection.

2.4. *Can Granular Effects Be Large Enough in Practice? A Calibration*

I now examine how large we can expect granular effects to be. For greater realism, I incorporate two extra features compared to the island economy: input–output linkages and the endogenous response in inputs to initial disturbances. I start with the impact of linkages.

2.4.1. *Economies With Linkages*

Consider an economy with N competitive firms buying intermediary inputs from one another. Let firm i have Hicks-neutral productivity growth $d\pi_i$. Hulten (1978) showed that the increase in aggregate TFP is⁹

$$(15) \quad \frac{d\text{TFP}}{\text{TFP}} = \sum_i \frac{\text{sales of firm } i}{\text{GDP}} d\pi_i.$$

This formula shows that, somewhat surprisingly, we can calculate TFP shocks without knowing the input–output matrix: the sufficient statistic for the impact of firm i is its size, as measured by its sales (i.e., gross output rather than net output). This helps simplify the analysis.¹⁰ In addition, the weights add up to more than 1. This reflects the fact that productivity growth of 1% in a firm

⁹For completeness, Appendix B rederives and generalizes Hulten's theorem.

¹⁰However, to study the propagation of shocks and the origin of size, the input–output matrix can be very useful. See Carvalho (2009) and Acemoglu, Ozdaglar, and Tahbaz-Salehi (2010), who studied granular effects in the economy viewed as a network.

generates an increase in produced values equal to 1% times its sales, not times its sales net of inputs (which would be the value added). The firm's sales are the proper statistic for that social value.

I now draw the implications for TFP volatility. Suppose productivity shocks $d\pi_i$ are uncorrelated with variance σ_π^2 . Then the variance of productivity growth is

$$(16) \quad \text{var} \frac{d\text{TFP}}{\text{TFP}} = \sum_i \left(\frac{\text{sales of firm } i}{\text{GDP}} \right)^2 \text{var}(d\pi_i)$$

and so the volatility of the growth of TFP is

$$(17) \quad \sigma_{\text{TFP}} = h\sigma_\pi,$$

where h is the sales herfindahl,

$$(18) \quad h = \left(\sum_{i=1}^N \left(\frac{\text{sales}_{it}}{\text{GDP}_t} \right)^2 \right)^{1/2}.$$

I now examine the empirical magnitude of the key terms in (17), starting with σ_π .

2.4.2. Large Firms Are Very Volatile

Most estimates of plant-level volatility find very large volatilities of sales and employment, with an order of magnitude $\sigma = 30\text{--}50\%$ per year (e.g., Caballero, Engel, and Haltiwanger (1997), Davis, Haltiwanger, and Schuh (1996)). Also, the volatility of firm size in Compustat is a very large, 40% per year (Comin and Mullani (2006)). Here I focus the analysis on the top 100 firms. Measuring firm volatility is difficult, because various frictions and identifying assumptions provide conflicting predictions about links between changes in total factor productivity and changes in observable quantities such as sales and employment. I consider the volatility of three measures of growth rates: $\Delta \ln(\text{sales}_{it}/\text{employees}_{it})$, $\Delta \ln \text{sales}_{it}$, and $\Delta \ln \text{employees}_{it}$. For each measure and each year, I calculate the cross-sectional variance among the top 100 firms of the previous year and take the average.¹¹ I find standard deviations of 12%, 12%, and 14% for, respectively, growth rates of the sales per employee, of sales, and of employees. Also, among the top 100 firms, the sample correlations are 0.023, 0.073, and 0.033, respectively, for each of the three measures.¹²

¹¹In other terms, for each year t , I calculate the cross-sectional variance of growth rates, $\sigma_t^2 = K^{-1} \sum_{i=1}^K g_{it}^2 - (K^{-1} \sum_{i=1}^K g_{it})^2$, with $K = 100$. The corresponding average standard deviation is $[\overline{T^{-1} \sum_{t=1}^T \sigma_t^2}]^{1/2}$.

¹²For each year, we measure the sample correlation $\rho_t = [\frac{1}{K(K-1)} \sum_{i \neq j} g_{it} g_{jt}] / [\frac{1}{K} \sum_i g_{it}^2]$, with $K = 100$. The correlations are positive. Note that a view that would attribute the major firm-level

Hence, the correlation between growth rates is small. At the firm level, most variation is idiosyncratic.

In conclusion, the top 100 firms have a volatility of 12% based on sales per employee. In what follows, I use $\sigma_\pi = 12\%$ per year for firm-level volatility as the baseline estimate.

2.4.3. Herfindahls and Induced Volatility

I next consider the impact of endogenous factor usage on GDP. Calling Λ TFP, many models predict that when there are no other disturbances, GDP growth dY/Y is proportional to TFP growth $d\Lambda/\Lambda$: $dY/Y = \mu d\Lambda/\Lambda$ for some $\mu \geq 1$ that reflects factor usage; alternatively, via (15),

$$(19) \quad \frac{dY}{Y} = \mu \sum_i \frac{\text{sales of firm } i}{Y} d\pi_i.$$

This gives a volatility of GDP equal to $\sigma_{\text{GDP}} = \mu \sigma_{\text{TFP}}$, and via (17),

$$(20) \quad \sigma_{\text{GDP}} = \mu \sigma_\pi h.$$

To examine the size of μ , I consider a few benchmarks. In a short-term model where capital is fixed in the short run and the Frisch elasticity of labor supply is ϕ , $\mu = 1/(1 - \alpha\phi/(1 + \phi))$, and if the supply of capital is flexible (e.g., via variable utilization or the current account), then $\mu = (1 + \phi)/\alpha$.¹³ With an effective Frisch elasticity of 2 (as recommend by Hall (2009) for an inclusive elasticity that includes movements in and out of the labor force), those values are $\mu = 1.8$ and $\mu = 4.5$. If TFP is a geometrical random walk, in the neoclassical growth model where only capital can be accumulated, in the long run, we have $\mu = 1/\alpha$, where α is the labor share; with $\alpha = 2/3$, this gives $\mu = 1.5$.¹⁴ I use the average of the three above values, $\mu = 2.6$.

Empirically, the sales herfindahl h is quite large: $h = 5.3\%$ for the United States in 2008 and $h = 22\%$ in an average over all countries.¹⁵ This means, parenthetically, that the United States is a country with relatively small firms (compared to GDP), where the granular hypothesis might be the hardest to establish.

movements to shocks to the relative demand for a firm's product compared to its competitors would counterfactually predict a negative correlation.

¹³This can be seen by solving $\max_L \Lambda K^{1-\alpha} L^\alpha - L^{1+1/\phi}$ or $\max_{K,L} \Lambda K^{1-\alpha} L^\alpha - rK - L^{1+1/\phi}$, respectively, which gives $Y \propto \Lambda^\mu$ for the announced value of μ . For this derivation, I use the local representation with a quasilinear utility function, but the result does not depend on that.

¹⁴If $Y_t = A_t K_t^{1-\alpha} L_t^\alpha$, $A_t \propto e^{\gamma t}$, and capital is accumulated, then in a balanced growth path, $Y_t \propto K_t \propto A_t^{1/\alpha}$. This holds also with stochastic growth.

¹⁵The U.S. data are from Compustat. The international herfindahls are from Acemoglu, Johnson, and Mitton (2009). They analyzed the Dun and Bradstreet data set, which has a good coverage of the major firms in many countries, though not a complete or homogeneous one.

I can now incorporate all those numbers, using $\sigma_\pi = 12\%$ seen above. Equation 20 yields a GDP volatility $\sigma_{\text{GDP}} = 2.6 \times 12\% \times 5.3\% = 1.7\%$ for the United States, and $\sigma_{\text{GDP}} = 2.6 \times 12\% \times 22\% = 6.8\%$ for a typical country. This is very much on the order of magnitude of GDP fluctuations. As always, further amplification mechanisms can increase the estimate. I conclude that idiosyncratic volatility seems quantitatively large enough to matter at the macroeconomic level.

2.5. *Extension: GDP Volatility When the Volatility of a Firm Depends on Its Size*

I now study the case where the volatility of a firm’s percentage growth rate decreases with firm size, which will confirm the robustness of the previous results and yield additional predictions. I examine the functional form $\sigma^{\text{firm}}(S) = kS^{-\alpha}$ from (21). If $\alpha > 0$, then large firms have a smaller standard deviation than small firms. Stanley, Amaral, Buldyrev, Havlin, Leschhorn, Maass, Salinger, and Stanley (1996) quantified the relation more precisely and showed that (21) holds for firms in Compustat, with $\alpha \simeq 1/6$.

It is unclear whether the conclusions from Compustat can generalize to the whole economy. Compustat only comprises firms traded on the stock market and these are probably more volatile than nontraded firms, as small volatile firms are more prone to seek outside equity financing, while large firms are in any case very likely to be listed in the stock market. This selection bias implies that the value of α measured from Compustat firms alone is presumably larger than in a sample composed of all firms. It is indeed possible α may be 0 when estimated on a sample that includes all firms, as random growth models have long postulated. In any case, any deviations from Gibrat’s law for variances appear to be small, that is, $0 \leq \alpha \leq 1/6$. If there is no diversification as size increases, then $\alpha = 0$. If there is full diversification and a firm of size S is composed of S units, then $\alpha = 1/2$. Empirically, firms are much closer to the Gibrat benchmark of no diversification, $\alpha = 0$.

The next proposition extends Propositions 1 and 2 to the case where firm volatility decreases with firm size.

