Aggregate shocks or aggregate information? Costly information and business cycle comovement ☆

Laura Veldkamp[a],*, Justin Wolfers[b]

[a]*Stern School of Business, New York University, 44 West Fourth St., Suite 7-180, New York, NY 10012, USA*
[b]*Wharton School of Business, University of Pennsylvania, USA*

Abstract

Synchronized expansions and contractions across sectors define business cycles. Yet synchronization is puzzling because productivity across sectors exhibits weak correlation. While previous explanations emphasized production complementarity, our analysis explores complementarity in information acquisition. Because information about future productivity has a high fixed cost of production and a low marginal cost of replication, sectors can share the cost of acquiring aggregate information, rather than each paying the full production cost to forecast their sector-specific productivity. Sectors with common, aggregate information make highly correlated production choices. By filtering out sector-specific shocks and transmitting aggregate ones, information markets amplify business-cycle comovement.
© 2007 Elsevier B.V. All rights reserved.

Business cycles involve similar and synchronized movements in output across sectors. While this comovement across industries in output is readily observed, its source is not. In canonical business cycle models, similar output across industries is the result of aggregate shocks hitting each industry. Yet the data reject this view: while output is highly correlated across industries, total factor productivity is much less so (Rebelo, 2005). While alternative aggregate shocks may explain the comovement of output, Cochrane (1994) argues that they are elusive: "we haven't found large, identifiable, exogenous shocks to account for the bulk of output fluctuations" (p. 296). This led him to characterize business cycles as driven by shocks to *endogenous* variables. Alternatively, it could be that complementarities across industries synchronize output. Yet, the search for quantitatively important production complementarities has not produced a consensus. Our paper proposes a new source of complementarity, based on efficient joint information acquisition about sectoral productivity shocks.

Complementarity in information acquisition is a natural market outcome due to the special characteristics of information. Because information has a high fixed cost of production and is non-rival in consumption (and hence has a low marginal cost of replication), competitive information producers must charge more for highly tailored research (forecasting firm- or sector-specific shocks), to cover their high average cost. By contrast, forecasts of macroeconomic aggregates, are relevant to many producers. In equilibrium, aggregate forecasts will be purchased widely and produced at low average cost. The market for information facilitates firms sharing the expense of acquiring that information that they will find jointly useful, while there are fewer opportunities for joint consumption of sector-specific information.

In our model, firms do not necessarily want to produce more when others are producing more. They simply want to acquire the same information others acquire; similar information leads to similar decisions, and this drives comovement in both inputs and outputs. The low equilibrium price of aggregate information induces some firms to use aggregate data to make inferences about their sector's productivity. When many firms' inferences are based on common information, expected productivity is more correlated than true productivity. Since production decisions depend only on variables in producers' information sets, our model predicts that production decisions are highly correlated across sectors, and hence output is more highly correlated than productivity is. While sectoral productivity shocks are the model's driving force, the information market strips out much of their sector-specific effects. The information market passes on aggregate shocks to beliefs that mimic the aggregate shocks to endogenous choice variables Cochrane (1994) observes in the data.

Section 1 characterizes the key empirical puzzle: output is much more highly correlated across industries than productivity. Section 2 illustrates our mechanism in the simplest setting: an island model of production, appended with an information sector that supplies forecasts at an endogenously determined price each period. The results in Section 3 highlight our main contribution—for most industries our model generates comovement in output in excess of comovement in driving forces, roughly in proportion to the facts outlined in Section 1.

A typical problem with information friction models is that their effects disappear when observable market prices can efficiently aggregate information.[1] Since real business cycles

---

[1]Examples include herding, global games, and (Lucas, 1972) island models.

models are premised on efficient markets, this is a serious concern. Our effect does not disappear with a fully revealing price. Comovement arises because agents have similar information; it does not require information asymmetry. To show this, Section 4 adds a market for labor that fully reveals all information. Even when information is a public good, a few large sectors still discover it. Their signals become the basis for projections by the remaining smaller sectors and produce strongly correlated beliefs about productivity.

However, this model does not generate comovement in output. The problem is shared by many standard business cycle models: Highly cyclical industries demand lots more labor in booms, driving up the wage. Less cyclical industries, seeing mild cyclical increases in productivity but large wage increases, decide to hire fewer workers. Their output declines. This Markets all the industries with lower-than-average cyclicality have negative output correlations with aggregate output. The average output correlation ends up near zero (Christiano and Fitzgerald, 1999). Macroeconomists have discovered many mechanisms that solve this cyclical wage problem, including home production (Benhabib et al., 1991), habit persistence (Boldrin et al., 2001), or capital adjustment costs (Jaimovich and Rebelo, 2006). But these models are based on aggregate productivity shocks. With sector-specific shocks, they generate only as much comovement in output as what there is in productivity (Section 4.2).

Solving the *excess comovement puzzle* therefore requires two ingredients: a solution to the aggregate shock problem and a solution to the cyclical wage problem. Our information friction produces aggregate shocks, while in Section 4.3 home production is added to the model, as in Benhabib et al. (1991), yielding less-cyclical wages. The model delivers output comovement that resembles the data.

The idea that common information is unduly influential also arises in models of global games. In Morris and Shin (2002), a coordination motive in actions leads firms to overweight and overvalue common information, relative to what is socially efficient. Our coordination motive in learning does not require any coordination motive in production. Instead of being more valuable, our common information is less costly and an efficient way to economize on information costs.

Any theory of business-cycle comovement should also be able to account for its long-run changes and cross-sectional differences (Section 5). As Comin and Philippon (2005) have documented, much of the long-run decline in business cycle volatility can be traced to a decline in comovement across industries. A decreasing cost of information can replicate this decline. With cheaper information, sectors purchase more sector-specific information, rely less on aggregates, and comove less with the aggregate economy. In the cross-section, data on industry output and asset prices are also informative because equity prices summarize the information available about firms' profitability. Abundant information corresponds with lower comovement and more efficient input decisions, both in the model and the data. The industry output data shows that industries with a higher model-predicted demand for information comove less. These findings all support our main premise, that industry comovement comes from incomplete information.

1. The facts about comovement

Our data—which come from Basu et al. (2006)—describe the evolution within the United States of sectoral value-added and inputs of capital, labor, energy and

materials within each of 29 private non-farm, (roughly two-digit SIC) industries from 1949–1996.[2]

The comovement puzzle is immediately evident in these data: The average correlation of detrended sectoral output (that is, value-added) with aggregate output is 0.51, while the average correlation of detrended sectoral TFP with aggregate TFP is only 0.17. The high comovement in output comes from inputs: the average correlation of detrended sectoral inputs with aggregate inputs is 0.57. Why are input decisions highly correlated when productivity is not?

Shea (2002) provides a simple way to account for the importance of cross-industry correlations in driving aggregate movements. He notes that the variance of aggregate output can be approximated by $\mathbf{wVw'}$, where $\mathbf{w}$ is a vector of industry shares, and $\mathbf{V}$ is the variance–covariance matrix of sectoral output. He proposes decomposing aggregate output variance into a term due to the diagonal elements of $\mathbf{V}$, and a "comovement term" due to the off-diagonal elements. Performing this exercise on our data suggests that 83% of the variance in aggregate output is due to industry comovement (and the corresponding proportion of variation in aggregate input use due to comovement is 85%). Performing the same decomposition as above, the proportion of the variation in aggregate TFP due to comovement is only 14%. It is the contrast of relatively strong comovement in output with weaker comovement in sectoral productivity that presents a challenge to existing models.

Fig. 1 shows the general pattern that our analysis is trying to explain. This figure allows us to distinguish our interpretation of *excess comovement* from simple measurement error in total factor productivity. Specifically, excess comovement would lead the data to be clustered above the 45-degree line, while classical measurement error would attenuate bivariate TFP correlations, leading the data to be clustered above the 45-degree line for industries whose productivity series is positively correlated with aggregate TFP, but clustering of data *below* the 45-degree line for industries with negatively correlated TFP. Yet the data are (roughly) uniformly clustered above the 45-degree line.

In our simulations, productivity shocks are constructed to match industry correlations with the common shock, and fit is assessed in terms of how well the model matches the correlation of output with aggregate output. Our goal is to develop a mechanism in which correlation in output is systematically larger than that in productivity growth.

2. An island model of information-driven comovement

In order to isolate the mechanism that is new to this paper, our basic model makes information the only source of interaction between industries. There are N industries, each of which is an island populated by a representative agent. There are no trade or production linkages between industries; workers on each island consume what they produce. Information matters because island productivity is unknown at the start of each period. Moreover, before the island's inhabitant makes her labor/leisure choice, she can purchase (imperfectly) informative forecasts of either island-specific, or aggregate productivity.

---

[2]A technical appendix available on the authors' websites and in NBER working paper 12557 contains details about the data, sector-by-sector analysis of comovement, as well as proofs of all the propositions.
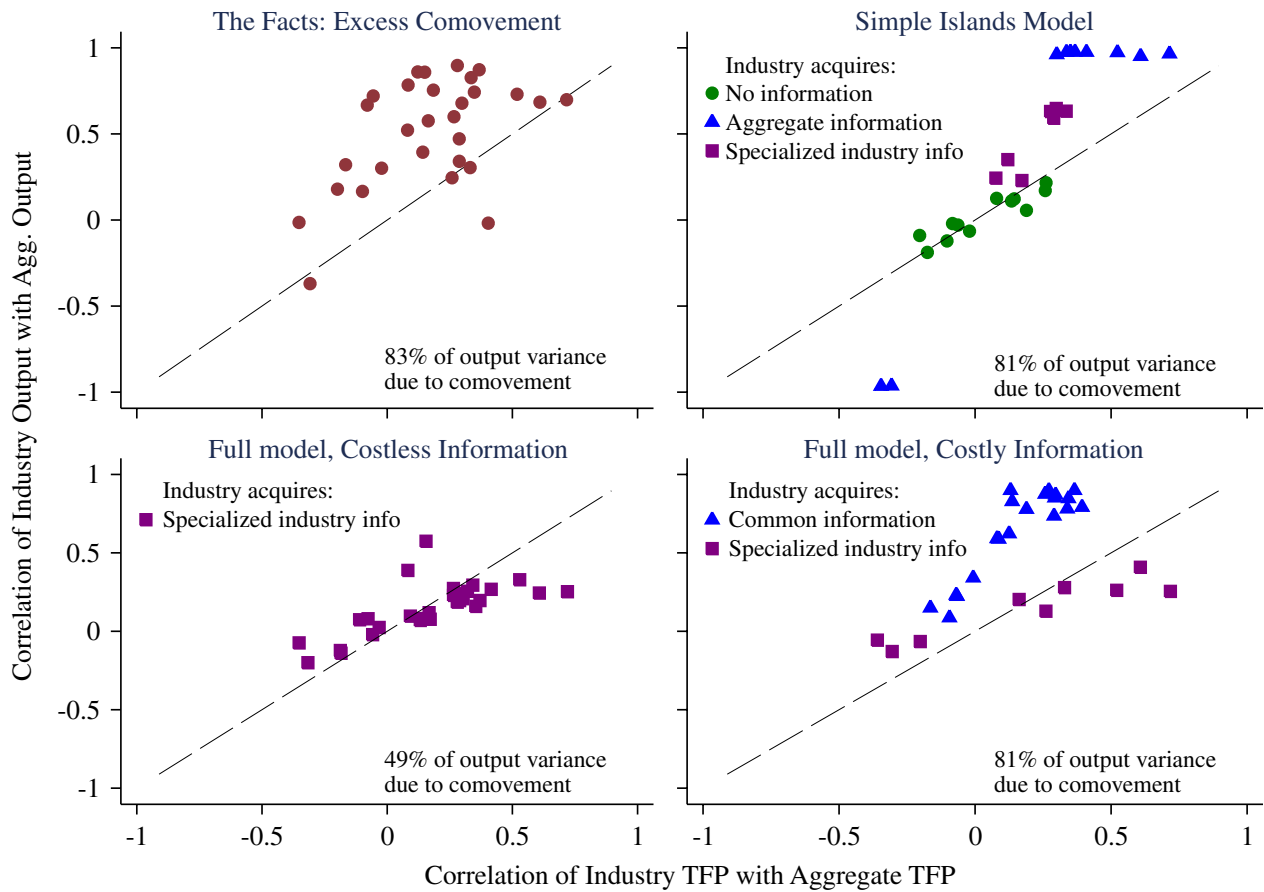
Fig. 1. Comovement of output and of total factor productivity: data and models. Excess comovement is represented by the cluster of correlations above the 45-degree line. The same industry has the same TFP correlation in the data and each model, by construction. Across the panels, a particular industry is identified by where it lies on the horizontal axis.