PROPOSITION 3: *Consider an islands economy, with N firms that have power-law distribution $\mathbb{P}(S > x) = (S_{\min}/x)^\zeta$ for $\zeta \in [1, \infty)$. Assume that the volatility of a firm of size S is*

$$(21) \quad \sigma^{\text{firm}}(S) = \bar{\sigma} \left(\frac{S}{S_{\min}} \right)^{-\alpha}$$

for some $\alpha \geq 0$ and the growth rate is $\Delta S/S = \sigma^{\text{firm}}(S)u$, where $\mathbb{E}[u] = 0$. Define $\zeta' = \zeta/(1 - \alpha)$ and $\alpha' = \min(1 - 1/\zeta', 1/2)$, so that $\alpha' = 1/2$ for $\zeta' \geq 2$. GDP fluctuations have the following form. If $\zeta > 1$,

$$(22) \quad \frac{\Delta Y}{Y} \sim N^{-\alpha'} \frac{\zeta - 1}{\zeta} \mathbb{E}[|u|^{\zeta'}]^{1/\zeta'} \bar{\sigma} g_{\zeta'}, \quad \text{if } \zeta' < 2,$$

$$(23) \quad \frac{\Delta Y}{Y} \sim N^{-\alpha'} \frac{\zeta - 1}{\zeta} \frac{\mathbb{E}[S^2 \sigma^{\text{firm}}(S)^2]^{1/2} \mathbb{E}[u^2]^{1/2}}{S_{\min}} g_2, \quad \text{if } \zeta' \geq 2,$$

where $g_{\zeta'}$ is a standard Lévy distribution with exponent ζ' . Recall that g_2 is simply a standard Gaussian distribution. If $\zeta = 1$,

$$(24) \quad \frac{\Delta Y}{Y} \sim \frac{N^{-\alpha'}}{\ln N} \mathbb{E}[|\varepsilon|^{\zeta'}]^{1/\zeta'} \bar{\sigma} g_{\zeta'}, \quad \text{if } \zeta' < 2,$$

$$(25) \quad \frac{\Delta Y}{Y} \sim \frac{N^{-\alpha'}}{\ln N} \frac{\mathbb{E}[S^2 \sigma^{\text{firm}}(S)^2]^{1/2} \mathbb{E}[u^2]^{1/2}}{S_{\min}} g_2, \quad \text{if } \zeta' \geq 2.$$

In particular, the volatility $\sigma(Y)$ of GDP growth decreases as a power-law function of GDP Y ,

$$(26) \quad \sigma^{\text{GDP}}(Y) \propto Y^{-\alpha'}.$$

To see the intuition for Proposition 3, we apply the case of Zipf's law ($\zeta = 1$) to an example with two large countries, 1 and 2, in which country 2 has twice as many firms as country 1. Its largest K firms are twice as large as the largest firms of country 1. However, scaling according to (21) implies that their volatility is $2^{-\alpha}$ times the volatility of firms in country 1. Hence, the volatility of country 2's GDP is $2^{-\alpha}$ times the volatility of country 1's GDP (i.e., (26)). Putting this another way, under the case presented by Proposition 3 and $\zeta = 1$, large firms are less volatile than small firms (equation (21)). The top firms in big countries are larger (in an absolute sense) than top firms in small countries. As the top firms determine a country's volatility, big countries have less volatile GDP than small countries (equation (26)).

Also, one can reinterpret Proposition 3 by interpreting a large firm as a "country" made up of smaller entities. If these entities follow a power-law distribution, then Proposition 3 applies and predicts that the fluctuations of the growth rate $\Delta \ln S_{it}$, once rescaled by $S_{it}^{-\alpha}$, follow a Lévy distribution with exponent $\min\{\zeta/(1-\alpha), 2\}$. Lee, Amaral, Meyer, Canning, and Stanley (1998) plotted this empirical distribution, which looks roughly like a Lévy stable distribution. It could be that the fat-tailed distribution of firm growth comes from the fat-tailed distribution of the subcomponents of a firm.¹⁶

A corollary of Proposition 3 may be worth highlighting.

COROLLARY 1—Similar Scaling of Firms and Countries: *When Zipf's law holds ($\zeta = 1$) and $\alpha \leq 1/2$, we have $\alpha' = \alpha$, that is, firms and countries should see their volatility scale with a similar exponent:*

$$(27) \quad \sigma^{\text{firms}}(S) \propto S^{-\alpha}, \quad \sigma^{\text{GDP}}(Y) \propto Y^{-\alpha}.$$

¹⁶See Sutton (2002) for a related model, and Wyart and Bouchaud (2003) for a related analysis, which acknowledges the contribution of the present article, which was first circulated in 2001.

Interestingly, Lee et al. (1998) presented evidence that supports (27), with a small exponent $\alpha \simeq 1/6$ (see also Koren and Tenreyro (2007)). A more systematic investigation of this issue would be interesting.

Finally, Proposition 3 adopts the point of view of an economist who would not know the sizes of firms in the country. Then the best guess is a Lévy distribution of GDP fluctuations. However, given precise knowledge of the size of firms, GDP fluctuations will depend on the details of the distribution of the microeconomic shocks u_i .

Before concluding this theoretical section, let me touch on another very salient feature of business cycles: firms and sectors comove. As seen by Long and Plosser (1983), models with production and demand linkages can generate comovement. Carvalho and Gabaix (2010) worked out such a model with purely idiosyncratic shocks and demand linkages. In that economy, the equilibrium growth rates of sales, employees, and labor productivity can be expressed as

$$(28) \quad g_{it} = a\varepsilon_{it} + bf_t, \quad f_t \equiv \sum_{j=1}^N \frac{S_{j,t-1}}{Y_{t-1}} \varepsilon_{jt},$$

where ε_{it} is the firm idiosyncratic productivity shock. Hence, the economy is a one-factor model, but, crucially, the common factor f_t is nothing but a sum of the idiosyncratic firm shocks. In their calibration, over 90% of output variance will be attributed to comovement, as in the empirical findings of Shea (2002). Hence, a calibrated granular model with linkages and only idiosyncratic shocks may account for a realistic amount of comovement. This arguably good feature of granular economies generates econometric challenges, as we shall now see.

3. TENTATIVE EMPIRICAL EVIDENCE FROM THE GRANULAR RESIDUAL

3.1. *The Granular Residual: Motivation and Definition*

This section presents tentative evidence that the idiosyncratic movements of the top 100 firms explain an important fraction (one-third) of the movement of total factor productivity (TFP). The key challenge is to identify idiosyncratic shocks. Large firms could be volatile because of aggregate shocks, rather than the other way around. There is no general solution for this “reflection problem” (Manski (1993)). I use a variety of ways to measure the share of idiosyncratic shocks.

I start with a parsimonious proxy for the labor productivity of firm i , the log of its sales per worker:

$$(29) \quad z_{it} := \ln \frac{\text{sales of firm } i \text{ in year } t}{\text{number of employees of firm } i \text{ in year } t}.$$

This measure is selected because it requires only basic data that are more likely to be available for non-U.S. countries, unlike more sophisticated measures

such as a firm-level Solow residual. Most studies that construct productivity measures from Compustat data use (29). I define the productivity growth rate as $g_{it} = z_{it} - z_{it-1}$. Various models (including the one in the National Bureau of Economic Research (NBER) working paper version of this article) predict that, indeed, the productivity growth rate is closely related to g_{it} .

Suppose that productivity evolves as

$$(30) \quad g_{it} = \beta' X_{it} + \varepsilon_{it},$$

where X_{it} is a vector of factors that may depend on firm characteristics at time $t - 1$ and on factors at time t (e.g., as in equation (28)). My goal is to investigate whether ε_{it} , the idiosyncratic component of the total factor productivity growth rate of large firms, can explain aggregate TFP. More precisely, I would like to empirically approximate the ideal granular residual Γ_t^* , which is the direct rewriting of (15):

$$(31) \quad \Gamma_t^* := \sum_{i=1}^K \frac{S_{i,t-1}}{Y_{t-1}} \varepsilon_{it}.$$

It is the sum of idiosyncratic firm shocks, weighted by size. I wish to see what fraction of the total variance of GDP growth comes from the granular residual, as the theory (19) predicts that GDP growth is $g_{Yt} = \mu \Gamma_t^*$.

I need to extract ε_{it} . To do so, I estimate (30) for the top $Q \geq K$ firms of the previous year, on a vector of observables that I will soon specify. I then form the estimate of idiosyncratic firm-level productivity shock as $\widehat{\varepsilon}_{it} = g_{it} - \widehat{\beta}' X_{it}$. I define the “granular residual” Γ_t as

$$(32) \quad \Gamma_t := \sum_{i=1}^K \frac{S_{i,t-1}}{Y_{t-1}} \widehat{\varepsilon}_{it}.$$

Identification is achieved if the measured granular residual Γ_t is close to the ideal granular residual Γ_t^* .

Two particularizations are useful, because they do not demand much data and are transparent. They turn out to do virtually as well as the more complicated procedures I will also consider. The simplest specification is to control for the mean growth rate in the sample, that is, to have $X_{it} = \bar{g}_t$, where $\bar{g}_t = Q^{-1} \sum_{i=1}^Q g_{it}$. Here, I take the average over the top Q firms. We could have $Q = K$ or take the average over more firms. In practice, I will calculate the granular residual over the top $K = 100$ firms, but take the averages for the controls over the top $Q = 100$ or 1000 firms. Then the granular residual is the weighted sum of the firm’s growth rate minus the average firm growth rate:

$$(33) \quad \Gamma_t = \sum_{i=1}^K \frac{S_{i,t-1}}{Y_{t-1}} (g_{it} - \bar{g}_t).$$

Another specification is to control for the mean growth $\bar{g}_{i,t}$, the equal-weighted average productivity growth rate among firms that are in i 's industry and among the top Q firms therein. Then $X_{it} = \bar{g}_{i,t}$. That gives

$$(34) \quad \Gamma_t = \sum_{i=1}^K \frac{S_{i,t-1}}{Y_{t-1}} (g_{it} - \bar{g}_{i,t}).$$

It is the weighted sum of the firm growth rates minus the growth rates of other firms in the same industry. The term $g_{it} - \bar{g}_{i,t}$ may be closer to the ideal ε_{it} than $g_{it} - \bar{g}_t$, as $\bar{g}_{i,t}$ may control better than \bar{g}_t for industry-wide disturbances, for examples, industry-wide real price movements.