Islanders have exponential utility, with constant absolute risk aversion, ρ, and preferences defined over consumption c and labor n:

$$U_i = -E[\exp(-\rho(c_i - \psi n_i))]. \tag{1}$$

Exponential utility makes the value for information analytically tractable. Production in each industry is linear in labor, with marginal product z_i:

$$y_i = z_i n_i. \tag{2}$$

Importantly, labor's marginal product z_i is unknown. The first-order condition for labor yields:

$$n_i^* = \frac{1}{\rho \, Var[z_i|\mathscr{I}_i]} (E[z_i|\mathscr{I}_i] - \psi), \tag{3}$$

where $\mathscr{I}_i$ describes the information available on island i. Thus, labor effort is a function of beliefs about productivity, and is increasing in expected productivity, but decreasing in the level of uncertainty.

Each agent has a prior belief about productivity, which she may update in light of signals purchased in the information market. More specifically, each industry's

productivity has an aggregate component which is known, μ_z; an aggregate component which is unknown but learnable, $\bar{z}$; an industry-specific component that is unknown but learnable, η_i; and an industry-specific component that unknown and unlearnable, e_i:

$$z_i = \mu_z + \beta_i \bar{z} + \eta_i + e_i, \tag{4}$$

where μ_z is common knowledge, $\bar{z} \sim N(0, \sigma_z^2)$, $\eta_i \sim N(0, \sigma_n^2)$ and $e_i \sim N(0, \phi_j^2)$. All three random variables are mutually independent and η_i and e_i are i.i.d. across industries. Like the β of a financial asset, an industry's β measures the extent to which it covaries with the aggregate. Each industry's β is common knowledge.

Information markets: Three features of information are crucial. First, information is produced according to a fixed-cost technology. A signal s_j can be discovered at the beginning of the period at a fixed cost χ. This can be interpreted as the cost of hiring an economist to make a productivity estimate for an industry, a set of industries, or the economy as a whole. Information producers can choose to either produce a signal about the aggregate economy, s_0, or about a specific sector s_i. The information, once discovered, can be distributed at zero marginal cost. Second, reselling purchased information is forbidden. In reality, intellectual property rights prohibit this. Third, there is free entry into the provision of information. Any agent (on any island) can invest in the production of information at any time.

For simplicity, consider information suppliers competing on price in a perfectly contestable market.[3] As such, profits from information discovery depend on the price charged and demand for information, given the pricing strategies of other agents. One way to ensure that the market is contestable is to force information producers to choose prices in a first stage and choose entry in a second stage. Let the number of the agents that demand signal s_j, be λ_j. This depends on the price information producer j charges, p_j, and on all other posted prices for information. The profit from information production is price times demand, minus the fixed cost: $\pi_j = p_j \lambda_j - \chi$.

If an information producer invests in learning about aggregate productivity, she learns a signal

$$s_0 = \bar{z} + e_0, \tag{5}$$

where the noise $e_0 \sim N(0, \phi_0^2)$. If she chooses to learn about industry i's productivity, she observes an industry-specific signal, s_i, that is the sum of the noisy signal of aggregate productivity, plus the learnable industry-specific shock:

$$s_i = \beta_i(\bar{z} + e_0) + \eta_i. \tag{6}$$

This information structure ensures that any specific sector will prefer to learn its sectoral forecast than the aggregate forecast and in turn, will prefer to learn the aggregate forecast over another industry's sectoral forecast. Moreover, conditional on an own-sectoral forecast, there is little useful information in the aggregate forecast, and conditional on the aggregate forecast, there is no useful information in another sector's forecast. Therefore sectors choose to buy either no signal, only the aggregate signal, or only their

---

[3]This market structure is used because it produces a simple pricing formula. Veldkamp (2006) shows that Cournot or monopolistically competitive markets also produce information prices that decrease in the quantity sold.

sector-specific signal. (While these assumptions about the signal structure can be relaxed and still generate our main results, they simplify the choice problem.)

Updating beliefs: Given the various forms of uncertainty shown in (4), a sector's prior beliefs about its productivity, are $z_i \sim N(\mu_z, \beta_i^2 \sigma_z^2 + \sigma_\eta^2 + \phi_i^2)$, $\forall i$. The information contained in a signal depends on what kind of signal it is. Own-industry signals contain the most precise information about industry-specific productivity: $s_i \sim N(z_i - \mu_z, \beta_i^2 \phi_0^2 + \phi_i^2)$. Aggregate signals contain strictly less precise information about industry-specific productivity: $\beta_i s_0 \sim N(z_i - \mu_z, \beta_i^2 \phi_0^2 + \sigma_\eta^2 + \phi_i^2)$.

To form posterior beliefs about the industry productivity, agents combine their prior beliefs and signals, according to Bayes' law. With normal variables, posterior means are simply a precision-weighted average of priors and signals. For the agents who do not observe any signal, their posterior beliefs are the same as their prior beliefs. Those who observe their industry-specific signal believe that productivity is normally distributed with mean:

$$E[z_i|s_i] = \mu_z + \frac{s_i(\beta_i^2 \phi_0^2)^{-1}}{(\beta_i^2 \sigma_z^2 + \sigma_\eta^2)^{-1} + (\beta_i^2 \phi_0^2)^{-1}} \qquad (7)$$

and variance $Var[z_i|s_i] = ((\beta_i^2 \sigma_z^2 + \sigma_\eta^2)^{-1} + (\beta_i^2 \phi_0^2)^{-1})^{-1} + \phi_i^2$. Firms that observe only the aggregate signal believe their industry productivity is normally distributed with variance $Var[z_i|s_0] = \beta_i^2 (\sigma_z^{-2} + \phi_0^{-2})^{-1} + \sigma_\eta^2 + \phi_i^2$ and mean[4]

$$E[z_i|s_0] = \mu_z + \frac{\beta_i s_0 \phi_0^{-2}}{\sigma_z^{-2} + \phi_0^{-2}}. \qquad (8)$$

Definition of equilibrium:

1. Information producers announce prices p_j, at which they are willing to sell each signal s_j. After observing the prices posted by their competitors, they choose whether or not to incur the fixed cost to produce each signal, taking as given the action of other agents.
2. Taking information prices and availability as given, agents choose whether to purchase no information, aggregate information or sector-specific information.
3. Agents choose consumption and labor. All decisions maximize (1), given all signals discovered or purchased, and subject to the budget constraint: agents in each industry consume what they produce, plus information production profits, minus information costs:

$$c_i = z_i n_i + \sum_j (\pi_j - \tau_j L_{ij}) \quad \text{where } L_{ij} = 1 \text{ if agent } i \text{ buys signal } j \text{ and 0 otherwise.}$$

3. Main results

The results are presented in three pieces: input choices, information choices and numerical results.

---

[4]Were firms to observe other-industry signals, they would update in a similar fashion to those observing the aggregate signal; for their purposes an other-industry signal is simply the aggregate signal plus noise. To adjust posterior mean and variance, replace the aggregate signal noise ϕ_0^2 with the higher-variance noise $(\phi_0^2 + \sigma_\eta^2/\beta_j^2)$ and the aggregate signal realization s_0 with the other-industry signal s_j.

3.1. Aggregate information generates comovement

The only stochastic component of the model is the level of productivity, z_i. While z_i is not an input into production decisions, $E[z_i | \mathscr{I}_i]$ is the key driving force (Eq. (3)). Thus, with no new information, the firm's decisions would be the same every period. The addition of the information market ensures that $E[z_i | \mathscr{I}_i]$ is shaped by the information purchased, and this is the source of the model's excess comovement. For two firms that both observe their industry-specific signal, the correlation of their conditional expectations is dictated by the correlation of their observable fundamentals. In contrast, when two firms observe the aggregate signal, their common source of information gives them more highly correlated conditional beliefs.

Proposition 1. *If $\beta_i \beta_j > 0$, the correlation of labor input for any two industries is higher if those industries observe the aggregate signal only, than it would be if they observed their industry-specific signal: $corr(n_i^a, n_j^a) \geqslant corr(n_i^i, n_j^j)$. This inequality is strict if $\sigma_\eta^2 > 0$.*

The only force causing labor input to vary over time is the realization of the common signal. Since the only driving force is a common one, the labor input of two of these aggregate-informed industries is perfectly correlated. An industry that observes only aggregate information makes employment decisions that depend only on the aggregate signal s_0, yielding: $n_i^A = \alpha_i(s_0 + \gamma_i)$ where $\gamma_i = (\mu_z - \psi)(\phi_0^2 + \sigma_z^2)/(\beta_i \sigma_z^2)$, $\alpha_i = \beta_i/(\rho Var[z_i|s_0]) * \sigma_z^2/(\phi_0^2 + \sigma_z^2)$, and $Var[z_i|s_0]$ is given in Section 2.

The resulting covariances of output are derived in the technical appendix. Because output is a product of two normal variables, labor and productivity, its second moments are far less tractable. While output covariances with aggregate and firm-specific information cannot be compared analytically, simulating the model in Section 3.3 yields useful numerical comparisons.

3.2. Equilibrium information provision

In equilibrium, the price of a piece of information is the fixed cost of discovering it, divided by the number of agents who demand the signal: $p_j = \chi/\lambda_j$. Free entry at the stage where information prices are set ensures zero profit for information suppliers. If they made profits, other suppliers would enter the market, which would not be in equilibrium. Zero profit means that the price of information times the quantity demanded equals the cost of discovery: $c_j \lambda_j = \chi$. Thus, each information supplier prices at average cost.