Before that, I state a result that establishes sufficient conditions for identification.

PROPOSITION 4: *Suppose that (i) decomposition (30) holds with a vector of observables X_{it} and that (ii) $\sum_{i=1}^{\infty} (\frac{S_{i,t-1}}{Y_{t-1}})^2 \mathbb{E}[|X_{it}|^2] < \infty$. Then, as the number of firms becomes large (in K or in $Q \geq K$), $\Gamma_t(K, Q) - \Gamma_t^*(K) \rightarrow 0$ almost surely, that is, the empirical granular residual Γ_t is close to the ideal granular residual Γ_t^* .*

Assumption (i) is the substantial one. Given that in practice I will have X_{it} made of \bar{g}_t and $\bar{g}_{i,t}$, and their interaction with firm size, I effectively assume that the average growth rate of firms and their industries, perhaps interacted with the firm size or such nonlinear transformation of it, span the vector of factors. In other terms, firms within a given industry respond in the same way to common shocks or respond in a way that is related to firm size as in (36) below. This is the case under many models, but they are not fully general. Indeed, without some sort of parametric restriction, there is no solution (Manski (1993)). A typical problematic situation would be the case where the top firm has a high loading on industry factors that is not captured by its size. Then, instead of the large firms affecting the common factor, the factor would affect the large firms. However, I do control for size and the interaction between size and industry, and aggregate effects, so in that sense I can hope to be reasonably safe.¹⁷

Assumption (ii) is simply technical and is easily verified. For instance, it is verified if $\mathbb{E}[X_{it}^2]$ is finite and the herfindahl is bounded. Formally, the herfindahl (which, as we have seen, is small anyway) is bounded if the total sales to out-

¹⁷The above reflects my best attempt with Compustat data. Suppose one had continuous-time firm-level data and could measure the beginning of a strike, the launch of a new product, or the sales of a big export contract. These events would be firm-level shocks. It would presumably take some time to reverberate in the rest of the economy. Hence, a more precise understanding would be achieved. Perhaps future data (e.g., using newspapers to approximate continuous-time information) will be able to systematically achieve this extra measure of identification via the time series.

put ratio is bounded by some amount B , as $\sum_{i=1}^{\infty} (\frac{S_{i,t-1}}{Y_{t-1}})^2 \leq (\sum_{i=1}^{\infty} \frac{S_{i,t-1}}{Y_{t-1}})^2 \leq B^2$. Note that here we do not need to assume a finite number of firms, and that in practice $B \simeq 2$ (Jorgensen, Gollop, and Fraumeni (1987)).

To complete the econometric discussion, let me also mention a small sample bias: The R^2 measured by a regression will be lower than the true R^2 , because the control by \bar{g}_t effectively creates an error in variables problem. This effect, which can be rather large (and biases the results against the granular hypothesis), is detailed in the Supplemental Material (Gabaix (2011)).

I would like to conclude with a simple economic example that illustrates the basic granular residual (equation (33)).¹⁸ Suppose that the economy is made of one big firm which produces half of output, and a million other very small firms, and that I have good data on 100 firms: the big firm and the top 99 largest of the very small firms. The standard deviation of all growth rates is 10%, and growth rates are given by $g_{it} = X_t + \varepsilon_{it}$, where X_t is a common shock. Suppose that, in a given year, GDP increases by 3% and that the big firm has growth of, say, 6%, while the average of the small ones is close to 0%. What can we infer about the origins of shocks? If one thinks of all this being generated by an aggregate shock of 3%, then the distribution of implied idiosyncratic shocks is 3% for the big firm and -3% on average for all small ones. The probability that the average of the i.i.d. small firms is -3% , given the law of large numbers for these firms, is very small. Hence, it is more likely that the average shock X_t is around 0%, and the economy-wide growth of 3% comes from an idiosyncratic shock to the large firm equal to 6%. The estimate of the aggregate shock is captured by \bar{g}_t , which is close to 0%, and the estimate of the contribution of idiosyncratic shocks is captured by the granular residual, $\Gamma = 3\%$.

3.2. Empirical Implementation

3.2.1. Basic Specification

I use annual U.S. Compustat data from 1951 to 2008. For the granular residual, I take the $K = 100$ largest firms in Compustat according to the previous year's sales that have valid sales and employee data for both the current and previous years and that are not in the oil, energy, or finance sectors.¹⁹ Industries are three-digit Standard Industrial Classification (SIC) codes. Compustat contains some large outliers, which may result from extraordinary events,

¹⁸I thank Olivier Blanchard for this example.

¹⁹For firms in the oil/energy sector, the wild swings in worldwide energy prices make (29) too poor a proxy of total factor productivity. Likewise, the "sales" of financial firms do not mesh well with the meaning ("gross output") used in the present paper; this exclusion has little impact, though is theoretically cleaner.

TABLE I
EXPLANATORY POWER OF THE GRANULAR RESIDUAL^a

	GDP Growth _t		Solow _t	
(Intercept)	0.018** (0.0026)	0.017** (0.0025)	0.011** (0.002)	0.01** (0.0021)
Γ_t	1.8* (0.69)	2.5** (0.69)	2.1** (0.54)	2.3** (0.57)
Γ_{t-1}	2.6** (0.71)	2.9** (0.67)	1.2* (0.55)	1.3* (0.56)
Γ_{t-2}		2.1** (0.71)		0.65 (0.59)
N	56	55	56	55
R^2	0.266	0.382	0.261	0.281
Adj. R^2	0.239	0.346	0.233	0.239

^aFor the year $t = 1952$ to 2008 , per capita GDP growth and the Solow residual are regressed on the granular residual Γ_t of the top 100 firms (equation (33)). The firms are the largest by sales of the previous year. Standard errors are given in parentheses.

such as a merger. To handle these outliers, I winsorize the extreme demeaned growth rates at 20%.²⁰

Table I presents regressions of GDP growth and the Solow residual on the simplest granular residual (33). These regressions are supportive of the granular hypothesis. The R^2 's are reasonably high, at 34.6% for the GDP growth and around 23.9% for the Solow residual when using two lags. We will soon see that the industry-demeaned granular residual does even better.

If only aggregate shocks were important, then the R^2 of the regressions in Table I would be zero. Hence, the good explanatory power of the granular residual is inconsistent with a representative firm framework. It is also inconsistent with the hypothesis that most firm-level volatility might be due to a zero-sum redistribution of market shares.

Let us now examine the results if we incorporate a more fine-grained control for industry shocks.

3.2.2. Controlling for Industry Shocks

I next control for industry shocks, that is, use specification (34). Table II presents the results, which are consistent with those in Table I. The adjusted

²⁰For instance, I construct (32) by winsorizing $\widehat{\varepsilon}_{it}$ at $M = 20\%$, that is by replacing it by $T(\widehat{\varepsilon}_{it})$, where $T(x) = x$ if $|x| \leq M$, and $T(x) = \text{sign}(x)M$ if $|x| > M$. I use $M = 20\%$, but results are not materially sensitive to the choice of that threshold.

TABLE II
EXPLANATORY POWER OF THE GRANULAR RESIDUAL WITH
INDUSTRY DEMEANING^a

	GDP Growth _{<i>t</i>}		Solow _{<i>t</i>}	
(Intercept)	0.019** (0.0024)	0.017** (0.0022)	0.011** (0.0019)	0.011** (0.0019)
I_t	3.4** (0.86)	4.5** (0.82)	3.3** (0.68)	3.7** (0.72)
I_{t-1}	3.4** (0.82)	4.3** (0.78)	1.5* (0.65)	1.9** (0.68)
I_{t-2}		2.7** (0.79)		0.77 (0.69)
N	56	55	56	55
R^2	0.356	0.506	0.334	0.372
Adj. R^2	0.332	0.477	0.309	0.335

^aFor the year $t = 1952$ to 2008 , per capita GDP growth and the Solow residual are regressed on the granular residual I_t of the top 100 firms (equation (34)), removing the industry mean within this top 100. The firms are the largest by sales of the previous year. Standard errors are given in parentheses.

R^2 's are a bit higher: about 47.7% for GDP growth and 33.5% for the Solow residual when using two lags.²¹

This table reinforces the conclusion that idiosyncratic movements of the top 100 firms seem to explain a large fraction (about one-third, depending on the specification) of GDP fluctuations. In addition, industry controls, which may be preferable to a single aggregate control on a priori grounds, slightly strengthen the explanatory power of the granular residual.

In terms of economics, Tables I and II indicate that the lagged granular residual helps explain GDP growth, and that the same-year “multiplier” μ is around 3.

3.2.3. Predicting GDP Growth With the Granular Residual

The above regressions attempt to explain GDP with the granular residual, that is, relating aggregate movement to contemporary firm-level idiosyncratic movements that may be more easily understood (as we will see in the narrative below). I now study forecasting GDP growth with past variables. In addition to the granular residual, I consider the main traditional predictors. I control for oil and monetary policy shocks by following the work of Hamilton (2003) and Romer and Romer (2004), which are arguably the leading way to control for oil and monetary policy shocks. I also include the 3-month nominal T-bill and the

²¹The similarity of the results is not surprising, as the correlation between the simple and industry-demeaned granular residual is 0.82.

TABLE III
 PREDICTIVE POWER OF THE GRANULAR RESIDUAL FOR TERM SPREAD,
 OIL SHOCKS, AND MONETARY SHOCKS^a

	1	2	3	4	5	6	7	8
(Intercept)	0.022** (0.0029)	0.02** (0.0029)	0.022** (0.0029)	0.026** (0.0057)	0.015 (0.0075)	0.015 (0.0079)	0.019** (0.0027)	0.021** (0.0073)
Oil _{<i>t</i>-1}	-0.00027* (0.00012)		-0.00024* (0.00012)			-8.7e-05 (0.00013)		-0.00017 (0.00012)
Oil _{<i>t</i>-2}	-0.00018 (0.00012)		-0.00017 (0.00012)			-6.9e-05 (0.00012)		-0.00012 (0.00011)
Monetary _{<i>t</i>-1}		-0.083 (0.057)	-0.08 (0.055)			-0.042 (0.055)		-0.051 (0.05)
Monetary _{<i>t</i>-2}		-0.059 (0.057)	-0.038 (0.056)			-0.024 (0.054)		0.043 (0.053)
<i>r</i> _{<i>t</i>-1}				-0.75** (0.2)	-0.6 (0.32)	-0.45 (0.37)		-0.41 (0.34)
<i>r</i> _{<i>t</i>-2}				0.65** (0.19)	0.56 (0.32)	0.43 (0.37)		0.39 (0.34)
Term spread _{<i>t</i>-1}					0.32 (0.6)	0.38 (0.64)		0.4 (0.58)
Term spread _{<i>t</i>-2}					0.45 (0.47)	0.27 (0.54)		-0.38 (0.53)
<i>Γ</i> _{<i>t</i>-1}							3.5** (0.96)	3.3** (1)
<i>Γ</i> _{<i>t</i>-2}							1.2 (0.92)	2.3* (0.97)
<i>N</i>	55	55	55	55	55	55	55	55
<i>R</i> ²	0.121	0.0764	0.175	0.22	0.288	0.312	0.215	0.463
Adj. <i>R</i> ²	0.0871	0.0409	0.109	0.191	0.231	0.192	0.185	0.341

^aFor the year *t* = 1952 to 2008, per capita GDP growth is regressed on the lagged values of the granular residual *I*_{*t*} of the top 100 firms (equation (34)), of the Hamilton (for oil) and Romer–Romer (for money) shocks, and the term spread (the government 5-year bond yield minus the 3-month yield). We see that the granular residual has good incremental predictive power even beyond the term spread. Standard errors are given in parentheses.

term spread (which is defined as the 5-year bond rate minus the 3-month bond rate), which is often found to be the a very good predictor of GDP (those two endogenous variables are arguably more “diagnostic” than “causal,” though). Table III presents the results.

The granular residual has an adjusted *R*² (called \bar{R}^2) equal to 18.5% (column 7). The traditional economic factors—oil and money shocks—have an \bar{R}^2 of 10.9% (column 3). Past GDP growth has a very small \bar{R}^2 of -0.3%, a number not reported in Table III to avoid cluttering the table too much. The traditional diagnostic financial factors—the interest rate and the term spread—have an \bar{R}^2 of 23.1% (column 5). Putting all predictors together, the \bar{R}^2 is 34.1%

(column 8) and the granular residual brings an incremental \bar{R}^2 of 14.9% (compared to column 6).

I conclude that the granular residual is a new and apparently useful predictor of GDP. This result suggests that economists might use the granular residual to improve not only the understanding of GDP, but also its forecasting.

3.3. Robustness Checks

An objection to the granular residual is that the control for the common factors may be imperfect. Table IV shows the explanatory power of the granular residual, controlling for oil and monetary shocks. The adjusted \bar{R}^2 is 47.7% for the granular residual (column 4), it is 8.2% and 2.3% for oil and monetary shocks, respectively (columns 1 and 2), and 49.5% for financial variables (interest rates and term spread, column 6). To investigate whether the granular residual does add explanatory power, the last column puts all those variables together (perhaps pushing the believable limit of ordinary least squares (OLS) because of the large number of regressors) and shows that the explanatory variables yield an \bar{R}^2 of 76.7%.

In conclusion, as a matter of “explaining” (in a statistical sense) GDP growth, the granular residual does nearly as well as all traditional factors together, and complements their explanatory power.

I report a few robustness checks in the Supplemental Material. For instance, among the explanatory variables of (30), I include not only \bar{g}_t or $\bar{g}_{I,t}$, but also their interaction with log firm size and its square. The impact of the control for size is very small. Using a number $Q = 1000$ of firms yields similar results, too. Finally, I could not regress g_{it} on GDP growth at time t because then by construction I would eliminate any explanatory power of ε_{it} .

I conclude that the granular residual has a good explanatory power for GDP, even controlling for traditional factors. In addition, it has good forecasting power, complementing other factors. Hence, the granular residual must capture interesting firm-level dynamics that are not well captured by traditional aggregate factors.

I have done my best to obtain “idiosyncratic” shocks; given that I do not have a clean instrument, the above results should still be considered provisional. The situation is the analogue, with smaller stakes, to that of the Solow residual. Solow understood at the outset that there are very strong assumptions in the construction of his residual, in particular, full capacity utilization and no fixed cost. But a “purified” Solow residual took decades to construct (e.g., Basu, Fernald, and Kimball (2006)), requires much better data, is harder to replicate in other countries, and relies on special assumptions as well. Because of that, the Solow residual still endures, at least as a first pass. In the present paper too, it is good to have a first step in the granular residual, together with caveats that may help future research to construct a better residual. The conclusion of this article contains some other measures of granular residuals that build on the

TABLE IV
EXPLANATORY POWER OF THE GRANULAR RESIDUAL FOR
OIL AND MONETARY SHOCKS, AND INTEREST RATES^a

	1	2	3	4	5	6	7	8
(Intercept)	0.023** (0.003)	0.02** (0.0029)	0.022** (0.003)	0.017** (0.0022)	0.019** (0.0023)	0.016* (0.0065)	0.02** (0.005)	0.023** (0.0048)
Oil _{<i>t</i>}	-9.8e-05 (0.00011)		-8.3e-05 (0.00012)		-4.6e-05 (8.6e-05)			-7.9e-05 (7.5e-05)
Oil _{<i>t-1</i>}	-0.00028* (0.00012)		-0.00026* (0.00012)		-0.00021* (8.8e-05)			-0.00019* (7.5e-05)
Oil _{<i>t-2</i>}	-0.00019 (0.00012)		-0.00019 (0.00012)		-0.00012 (8.9e-05)			-4.3e-05 (6.8e-05)
Monetary _{<i>t</i>}		-0.0088 (0.059)	-0.03 (0.058)		-0.057 (0.043)			-0.044 (0.032)
Monetary _{<i>t-1</i>}		-0.08 (0.061)	-0.065 (0.059)		0.012 (0.047)			-0.013 (0.033)
Monetary _{<i>t-2</i>}		-0.061 (0.059)	-0.048 (0.058)		0.031 (0.046)			0.095** (0.033)
Γ_t				4.5** (0.82)	4.2** (0.88)		3.7** (0.69)	4** (0.66)
Γ_{t-1}				4.3** (0.78)	4.5** (0.85)		2.8** (0.71)	3.6** (0.68)
Γ_{t-2}				2.7** (0.79)	2.7** (0.8)		2.6** (0.69)	2.8** (0.63)
r_t						0.66* (0.26)	0.69** (0.2)	0.83** (0.19)
r_{t-1}						-1.6** (0.35)	-1.5** (0.28)	-1.5** (0.27)
r_{t-2}						1** (0.29)	0.85** (0.23)	0.7** (0.22)
Term spread _{<i>t</i>}						-0.49 (0.52)	-0.11 (0.41)	-0.13 (0.38)
Term spread _{<i>t-1</i>}						0.17 (0.52)	-0.34 (0.41)	-0.37 (0.42)
Term spread _{<i>t-2</i>}						0.31 (0.39)	-0.02 (0.32)	-0.18 (0.33)
<i>N</i>	55	55	55	55	55	55	55	55
<i>R</i> ²	0.133	0.0768	0.189	0.506	0.582	0.551	0.755	0.832
Adj. <i>R</i> ²	0.0824	0.0225	0.0878	0.477	0.498	0.495	0.706	0.767

^aFor the year $t = 1952$ to 2008 , per capita GDP growth is regressed on the granular residual Γ_t of the top 100 firms (equation (34)), and the contemporaneous and lagged values of the Hamilton (for oil) shocks, and Romer–Romer (for money) shocks. The firms are the largest by sales of the previous year. Standard errors are given in parentheses.

To obtain a manageable number of important episodes, I report the events with $|g_{Yt}| \geq 0.7\sigma_Y$, and in those years, report the firms for which $|F_{it}/g_{Yt}| \geq 0.14$. I also consider all the most extreme fifth of the years for F_t . I avoid, however, most points that are artefacts of mergers and acquisitions (more on that later). To avoid boring the reader with too many tales of car companies, I add a few non-car events that I found interesting economically or methodologically.

A general caveat is that the direction of the causality is hard to assess definitively, as the controls \bar{g}_{it} for industry-wide movements are imperfect. With that caveat in mind, we can start reading Table V.