Proposition 2. *A firm will purchase the aggregate signal s_0 if two conditions are satisfied:*

1. *Buying the aggregate signal at price χ/λ yields higher utility than buying the industry-specific signal, at the higher cost χ: $\log(Var(z_i)/Var(z_i|s_0)) - 2\rho\chi/\lambda \geqslant \log(Var(z_i)/Var(z_i|s_i)) - 2\rho\chi$.*
2. *Buying the aggregate signal yields higher utility than not purchasing any information: $\log(Var(z_i)/Var(z_i|s_0)) - 2\rho\chi/\lambda \geqslant 0$.*

As the economy becomes large ($N \to \infty$) or information becomes expensive ($\chi \to \infty$), industry-specific information becomes more and more costly and its demand falls to zero. All agents purchase either aggregate information (at a price χ/λ which is stationary for a

given fraction of islands purchasing information λ/N), or no information. With high variance of productivity z_i, information will be valuable and not learning will be a costly choice. This is the type of environment where many agents choose to observe aggregate information and where the strongest aggregate shocks emerge.

3.3. Simulating the model

Choosing parameter values: Each of our islands has the characteristics of one of the two-digit industries described in Section 1. Specifically, the aggregate TFP process is constructed to match the variance of aggregate TFP in the data ($\sigma_z = 0.4$ log points).[5] The industry TFP processes are constructed to match the variance of each industry's TFP in the data, and also to ensure that the covariance of each industry's TFP with aggregate TFP matches the data. This involved estimating a one-factor model, $TFP_{it} = \alpha_i + \beta_i T\tilde{F}P_{agg,t}$, and equating two features of the model and data: the loading on the aggregate factor β_i and the sum of the firm-specific shock variances ($\sigma_\eta + \phi_{ii}$). The relevant parameters for each industry are described in the on-line appendix. Thus, basic patterns in the data drive the pattern of shocks in our simulations.

One important feature of the model is how much of each shock is learnable. Arouba (2006) studies BLS estimates of labor productivity growth. He finds a signal-to-noise ratio of 1.06, for the initial estimates with respect to the final numbers. Therefore, informed agents are assumed to be as uncertain about current productivity as the official estimates are by setting all our signal-to-noise ratios to 1. For the aggregate shock, the variance of the noise is equal to the variance of the true shock ($\phi_0^2 = \sigma_z^2$). The equivalent assumption for the industry-specific shock is that the observable shock and the unobservable shock have equal variance ($\sigma_\eta^2 = \phi_{ii}^2$).

The remaining parameters are the disutility of labor ($\psi = 0.96$), which matches the relative volatility of labor hours in the data $std(n)/std(y) = 0.8$ (Benhabib et al., 1991), and risk aversion and the information cost, which affect the correlation of output only through their product ($\rho\chi$). Absolute and relative risk aversion are interchangeable in the labor market model because the average output per person is 1.01. Therefore, agents are assigned a standard risk aversion coefficient of $\rho = 3$. The information cost ($\chi = 0.17$) matches the average 0.16 R-squared when each industry's labor input is regressed on its own productivity shock. This makes sense as a measure of information because what makes information valuable is that it enables industries to choose labor that can covary with their productivity. However, this approach in no way guarantees success for the aggregate covariance. Industries with a low productivity-labor R-squared could have low-output correlation with aggregate output (if there is idiosyncratic noise in their choices) or high-output correlation (if labor choice has common aggregate shocks). The commonality of shocks to labor inputs is endogenous. This is not a large friction. Only 0.025% of total output is devoted to information payments in equilibrium. Section 5.5 explores other information costs.

Results: Fig. 1 illustrates the comovement in each simulated sector. The important feature of this graph is that the industries that observe the aggregate signal have

---

[5]Aggregate TFP is constructed as the weighted average of industry TFP. Aggregate output will also be a weighted average, where each industry's output is first scaled to have the same standard deviation as the output data and then weighted. In both cases, the weights are the industry's share in total value-added.

significantly higher-output correlations, in excess of their TFP correlations. The average correlation of an industry with aggregate output is 52%, compared to 51% in the data. For productivity, this average correlation is 17%, equal to the data by construction.

One obvious failing of this model is its prediction that industries with counter-cyclical TFP $\beta_i < 0$ have strongly counter-cyclical output. In the data, more than half of these industries have pro-cyclical TFP. What this model does is make a heterogeneous-shock economy more closely resemble an economy with one aggregate shock. In that sense, it achieves its goal. But a one-shock model would also predict that the output of such industries is counter-cyclical. One way to remedy this problem is to introduce a second aggregate shock. Then, even firms with counter-cyclical TFP can potentially have pro-cyclical output. An example of this arises in Section 4.

One of the reasons that it is important to understand comovement is because it helps to explain the volatility of business cycles. Fluctuations in output are more volatile than measured changes in productivity can account for. If output covaries highly across sectors, then the sum, aggregate output, will have higher volatility. This initial model produces additional volatility. In fact, it produces too much of it, yielding a 9.5% standard deviation of log output. There are at least two simple reasons model output is too volatile. First, there are no decreasing returns. The linearity of the production function simplifies the analysis, but makes labor fluctuate excessively. Second, there is no capital which, unlike labor, is a slow-moving stock that must be built up over time. The presence of capital would dampen output fluctuations.

Although the model's output variance does not match the data, the fraction of this variance that comes from sectoral covariance does match. Using the Shea (2002) decomposition, the fraction of variance in aggregate output due to sectoral comovement is 83% in the data and 81% in the model simulation.

This model illustrates how information markets dampen the effect of heterogeneous information and transmit more aggregate information. In this instance, the aggregate information is aggregate productivity. But the mechanism is more general than that. It filters out all kinds of industry-specific and firm-specific information and delivers aggregate shocks to beliefs. These shocks to beliefs show up in the data as shocks to endogenous choice variables, just as they do in the data.

4. Adding a market for labor

The island model of Section 2 illustrates why the non-rival nature of information produces complementarity in information acquisition, and why this can induce comovement in industry output. However, that model is missing many realistic features of the aggregate economy such as labor markets or tradeable goods. This might be a cause for concern: market prices that aggregate information have been the Achilles' heel of other information-friction models. It also leaves unexplored how inter-island linkages can affect information transmission and output comovement.