To interpret the table, let me take a salient and relatively easy year, 1970. This year features a major strike at General Motors, which lasted 10 weeks (September 15 to November 20). The 1970 row of Table V shows that GM's sales fell by 31% and employment fell by 13%. Its labor productivity growth rate is thus -17.9% and, controlling for the industry mean productivity growth of 2.6% that year, GM's demeaned growth rate is -20.5% . Given that GM's sales the previous year were 2.47% of GDP, GM's granular residual is $F_{it} = -0.20 \times 2.47\% = -0.49\%$. That means the direct impact of this GM event is a change in GDP by -0.49% that year. Note also that with a productivity multiplier of $\mu \simeq 3$, the imputed impact of GM on GDP is -1.47% . As GDP growth that year was 3% below trend ($g_{Yt} = -3\%$), the direct share of the GM event is $0.49\%/3\% = 0.16$ and its full imputed share is $1.47\%/3\% = 0.49$. In some mechanical sense, the GM event appears to account for a fraction 0.17 of the GDP movement directly and, indirectly, for about 0.5 of the GDP innovation that year. It also accounts for a fraction 0.76 of the granular residual. Hence, it is plausible to interpret 1970 as a granular year, whose salient event was the GM strike and the turmoils around it.²³ This example shows how the table is organized. Let me now present the rest of the narrative.

1952–1953: U.S. Steel faces a strike from about April 1952 to August 1952. U.S. Steel's production falls by 13.1% in 1952 and rebounds by 19.5% in 1953. The 1953 events explains a share of 3.99 of the granular residual and 0.06 of excess GDP growth.

1955 experiences a high GDP growth, and a reasonably high granular residual. The likely microfoundation is a boom in car production. Two main specific factors seem to explain the car boom: the introduction of new models of cars and the fact that car companies engaged in a price war (Bresnahan (1987)). The car sales of GM increase by 21.9%, while employment increases by 7.9%. The demeaned growth rate is $g'_{it} = 17.8\%$. GM accounts for 81% of the gran-

²³Temin (1998) noted that the winding down of the Vietnam War (which ended in 1975) may also be responsible for the slump of 1970. This is in part the case, as during 1968–1972 the ratios of defense outlays to GDP were 9.5, 8.7, 8.1, 7.3, and 6.7%. On the other hand, the ratio of total government outlays to GDP were, respectively, 20.6, 19.4, 19.3, 19.5, and 19.6% (source: Council of Economic Advisors (2005, Table B-79)). Hence the aggregate government spending shock was very small in 1970.

TABLE V
NARRATIVE^a

Year	Firm	Share of GDP $\frac{S_{it,t-1}}{Y_{t-1}}$ in %	Labor Prod. Growth g_{it} [$\Delta \ln S_{it}, \Delta \ln L_{it}$]	Demeaned Growth $g_{it} - \bar{g}_{it}$ in %	Gran. Res. I_{it} in %	Share of GR $\frac{I_{it}}{Y_t}$	Direct Share of g_{Yt} $\frac{I_{it}}{g_{Yt}}$	Imputed Share of g_{Yt} $\frac{\mu I_{it}}{g_{Yt}}$	Brief Explanation
1952	U.S. Steel	1.03	-10.75 [-13.10, -2.35]	-3.56	-0.037	0.061	-0.810	-2.430	Strike
1953	U.S. Steel	0.87	17.06 [19.51, 2.45]	5.86	0.051	3.985	0.060	0.180	Rebound from strike
1955	GM	2.58	14.00 [21.89, 7.88]	17.84	0.461	0.808	0.142	0.426	Boom in car production: New models and price war
1956	Ford	1.35	-20.95 [-21.96, -1.01]	-20.72	-0.270	0.407	0.145	0.435	End of price war
1956	GM	3.00	-13.55 [-17.61, -4.06]	-13.32	-0.400	0.603	0.215	0.645	End of price war
1957	GM	2.47	0.36 [-1.50, -1.85]	-12.38	-0.305	2.201	0.167	0.501	End of price war (aftermath)
1961	Ford	1.00	25.12 [23.64, -1.48]	27.03	0.199	4.131	-0.147	-0.441	Success of compact Falcon (rebound from Edsel failure)
1965	GM	2.56	7.45 [18.06, 10.61]	11.10	0.284	0.600	0.092	0.276	Boom in new-car sales
1967	Ford	1.55	-19.84 [-18.23, 1.61]	-14.91	-0.232	2.461	0.379	1.137	Strike
1970	GM	2.47	-17.85 [-31.06, -13.20]	-20.52	-0.493	0.757	0.165	0.495	Strike
1971	GM	1.81	25.58 [36.15, 10.57]	23.35	0.361	0.516	7.344	22.032	Rebound from strike

(Continues)

TABLE V—Continued

Year	Firm	Share of GDP $\frac{S_{it,t-1}}{Y_{t-1}}$ in %	Labor Prod. Growth g_{it} [$\Delta \ln S_{it}, \Delta \ln L_{it}$]	Demeaned Growth $g_{it} - \bar{g}_{it}$ in %	Gran. Res. I_{it} in %	Share of GR $\frac{I_{it}}{I_t}$	Direct Share of g_{Yt} $\frac{I_{it}}{g_{Yt}}$	Imputed Share of g_{Yt} $\frac{\mu I_{it}}{g_{Yt}}$	Brief Explanation
1972	Chrysler	0.71	16.76 [15.64, -1.13]	17.80	0.126	0.234	0.058	0.174	Rush of sales for subcompacts (Dodge Dart and Plymouth Valiant)
1972	Ford	1.46	14.18 [16.36, 2.18]	15.22	0.222	0.411	0.103	0.309	Rush of sales for subcompacts (Ford Pinto)
1974	GM	2.59	-11.31 [-21.28, -9.97]	-15.23	-0.394	0.913	0.115	0.345	Cars with poor gas mileage hit by higher oil price
1983	IBM ^b	1.06	10.46 [11.76, 1.29]	10.52	0.111	0.177	0.071	0.213	Launch of the IBM PC
1987	GE ^b	0.79	25.62 [8.33, -17.29]	21.46	0.158	1.110	0.357	1.071	Moving out of manufacturing and into finance and high-tech
1988	GE ^b	0.83	21.42 [20.08, -1.33]	16.55	0.137	0.441	0.117	0.351	Moving out of manufacturing and into finance and high-tech
1996	AT&T	1.08	38.97 [-44.11, -83.08]	32.45	0.215	0.471	0.446	1.338	Spin-off of NCR and Lucent
2000	GE	1.20	20.56 [12.29, -8.27]	33.04	0.239	9.934	0.468	1.404	Sales topped \$111bn, expansion of GE Medical Systems
2002	Walmart	2.16	8.61 [9.83, 1.22]	6.39	0.138	3.219	-0.099	-0.297	Success of lean distribution model

^aGE and GM are General Electric and General Motors, respectively. For each firm i , g_{it} , $\Delta \ln S_{it}$, and $\Delta \ln L_{it}$ denote productivity, sales, and employment growth rates, respectively, $g_{it} - \bar{g}_{it}$ denotes industry-demeaned growth, and $S_{it,t-1}/Y_{t-1}$ is the sales share of GDP. The firm granular residual is $I_{it} = \frac{S_{it,t-1}(g_{it} - \bar{g}_{it})}{Y_{t-1}}$, and I_{it}/I_t is the respective share of the granular residual. I_{it}/g_{Yt} is the direct share of the firm shock on demeaned GDP growth. The full share would be equal to $\mu I_{it}/g_{Yt}$, where $\mu = 3$ is the typical productivity multiplier estimated from Tables I and II.

^bThere is just one firm in this industry in the top 100, hence \bar{g}_{it} was replaced by \bar{g}_t .

ular residual, a direct fraction 0.14 of excess GDP growth, and an imputed fraction of 0.43 of excess GDP growth.

1956–1957: In 1956, the price war in cars ends, and sales drop back to their normal level (the sales of General Motors decline by 17.6%; those of Ford decline by 22%). The granular residual is -0.66% , of which 60% is due to General Motors. Hence, one may provisionally conclude the 1955–1956 boom–bust episode was in large part a granular event driven by new models and a price war in the car industry.²⁴ In Figure 2, the 56 point is actually the sum of 1955 (granular boom) and 1956 (granular bust), and is unremarkable, but the bust is reflected in the 1957 point, which is the most extreme negative point in the granular residual.

1961: In previous years, Ford cancelled the Edsel brand and introduces to great success the Falcon, the leading compact car of its time. Ford's demeaned growth rate is $g'_{it} = 27\%$ and its firm granular residual explains a fraction -0.15 of excess GDP growth. That is, without Ford's success, the recession would have been worse.

1965 is an excellent year for GM, with the great popularity of its Chevrolet brand.

1967: Ford experiences a 64-day strike and a terrible year. Its demeaned growth rate is -14.9% and its granular residual is -0.23% . It explains a fraction 2.5 of the granular residual and 0.38 of GDP growth.

1970 is the GM year described above.

1971, which appears in Figure 2 as label “72,” representing the sum of the granular residuals in 1971 and 1972, is largely characterized by the rebound from the negative granular 1970 shock. Hence, the General Motors strike may explain the very negative 70 (1969 + 1970) point and the very positive 72 (1971 + 1972) point. Sales increase by 36.2% and employment increases by 10.6%. The firm granular residual is $\Gamma_{it} = 0.36\%$ for a fraction of granular residual of 0.52. Another interesting granular event takes place in 1971. The Council of Economic Advisors (1972, p. 33) reports that “prospects of a possible steel strike after July 31st [1971], the expiration day of the labor contracts, caused steel consumers to build up stock in the first seven months of 71, after which these inventories were liquidated.” Here, a granular shock—the possibility of a steel strike—creates a large swing in inventories. Without exploring inventories here, one notes that such a plausibly orthogonal inventory shock could be used in future macroeconomic studies.

1972 is a very good year for Ford and Chrysler. Ford has an enormous success with its Pinto. At Chrysler, there is a rush of sales for the compact Dodge Dart and Plymouth Valiant (low-priced subcompacts). For those two firms, $\Gamma_{it} = 0.22\%$ and $\Gamma_{it} = 0.13\%$, respectively.

1974 is probably not a granular year, because the oil shock was common to many industries. Still, the low value of the granular residual reflects the

²⁴To completely resolve the matter, one would like to control for the effect of the Korean war.

fact that the top three car companies, and particularly General Motors, were disproportionately affected by the shock. It is likely that if large companies were producing more fuel efficient cars, the granular residual would have been closer to 0, and the slump of 1974 could have been much more moderate. For instance, GM's granular contribution is -0.39% , and its multiplier-adjusted contribution -1.18% .

1983 is an excellent year for IBM, with the launch of the IBM PC. Its $g_{it} = 10.5\%$, so that its granular residual is 0.11% .

1987–1988 is an instructive year, in part for methodological reasons. After various investments and mergers and acquisitions in 1986–1987 (acquisition of financial services providers, e.g., KidderPeabody, and high-tech companies such as medical diagnostics business), the clear majority of GE's earnings (roughly 80% , compared to 50% 6 years earlier) were generated in services and high technology. Its g_{it} is 26% and 21% in 1987 and 1988, respectively. Its fraction of the granular residual is 1.11 and 0.44 , and its imputed growth fraction is 1.07 and 0.35 . This episode can be viewed either as a purely formal reallocation of titles in economic activity (in which case it arguably should be discarded) or as a movement of "structural change" where this premier firm's efforts (human and physical) are reallocated toward higher value-added activities, thereby potentially increasing economic activity.²⁵ The same can be said about the next event.

1996: There is an intense restructuring at AT&T, with a spin-off of NCR and Lucent. AT&T recenters to higher productivity activities, and as a result its measured g'_{it} is 32.5% . This movement explains a fraction 0.47 of the granular residual and 0.45 of GDP growth.

2000 is a year of great productivity growth for GE, in particular via the expansion of GE Medical Systems. Its g_{it} is 20.6% and its firm granular residual is $\Gamma_{it} = 0.24\%$.

2002 sees a surge in sales for Walmart, a vindication of its lean distribution model. The company's share of the U.S. GDP in 2001 was 2.2% . This approached the levels reached by GM (3% in 1956) and U.S. Steel Corp. (2.8% in 1917) when these firms were at their respective peaks. Its g'_{it} is 6.4% and its fraction of the granular residual is 3.22 , while its fraction of demeaned GDP growth is -0.10 .

We arrive at the limen of the financial crisis. 2007 sees three interesting granular events (not reported in the table) if one is willing to accept the "sales" of financial firms as face value (it is unclear they should be). The labor productivity growth of AIG, Citigroup, and Merrill Lynch is -15% , -9% , and -25% , respectively, which gives them granular contributions of -0.09% , -0.18% ,

²⁵Under the first interpretation, it would be interesting to build a more "purified" granular residual that filters out corporate finance events. Of course, to what extent those events should be filtered out is debatable.

and -0.10% . It would be interesting to exploit the hypothesis that the financial crisis was largely caused by the (ex post) mistakes of a few large firms, e.g., Lehman and AIG. Their large leverage and interconnectedness amplified into a full-fledged crisis instead of what could have been a run-of-the-mill that would have affected in a diffuse way the financial sector. But doing justice to this issue would require another paper.

Figure 2 reveals that, in the 1990's, granular shocks are smaller. Likewise, GDP volatility is smaller—reflecting the “great moderation” explored in the literature (e.g., McConnell and Perez-Quiros (2000)). Carvalho and Gabaix (2010) explored this link in more depth, and proposed that indeed the decrease in granular volatility explains the great moderation of GDP and its demise.

Finally, the bottom of Figure 2 contains three outliers that are not granular years. They have conventional “macro” interpretations. 1954 is often attributed to the end of the Korean War, and 1958 and 1982 (the “Volcker recession”) are attributed to tight monetary policy aimed at fighting inflation.

This narrative shows the importance of two types of events: some (e.g., a strike) inherently have a negative autocorrelation, while some others (e.g., new models of cars) do not. It is conceivable that forecasting could be improved by taking into account that distinction.

4. CONCLUSION

This paper shows that the forces of randomness at the microlevel create an inexorable amount of volatility at the macro level. Because of random growth at the microlevel, the distribution of firm sizes is very fat-tailed (Simon (1955), Gabaix (1999), Luttmer (2007)). That fat-tailedness makes the central limit theorem break down, and idiosyncratic shocks to large firms (or, more generally, to large subunits in the economy such as family business groups or sectors) affect aggregate outcomes.

This paper illustrates this effect by taking the example of GDP fluctuations. It argues that idiosyncratic shocks to the top 100 firms explain a large fraction (one-third) of aggregate volatility. While aggregate fluctuations such as changes to monetary, fiscal, and exchange rate policy, and aggregate productivity shocks are clearly important drivers of macroeconomic activity, they are not the only contributors to GDP fluctuations. Using theory, calibration, and direct empirical evidence, this paper makes the case that idiosyncratic shocks are an important, and possibly the major, part of the origin of business-cycle fluctuations.

The importance of idiosyncratic shocks in aggregate volatility leads to a number of implications and directions for future research. First, and most evidently, to understand the origins of fluctuations better, one should not focus exclusively on aggregate shocks, but concrete shocks to large players, such as General Motors, IBM, or Nokia.

Second, shocks to large firms (such as a strike, a new innovation, or a CEO change), initially independent of the rest of the economy, offer a rich source of

shocks for vector autoregressions (VARs) and impulse response studies—the real-side equivalent of the Romer and Romer shocks for monetary economics. As a preliminary step in this direction, the granular residual, with a variety of specifications, is available in the Supplemental Material.

Third, this paper gives a new theoretical angle for the propagation of fluctuations. If Apple or Walmart innovates, its competitors may suffer in the short term and thus race to catch up. This creates rich industry-level dynamics (that are already actively studied in the industrial organization literature) that should be useful for studying macroeconomic fluctuations, since they allow one to trace the dynamics of productivity shocks.

Fourth, this argument could explain the reason why people, in practice, do not know the state of the economy. This is because “the state of the economy” depends on the behavior (productivity and investment behavior, among others) of many large and interdependent firms. Thus, the integration is not easy and no readily accessible single number can summarize this state. This contrasts with aggregate measures, such as GDP, which are easily observable. Conversely, agents who focus on aggregate measures may make potentially problematic inferences (see Angeletos and La’O (2010) and Veldkamp and Wolfers (2007) for research along those lines). This paper could therefore offer a new mechanism for the dynamics of “animal spirits.”

Finally this mechanism might explain a large part of the volatility of many aggregate quantities other than output, for instance, inventories, inflation, short- or long-run movements in productivity, and the current account. Fluctuations of exports due to granular effects are explored in Canals et al. (2007) and di Giovanni and Levchenko (2009). The latter paper in particular finds that lowering trade barriers increases the granularity of the economy (as the most productive firms are selected) and implies an increase in the volatility of exports. Blank, Buch, and Neugebauer (2009) constructed a “banking granular residual” and found that negative shocks to large banks negatively impact small banks. Malevergne, Santa-Clara, and Sornette (2009) showed that the granular residual of stock returns (the return of a large firm minus a return of the average firm) is an important priced factor in the stock market and explained the performance of Fama–French factor models. Carvalho and Gabaix (2010) found that the time-series changes in granular volatility predict well the volatility of GDP, including the “great moderation” and its demise.

In sum, this paper suggests that the study of very large firms can offer a useful angle of attack on some open issues in macroeconomics.

APPENDIX A: LÉVY’S THEOREM

Lévy’s theorem (Durrett (1996, p. 153)) is the counterpart of the central limit theorem for infinite-variance variables.

THEOREM 1—Lévy’s Theorem: *Suppose that X_1, X_2, \dots are i.i.d. with a distribution that satisfies (i) $\lim_{x \rightarrow \infty} \mathbb{P}(X_1 > x) / \mathbb{P}(|X_1| > x) = \theta \in [0, 1]$ and*

(ii) $\mathbb{P}(|X_1| > x) = x^{-\zeta}L(x)$, with $\zeta \in (0, 2)$ and $L(x)$ slowly varying.²⁶ Let $s_n = \sum_{i=1}^n X_i$, $a_n = \inf\{x : \mathbb{P}(|X_1| > x) \leq 1/n\}$, and $b_n = nE[X_1 1_{|X_1| \leq a_n}]$. As $n \rightarrow \infty$, $(s_n - b_n)/a_n$ converges in distribution to a nondegenerate random variable Y , which follows a Lévy distribution with exponent ζ .

The most typical use of Lévy’s theorem is the case of a symmetrical distribution with zero mean and power-law distributed tails, $\mathbb{P}(|X_1| > x) \sim (x/x_0)^{-\zeta}$. Then $a_n \sim x_0 n^{1/\zeta}$, $b_n = 0$, and $(x_0 N^{1/\zeta})^{-1} \sum_{i=1}^N X_i \xrightarrow{d} Y$, where Y follows a Lévy distribution. The sum $\sum_{i=1}^N X_i$ does not scale as $N^{1/2}$, as it does in the central limit theorem, but it scales as $N^{1/\zeta}$. This is because the size of the largest units X_i scales as $N^{1/\zeta}$.

A symmetrical Lévy distribution with exponent $\zeta \in (0, 2]$ has the distribution $\lambda(x, \zeta) = \frac{1}{\pi} \int_0^\infty e^{-kx} \cos(kx) dk$. For $\zeta = 2$, a Lévy distribution is Gaussian. For $\zeta < 2$, the distribution has a power-law tail with exponent ζ .