In this section, workers can now supply labor to other islands, at a market wage. Because variation in sectoral employment is driven by differences in information, prices and quantities in the labor market reveal all known information. But the information mechanism that filters out heterogeneous information and transmits aggregate shocks to choice variables still functions, albeit with some modifications, in the presence of a

perfectly efficient market price. Thus, the model's mechanism is robust to both a labor market and the unravelling of all private information.

Setup: Each agent i is endowed with a technology z_i and 1 unit of time, which she supplies to a labor market, or consumes as leisure $l_i < 1$. Agents have exponential preference for consumption c:

$$U = -E[\tilde{\xi}(1 - l_i)^2 \exp(-\rho c)]. \tag{9}$$

The form of utility over leisure needs to change from the previous model because the payoff to supplying labor is no longer random. Whereas before, the effective wage was the marginal product z_i, which was unknown, now the wage w is known when l_i, n_i are chosen. Using the previous utility formulation would yield an indeterminate labor supply with a non-random payoff. This form of utility ensures the labor supply is a linear function of the wage and keeps the model tractable.

Production requires combining labor n, which can be hired at an endogenous wage w, and technology to get output of sector i: $y_i = z_i n_i$.

Agents can also decide whether to discover information about their sector $\lambda_i = 1$, or not $\lambda_i = 0$. Discovering information about your sector means that the signal $s_i = \beta_i \bar{z} + \eta_i$ becomes common knowledge. The public good nature of information rules out the possibility of sharing this cost through an information market: no one will share the cost of an aggregate or other-industry signal they will otherwise observe for free. However, a few large industries will underwrite public provision of their own-industry information. Other industries infer their productivity from this other-industry information, for free.

The agent's budget constraint is that she can eat what she produces, minus the cost of the labor, net of labor supplied, and the cost of information discovery:

$$c_i = z_i n_i - w n_i + w(1 - l_i) - \chi \lambda_i. \tag{10}$$

Equilibrium: An equilibrium is a set of labor $\{n_i \geqslant 0\}$, leisure $\{l_i\} \in [0, 1]$, consumption $\{c_i \geqslant 0\}$, and information choices $\{\lambda_i\} \in \{0, 1\}$ that maximize (9) subject to (10), and a market wage w that clears the labor market: $\sum_i n_i = \sum_i (1 - l_i)$.

4.1. Costly information but no home production

The first-order condition tells us that the optimal labor demand in industry i is

$$n_i^* = \max\left\{\frac{E[z_i | \mathcal{I}_i] - w}{\rho \, Var[z_i | \mathcal{I}_i]}, 0\right\}. \tag{11}$$

The first-order condition for leisure dictates the labor supply:

$$1 - l_i = \frac{\rho}{2\tilde{\xi}} w \equiv \xi w, \tag{12}$$

where $\xi \equiv \rho/(2\tilde{\xi})$. Since all information is public, the wage is known to all agents. Substituting (11) and (12) into the labor market clearing condition delivers the equilibrium wage:

$$w = \frac{\sum_{i=1}^{N} E[z_i | \mathcal{I}_i] / Var[z_i | \mathcal{I}_i]}{\sum_{i=1}^{N} 1 / Var[z_i | \mathcal{I}_i] + \rho\xi} + \zeta. \tag{13}$$

where $\zeta = 0$ if the non-negativity constraints on n_i do not bind.

An agent that discovers their industry signal believes that their productivity is $z_i \sim \mathrm{N}(\mu_z + s_i \phi_i^{-2} v_i, v_i)$, where $v_i = 1/(\phi_i^{-2} + (\sigma_z^2 + \sigma_\eta^2)^{-1})$. For an agent who infers their sector productivity from others' signals, the precision of their information depends on how many of those signals were observed. Suppose that S signals are in the public domain: $\{s_1, \ldots, s_S\}$. Each one is an independent signal about $\bar{z}$ with variance σ_η^2/β_i^2. Thus, $E[\bar{z}|s_1, \ldots, s_S] = \mu_z + \sigma_\eta^{-2}(\sum_{i=1}^S \beta_i^2 s_i) Var[\bar{z}|s_1, \ldots, s_S]$ and $Var[\bar{z}|s_1, \ldots, s_S] = (\sigma_\eta^{-2} \sum_{i=1}^S \beta_i^2 + \sigma_z^{-2})^{-1}$. Since agents who did not discover their industry signal have no information about the industry-specific component of their productivity, their expected value of their own z_i and $\bar{z}$ are the same. But their beliefs about their own industry have higher variance: $Var[z_i|s_1, \ldots, s_S] = Var[\bar{z}|s_1, \ldots, s_S] + \sigma_\eta^2 + \phi_i^2$.

Since all agents that do not discover information have identical beliefs about their productivity, they have highly correlated labor decisions, just like in Section 2. Whether labor decisions are highly positively or negatively correlated depends on whether the industry's TFP is more or less pro-cyclical than the wage.

Proposition 3 (*Correlation of labor inputs*). *Suppose that the non-negativity constraint on labor does not bind ($\xi = 0$). If only one industry l chooses to observe its industry-specific productivity, then $corr(n_i, n_j) = 1$ or -1, for all firms i and j.*

If more than one industry chooses to observe its industry-specific productivity, but industry i and industry j both choose not to, then $corr(n_i, n_j) = 1$ iff $\beta_i = \beta_j$.

When there is only one firm learning in the economy, there is only one shock that determines all firms' beliefs and the wage. Since $E[z_i|\mathscr{I}_i] - w$ is proportional to the same random variable, for every firm, all firms have perfectly correlated labor inputs. When more than one firm learns, the wage and the posterior beliefs depend on the realized signals in slightly different ways. Therefore, each firm's labor inputs are a different linear combination of these two shocks. That linear combination is the same only when both firms are uninformed and have the same loading on aggregate productivity.