APPENDIX B: LONGER DERIVATIONS

PROOF OF PROPOSITION 3: To simplify notations, using homogeneity, I consider the case $\bar{\sigma} = S_{\min} = 1$. As $\Delta S_i/S_i = S_i^{-\alpha} u_i$,

$$(35) \quad \frac{\Delta Y}{Y} = \frac{\sum_{i=1}^N \Delta S_i}{Y} = \frac{\sum_{i=1}^N S_i^{1-\alpha} u_i}{\sum_{i=1}^N S_i}.$$

When $\zeta > 1$, by the law of large numbers, $N^{-1}Y = N^{-1} \sum_{i=1}^N S_i \rightarrow \bar{S}$. To tackle the numerator, I observe that $S_i^{1-\alpha}$ has power-law tails with exponent $\zeta' = \zeta/(1 - \alpha)$. I consider two cases. First, if $\zeta' < 2$, define $x_i = S_i^{1-\alpha} u_i$, which has power-law tails with exponent ζ' , and prefactor given by, for $x \rightarrow \infty$,

$$\begin{aligned} \mathbb{P}(|S_i^{1-\alpha} u_i| > x) &= \mathbb{P}\left(S_i > \left(\frac{x}{|u_i|}\right)^{1/(1-\alpha)}\right) \\ &\sim \mathbb{E}[|u_i|^{\zeta/(1-\alpha)}] x^{-\zeta/(1-\alpha)} = \mathbb{E}[|u_i|^{\zeta'}] x^{-\zeta'}. \end{aligned}$$

Hence in Lévy’s theorem, the a_N factor is $a_N \sim N^{1/\zeta'} \mathbb{E}[|u_i|^{\zeta'}]^{1/\zeta'}$ and

$$\Delta Y = \sum_{i=1}^N S_i^{1-\alpha} u_i \sim N^{1/\zeta'} \mathbb{E}[|u_i|^{\zeta'}]^{1/\zeta'} g_{\zeta'},$$

²⁶ $L(x)$ is said to be slowly varying if $\forall t > 0, \lim_{x \rightarrow \infty} L(tx)/L(x) = 1$. Prototypical examples are $L(x) = c$ and $L(x) = c \ln x$ for a nonzero constant c .

where $g_{\zeta'}$ is a Lévy with exponent ζ' . Hence, given $\bar{S} = \mathbb{E}[S] = S_{\min}\zeta/(\zeta - 1)$,

$$\frac{\Delta Y}{Y} \sim \frac{\mathbb{E}[|u_i|^{\zeta'}]^{1/\zeta'} N^{1/\zeta'}}{\bar{S}N} g_{\zeta'} = \frac{\zeta - 1}{\zeta} \mathbb{E}[|u_i|^{\zeta'}]^{1/\zeta'} N^{1/\zeta' - 1} g_{\zeta'}.$$

If $\zeta' = 2$, $S_i\sigma(S_i)u_i$ has finite variance equal to $\mathbb{E}[S^2\sigma(S)^2]\mathbb{E}[u_i^2]$. By the central limit theorem, $\Delta Y \sim \sqrt{N}\mathbb{E}[S^2\sigma(S)^2]^{1/2}\mathbb{E}[u_i^2]g_2$, where g_2 is a standard Gaussian distribution and

$$\frac{\Delta Y}{Y} \sim \frac{\Delta Y}{N\bar{S}} \sim \frac{1}{\sqrt{N}} \frac{\mathbb{E}[S^2\sigma(S)^2]^{1/2}\mathbb{E}[u_i^2]^{1/2}}{\bar{S}} g_2$$

as announced. When $\zeta = 1$, a $\ln N$ correction appears because of (14), but the reasoning is otherwise the same. Q.E.D.

Hulten’s Theorem With and Without Instantaneous Reallocation of Factors

For clarity, I rederive and extend Hulten’s (1978) result, which says that a Hicks-neutral productivity shock $d\pi_i$ to firm i causes an increase in TFP given by equation (15) (see also Jones (2011) for consequences of this result). There are N firms. Firm i produces good i and uses a quantity X_{ij} of intermediary inputs from firm j . It also uses L_i units of labor and K_i units of capital. It has productivity π_i . If production is: $Q_i = e^{\pi_i} F^i(X_{i1}, \dots, X_{iN}, L_i, K_i)$, the representative agent consumes C_i of good i and has a utility function of $U(C_1, \dots, C_N)$. Production of firm i serves as consumption and intermediary inputs, so $Q_i = C_i + \sum_k X_{ki}$. The optimum in this economy reads $\max_{C_i, X_{ik}, L_i, K_i} U(C_1, \dots, C_N)$ subject to

$$C_i + \sum_k X_{ki} = e^{\pi_i} F^i(X_{i1}, \dots, X_{iN}, L_i, K_i);$$

$$\sum_i L_i = L; \quad \sum_i K_i = K.$$

The Lagrangian is

$$\begin{aligned} W = & U(C_1, \dots, C_N) \\ & + \sum_i p_i \left[e^{\pi_i} F^i(X_{i1}, \dots, X_{iN}, L_i, K_i) - C_i - \sum_k X_{ki} \right] \\ & + w \left[L - \sum_i L_i \right] + r \left[K - \sum_i K_i \right]. \end{aligned}$$

Assume marginal cost pricing. GDP in this economy is $Y = wL + rK = \sum_i p_i C_i$. The value added of firm i is $wL_i + rK_i$ and its sales are $p_i Q_i$.

If each firm i has a shock $d\pi_i$ to productivity, welfare changes as

$$\begin{aligned} dW &= \sum_i p_i [e^{\pi_i} F^i(X_{i1}, \dots, X_{iN}, L_i, K_i) d\pi_i] \\ &= \sum_i (\text{sales of firm } i) d\pi_i. \end{aligned}$$

As dW can be rewritten $dW = \sum_i \frac{\partial U}{\partial C_i} dC_i = \sum_i p_i dC_i = dY$, we obtain equation (15).

The above analysis shows that Hulten's theorem holds even if, after the shock, the capital, labor and material inputs are not reallocated. This is a simple consequence of the envelope theorem. Hence Hulten's result also holds if there are frictions in the adjustment of labor, capital, or intermediate inputs.

PROOF OF PROPOSITION 4: We have

$$\begin{aligned} \Gamma_t(K, Q) - \Gamma_t^*(K) &= \sum_{i=1}^K \frac{S_{i,t-1}}{Y_{t-1}} (\widehat{\varepsilon}_{it} - \varepsilon_{it}) \\ &= \sum_{i=1}^K \frac{S_{i,t-1}}{Y_t} ((g_{it} - \widehat{\beta}' X_{it}) - (g_{it} - \beta' X_{it})) \\ &= (\beta - \widehat{\beta})' \sum_{i=1}^K \frac{S_{i,t-1}}{Y_t} X_{it}. \end{aligned}$$

We have $\beta - \widehat{\beta} \rightarrow 0$ almost surely (a.s.) when K or $Q \rightarrow \infty$ (by standard results on OLS), and $\sum_{i=1}^K \frac{S_{i,t-1}}{Y_t} X_{it} \rightarrow \sum_{i=1}^{\infty} \frac{S_{i,t-1}}{Y_t} X_{it}$ in the L^2 sense by assumption (ii). Hence $\Gamma_t(K, Q) - \Gamma_t^*(K, Q) \rightarrow 0$ a.s. *Q.E.D.*

APPENDIX C: DATA APPENDIX

Firm-Level Data

The firm-level data come from the Fundamentals Annual section of the Compustat North America database on WRDS. The data consist of year by firm observations from 1950 to 2008 of the following variables: SIC code (SIC), net sales in \$MM (DATA 12), and employees in M (DATA 29). I exclude foreign firms based in the United States, restricting the data set to firms whose fic and loc codes are equal to USA. I filter out oil and oil-related companies (SIC codes 2911, 5172, 1311, 4922, 4923, 4924, and 1389), and energy companies (SIC code between 4900 and 4940), as fluctuations of their sales come mostly from worldwide commodity prices, rather than real productivity shocks, and

financial firms (SIC code between 6000 and 6999), because their sales do not mesh well with the meaning used in the present paper.²⁷

An important caveat is in order for U.S. firms. With Compustat, the sales of Ford, say, represent the worldwide sales of Ford, not directly the output produced by Ford in the United States. There is no simple solution to this problem, especially if one wants a long time series. An important task of future research is to provide a version of Compustat that corrects for multinationals.

Macroeconomic Data

The real GDP, GDP per capita, and inflation index data series all come from the Bureau of Economic Analysis. The Solow residual is the multifactor productivity of the private business sector reported by the Bureau of Labor Studies. The term spread and real interest rate data are from the Fama–Bliss Center for Research in Security Prices (CRSP) Monthly Treasury data base on WRDS.

The data for the Romer and Romer (2004) monetary policy shocks come from David Romer’s web page. Their original series (RESID) is monthly from 1969 to 1996. Here the yearly Romer–Romer shock is the sum of the 12 monthly shocks in that year. For the years not covered by Romer and Romer, the value of the shock is assigned to be 0, the mean of the original data. This assignment does not bias the regression coefficient under simple conditions, for instance if the data are i.i.d. It does lower the R^2 by the fraction of missed variance; fortunately, most large monetary shocks (e.g., of the 1970’s and 1980’s) are in the data set.

The data for the Hamilton (2003) oil shocks primarily come from James Hamilton’s web page. This series is quarterly and runs until 2001. It is defined as the amount by which the current oil price exceeds the maximum value over the past year. This paper’s yearly shock is the sum of the quarterly Hamilton shocks. The spot price for oil reported by the St. Louis Federal Reserve is used to extend the series to the present.

REFERENCES

- ABRAMOVITZ, M. (1956): “Resource and Output Trends in the United States Since 1870,” *American Economic Review*, 46, 5–23. [735]
- ACEMOGLU, D., S. JOHNSON, AND T. MITTON (2009): “Determinants of Vertical Integration: Financial Development and Contracting Costs,” *Journal of Finance*, 64, 1251–1290. [746]
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2010): “Cascades in Networks and Aggregate Volatility,” Working Paper, MIT. [744]
- ANGELETOS, G. M., AND J. LA’O (2010): “Noisy Business Cycles,” in *NBER Macroeconomics Annual*. Cambridge: MIT Press. [765]
- AXTELL, R. (2001): “Zipf Distribution of U.S. Firm Sizes,” *Science*, 293, 1818–1820. [739]

²⁷The results with financial firms are very similar.

- BAK, P., K. CHEN, J. SCHEINKMAN, AND M. WOODFORD (1993): "Aggregate Fluctuations From Independent Sectoral Shocks: Self-Organized Criticality in a Model of Production and Inventory Dynamics," *Ricerche Economiche*, 47, 3–30. [736,737]
- BASU, S., J. FERNALD, AND M. KIMBALL (2006): "Are Technology Improvements Contractionary?" *American Economic Review*, 96, 1418–1448. [756]
- BLANK, S., C. BUCH, AND K. NEUGEBAUER (2009): "Shocks at Large Banks and Banking Sector Distress: The Banking Granular Residual," Discussion Paper 04/2009, Deutsche Bundesbank. [765]
- BRESNAHAN, T. (1987): "Competition and Collusion in the American Automobile Market: The 1955 Price War," *Journal of Industrial Economics*, 35, 457–482. [759]
- CABALLERO, R., E. M. R. A. ENGEL, AND J. C. HALTIWANGER (1997): "Aggregate Employment Dynamics: Building From Microeconomic Evidence," *American Economic Review*, 87, 115–137. [737,745]
- CANALS, C., X. GABAIX, J. VILARRUBIA, AND D. WEINSTEIN (2007): "Trade Patterns, Trade Balances, and Idiosyncratic Shocks," Working Paper, Columbia University. [734,765]
- CARVALHO, V. (2009): "Aggregate Fluctuations and the Network Structure of Intersectoral Trade," Working Paper, CREI. [744]
- CARVALHO, V., AND X. GABAIX (2010): "The Great Diversification and Its Undoing," Working Paper 16424, NBER. [749,764,765]
- CLAESSENS, S., S. DJANKOV, AND L. H. P. LANG (2000): "The Separation of Ownership and Control in East Asian Corporations," *Journal of Financial Economics*, 58, 81–112. [737]
- COCHRANE, J. (1994): "Shocks," *Carnegie-Rochester Conference Series on Public Policy*, 41, 295–364. [733]
- COMIN, D., AND S. MULANI (2006): "Diverging Trends in Macro and Micro Volatility: Facts," *Review of Economics and Statistics*, 88, 374–383. [745]
- CONLEY, T., AND B. DUPOR (2003): "A Spatial Analysis of Sectoral Complementarity," *Journal of Political Economy*, 111, 311–352. [737]
- COUNCIL OF ECONOMIC ADVISORS (1972): *Economic Report of the President Transmitted to the Congress*. Washington, DC: U.S. Government Printing Office. [762]
- (2005): *Economic Report of the President Transmitted to the Congress*. Washington, DC: U.S. Government Printing Office. [759]
- DAVIS, S., J. C. HALTIWANGER, AND S. SCHUH (1996): *Job Creation and Destruction*. Cambridge, MA: MIT Press. [745]
- DI GIOVANNI, J., AND A. LEVCHENKO (2009): "International Trade and Aggregate Fluctuations in Granular Economies," Working Paper, University of Michigan. [734,765]
- DUPOR, W. (1999): "Aggregation and Irrelevance in Multi-Sector Models," *Journal of Monetary Economics*, 43, 391–409. [737,743]
- DURLAUF, S. (1993): "Non Ergodic Economic Growth," *Review of Economic Studies*, 60, 349–366. [736]
- DURRETT, R. (1996): *Probability: Theory and Examples*. Belmont, CA: Wadsworth. [765]
- FACCIO, M., AND L. H. P. LANG (2002): "The Ultimate Ownership of Western European Corporations," *Journal of Financial Economics*, 65, 365–395. [737]
- FOERSTER, A., P.-D. SARTE, AND M. WATSON (2008): "Sectoral vs. Aggregate Shocks: A Structural Factor Analysis of Industrial Production," Working Paper 08-07, Federal Reserve Bank of Richmond. [758]
- GABAIX, X. (1999): "Zipf's Law for Cities: An Explanation," *Quarterly Journal of Economics*, 114, 739–767. [764]
- (2009): "Power Laws in Economics and Finance," *Annual Review of Economics*, 1, 255–293. [740]
- (2011): "Supplement to 'The Granular Origins of Aggregate Fluctuations'," *Econometrica Supplemental Material*, 79, http://www.econometricsociety.org/ecta/Supmat/8769_extensions.pdf; http://www.econometricsociety.org/ecta/Supmat/8769_data and programs. zip.zip. [752]

- HALL, R. (2009): "By How Much Does GDP Rise If the Government Buys More Output?" *Brookings Papers on Economic Activity*, 2, 183–231. [746]
- HAMILTON, J. D. (2003): "What Is an Oil Shock?" *Journal of Econometrics*, 113, 363–398. [754, 769]
- HORVATH, M. (1998): "Cyclicality and Sectoral Linkages: Aggregate Fluctuations From Sectoral Shocks," *Review of Economic Dynamics*, 1, 781–808. [737,743]
- (2000): "Sectoral Shocks and Aggregate Fluctuations," *Journal of Monetary Economics*, 45, 69–106. [737,743]
- HULTEN, C. (1978): "Growth Accounting With Intermediary Inputs," *Review of Economic Studies*, 45, 511–518. [744,767]
- JONES, C. I. (2011): "Intermediate Goods and Weak Links: A Theory of Economic Development," *American Economic Journal: Macroeconomics* (forthcoming). [767]
- JORGENSEN, D. W., F. M. GOLLOP, AND B. M. FRAUMENI (1987): *Productivity and U.S. Economic Growth*. Cambridge, MA: Harvard University Press. [752]
- JOVANOVIC, B. (1987): "Micro Shocks and Aggregate Risk," *Quarterly Journal of Economics*, 102, 395–409. [736]
- KOREN, M., AND S. TENREYRO (2007): "Volatility and Development," *Quarterly Journal of Economics*, 122, 243–287. [749]
- KYDLAND, F. E., AND E. C. PRESCOTT (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, 50, 1345–1370. [735]
- LEE, Y., L. A. N. AMARAL, M. MEYER, D. CANNING, AND H. E. STANLEY (1998): "Universal Features in the Growth Dynamics of Complex Organizations," *Physical Review Letters*, 81, 3275–3278. [748,749]
- LONG, J., AND C. PLOSSER (1983): "Real Business Cycles," *Journal of Political Economy*, 91, 39–69. [737,749]
- LUTTMER, E. G. J. (2007): "Selection, Growth, and the Size Distribution of Firms," *Quarterly Journal of Economics*, 122, 1103–1144. [764]
- MALEVERGNE, Y., P. SANTA-CLARA, AND D. SORNETTE (2009): "Professor Zipf Goes to Wall Street," Working Paper 15295, NBER. [765]
- MANSKI, C. F. (1993): "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 60, 531–542. [749,751]
- MCCONNELL, M., AND G. PEREZ-QUIROS (2000): "Output Fluctuations in the United States: What Has Changed Since the Early 1980's," *American Economic Review*, 90, 1464–1476. [764]
- NIREI, M. (2006): "Threshold Behavior and Aggregate Critical Fluctuations," *Journal of Economic Theory*, 127, 309–322. [736]
- OECD (2004): *Economic Survey of Finland*, Issue 14. Paris: OECD Publishing. [733]
- ROMER, C., AND D. ROMER (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94, 1055–1084. [754,769]
- SHEA, J. (2002): "Complementarities and Comovements," *Journal of Money, Credit and Banking*, 42, 412–433. [749]
- SIMON, H. (1955): "On a Class of Skew Distribution Functions," *Biometrika*, 42, 425–440. [764]
- SORNETTE, D. (2006): *Critical Phenomena in Natural Sciences*. New York: Springer. [740]
- STANLEY, M. H. R., L. A. N. AMARAL, S. V. BULDYREV, S. HAVLIN, H. LESCHHORN, P. MAASS, M. A. SALINGER, AND H. E. STANLEY (1996): "Scaling Behaviour in the Growth of Companies," *Nature*, 379, 804–806. [747]
- STOCK, J., AND M. WATSON (2002): "Forecasting Using Principal Components From a Large Number of Predictors," *Journal of the American Statistical Association*, 97, 1167–1179. [758]
- SUTTON, J. (2002): "The Variance of Firm Growth Rates: The 'Scaling' Puzzle," *Physica A*, 312, 577–590. [748]
- TEMIN, P. (1998): "The Causes of American Business Cycles: An Essay in Economic Historiography," in *Beyond Shocks*, ed. by J. C. Fuhrer and S. Schuh. Boston: Federal Reserve Bank. [758,759]

- VELDKAMP, L., AND J. WOLFERS (2007): "Aggregate Shocks or Aggregate Information? Costly Information and Business Cycle Comovement," *Journal of Monetary Economics*, 54, 37–55. [765]
- WYART, M., AND J.-P. BOUCHAUD (2003): "Statistical Models for Company Growth," *Physica A*, 326, 241–255. [748]
- ZIPF, G. (1949): *Human Behavior and the Principle of Least Effort*. Cambridge, MA: Addison-Wesley. [740]

Stern School of Finance, New York University, 44 West 4th Street, Suite 9-190, New York, NY 10012, U.S.A. and CEPR and NBER; xgabaix@stern.nyu.edu.

Manuscript received August, 2009; final revision received October, 2010.