Information choice and the free-rider effect: When labor prices and quantities fully reveal the information others observe, information becomes a public good. This introduces a free-rider problem that was not present in the baseline model. But this 'problem' actually strengthens our result. Free-riding makes firms not want to learn, when other firms are learning. It is the firms that do not learn that base their actions on aggregate information and exhibit the excess comovement. Thus, the more industries free-ride off the information of others, the stronger comovement becomes.

Simulation: With one exception, the parameter values are identical to those in Section 3.3. The information cost needs to be lower ($\chi = 0.0015$) to achieve the same targeted relationship between labor and productivity. The implied information cost changes for two reasons: The benefit from information that is publicly revealed is lower than the benefit from the private information in Section 2, and the level of output changes when the preferences change.

This model fails. Output comovement is near zero (-4%), on average. The problem is that highly cyclical (high-β) industries demand lots more labor in booms, driving up the wage. Less cyclical (low-β) industries, seeing mild cyclical increases in productivity but large wage increases, decide to hire fewer workers, which results in declining output. This pushes all the industries whose expected productivity is less cyclical than the wage to have

negative output correlations with aggregate output. This problem is common to standard RBC models, even ones with aggregate shocks.

4.2. Home production and costless information

A simple solution to the cyclical wage problem is a large, not highly cyclical, home production sector (Benhabib et al., 1991). Because this sector has productivity shocks that are not very correlated with the aggregate productivity, it makes the wage less cyclical. It absorbs lots of labor when aggregate productivity is low, keeping the wage up, and sheds workers when aggregate productivity is high, keeping the wage from rising.

To illustrate the effect of adding this home production sector, this subsection turns off the information friction by making information costless ($\chi = 0$).

Simulation: The introduction of one new sector leaves two new parameters to be chosen. The loading of the home production industry on aggregate productivity ($\beta_h = 0.05$) matches the standard deviation of aggregate market labor in the data (0.8, as reported in Benhabib et al., 1991). The industry-specific variance $\phi_h^2 = 0.34\%$ equates the time spent on home production and on market labor, similar to the division of time in the data (Benhabib et al., 1991).

The average correlation of each industry's output with the aggregate (18%) is lower than the 51% correlation in the data. Furthermore, less than half (49%) of the variation in output comes from comovement, compared to 83% in the data. Adding home production is clearly an improvement because without it, average comovement is close to zero. But, home production alone cannot explain the large output correlations in the data.

4.3. Home production and costly information

Replicating the high-output comovement in the data requires both solving the cyclical wage problem and generating some aggregate shocks. This section puts the two pieces together. The setting is exactly the same as before, but with a positive information cost ($\chi = 0.01$).

For many industries, output correlation in the model matches or exceeds that in the data (Fig. 1). The strength of information frictions in generating comovement is just as strong with fully informative prices, as it is in the island model. The average sector's correlation with aggregate output is 52%, slightly more than the 51% in the data. Likewise, the Shea decomposition tells us that fraction of aggregate output variance from comovement is 81%, compared to 83% in the data. The information friction generating these effects is not a large one; the economy-wide expenditures on information are only 0.05% of the value of total output.

Correlation sign reversals: There is one common factor $\bar{z}$ that every sector's productivity depends on. But one factor, by itself, cannot explain why TFP and output correlations sometimes have different signs. Information markets can only magnify or dampen their correlation. Yet, Fig. 1 shows industries in the top-left quadrant with negative productivity correlations and positive output correlations. The reason for this sign reversal lies in the home-production sector. Because the home sector is so large, a change in its productivity affects the wage. This is a form of aggregate shock that all market industries load on with the same sign: when wages are higher, every industry produces less than they otherwise would.

This result is not meant to persuade the reader that home-production shocks are an important driver of business cycles. Rather, it illustrates how information markets can amplify all sorts of aggregate shocks, not just productivity shocks, by transmitting them and suppressing reactions to idiosyncratic shocks. Those other aggregate shocks are likely the usual suspects: inflation risk, oil prices, shocks to foreign markets. Such events are transmitted cheaply to many readers who react to them in a coordinated way, causing output in many sectors to comove. Similar to Gabaix (2005), what moves aggregate output is aggregate productivity plus the shocks to the economy's largest firms. But instead of relying on input–output linkages to propagate these shocks, our information markets transmit them and, in the process, create comovement in the beliefs of all firms.

5. Distinguishing this theory from the alternatives

The challenge in evaluating a theory based on informational scarcity is that the signals agents chose to acquire are not directly observable or measurable. Instead, only indirect measures of information quantity are available. One such measure is the information content of a firm's equity price. Another approach is to let the model predict which firms should learn more and how information choice should evolve over time. Evidence from both approaches supports the theory and distinguishes it from the real complementarity alternatives.

5.1. Theories of comovement

The leading explanations of business-cycle comovement employ input–output linkages to generate comovement. Because one industry uses the product of another, this real complementarity propagates industry-specific shocks across industries. Originally proposed by Long and Plosser (1983), work by Hornstein and Praschnik (1997), Horvath (1998), Horvath (2000), and Shea (2002) argues that these linkages can explain the extent of comovement.

However, other authors dispute the explanatory power of linkages. Christiano and Fitzgerald (1999) look across sectors and find considerable variation in linkages among sectors that all display strong business-cycle comovement. Dupor (1999) argues that input–output matrices are too sparse to constitute a major source of comovement. Rebelo (2005) concurs, calling comovement an ''under-researched topic in business-cycle research.'' Our paper offers a new mechanism to either challenge or complement sectoral linkages. Exploring this alternative may potentially improve our understanding of the sources of business-cycle comovement.

5.2. Decomposing excess output comovement

A simple statistical exercise shows that amplification of common productivity shocks is a key driver of comovement in output.[6] There are two reasons why comovement in output might exceed comovement in driving forces: (1) Output comovement might reflect amplification of common productivity shocks or, (2) after controlling for aggregate productivity shocks, residual variation in output may be highly correlated across sectors.

---

[6]Thanks to Lucrezia Reichlin for suggesting this approach.

Our costly information model provides an example of a theory that amplifies common productivity shocks, while models based on intersectoral linkages emphasize the high degree of correlation of residual output across sectors. For example, if a TFP shock in one sector causes relative prices and output of many complementary sectors to rise, this would yield an increase in output not fully explained by aggregate or own-industry TFP, and would show up as correlated residual output in the complementary sectors.

For each industry, the exercise is to regress the industry output gap against four lags of both the common TFP shock, and the idiosyncratic industry-specific TFP shock. The predicted values from this regression, which reflect the influence of aggregate and industry-specific TFP shocks, generate much of the output comovement in the data. Whereas the average correlation of sectoral output with aggregate output is 0.51, the predicted values have an average correlation of 0.42. This far exceeds the empirical average correlation of sectoral TFP with aggregate TFP of 0.17. If one instead regressed output on the industry TFP shocks, which include both their aggregate and idiosyncratic components equally weighted, the average correlation of each industry's predicted output with aggregate output would be only 0.13. Thus, most of the comovement puzzle can be explained by the extent to which output loads more heavily on aggregate TFP than on its industry-specific component. Moreover, these predicted outputs also successfully pinpoint which industries show the most substantial excess comovement.

This exercise suggests that a mechanism such as an information market, which amplifies the effects of common productivity shocks, is important to explain the "excess comovement" puzzle. Even so, there is still a role for alternative theories, as the residuals from this regression are highly correlated. It is the addition of these residuals back onto the predicted values that push the correlation of sectoral output with the aggregate from 0.42 up to the observed level of 0.51.

5.3. Output comovement: evidence from manufacturing data

According to Proposition 3.2, sectors with more sector-specific productivity variation (high $Var(z_i|s_0)$) value sector-specific information more. If such sectors acquire more sector-specific information, then their excess comovement will be low. This is not saying that more idiosyncratic variation in industry TFP implies lower correlation between sectoral and aggregate TFP: that is simply a mechanical relationship. Rather, the model predicts that such sectors should have lower output correlation, after accounting for the lower correlation of productivity.

This prediction is borne out in US data on output and TFP across the 459 4-digit manufacturing sectors covered in the NBER-CES Manufacturing Industry Database.[7] Our estimates of $Var(z_i|s_0)$ come from taking the industry-specific residuals from an estimated one-factor TFP model. Analyzing detrended annual data from 1958 to 1996 yielded the following estimate:

$$corr(GDP_i, GDP_{\text{agg}}) - corr(TFP_i, TFP_{\text{agg}}) = 0.048 - 5.46 \cdot Var(TFP_i|TFP_{\text{agg}}), \quad (14)$$

---

[7]These data come from the NBER-CES Manufacturing Industry Database, by Eric J. Bartelsman, Randy A. Becker, and Wayne B. Gray, June 2000 (http://www.nber.org/nberces/nbprod96.htm). These detailed sectoral data are our focus because the two-digit data emphasized earlier in the paper offered too few data points to estimate any effect precisely.

where the standard deviations of the estimated coefficients are (0.006) and (2.89), respectively.

High levels of idiosyncratic TFP variation—which in the model would be associated with greater demand for industry-specific information—are associated with a smaller gap between the correlation between industry and aggregate output, and the correlation between industry and aggregate TFP. To get a sense of the relevant magnitudes for instance, note that the difference between a firm at the 25th percentile of idiosyncratic TFP variance (0.04%) and a firm at the 75th percentile (0.15%) is about a 12% decline in excess comovement. This effect is also statistically significant at the 90% confidence level. Finally, a simulation exercise shows that if there is measurement error in TFP, our estimates understate the true relationship between idiosyncratic variance and excess comovement.

5.4. Information comovement: evidence from financial markets

The efficient markets hypothesis says that equity prices summarize available information regarding a firm's profitability. Since equity prices inform us about what information firms might have, the precision with which prices predict future earnings is a proxy for firm information. This idea that equity prices inform firms' input choices is supported by Chen et al. (2006).

Three findings in the empirical finance literature confirm predictions of our costly information model. First, the model predicts that cheap, aggregate information should be more abundant than expensive industry- or firm-specific information. Empirical finance studies confirms this prediction: equity prices of different firms and sectors comove more than corporate earnings. That is, prices are driven more by aggregate than firm-specific information.

Second, the model predicts that when information is scarce, it has more aggregate content. When information is abundant, firms have access to signals with more firm-specific content. In Durnev et al. (2003), the prices most strongly driven by their aggregate component have the least information; they are the worst predictors of future earnings. At the industry level, greater industry-specific price variation (less reliance on aggregate signals) is also correlated with a larger number information-producing analysts (Piotroski and Roulstone, 2003).

Finally, the model predicts that firms whose signals contain more industry-specific information make more efficient labor input decisions. In the data, firms whose equity prices show little firm-specific variation also make less efficient input choices. Durnev et al. (2004) find that firms with more firm-specific equity price variation make more efficient real investments.

5.5. The long-run fall in comovement

Over the last 30 years, firm-level volatility increased and aggregate volatility decreased. Comin and Philippon (2005) show these two facts imply a decrease in firm comovement. Our model could explain this trend from a decreasing cost of information over time. A falling information cost causes more firms to acquire firm- or industry-specific signals. This moves the economy closer to the full-information economy, which exhibits no excess comovement. Fig. 2 illustrates how small changes in information cost can replicate the observed comovement decline.
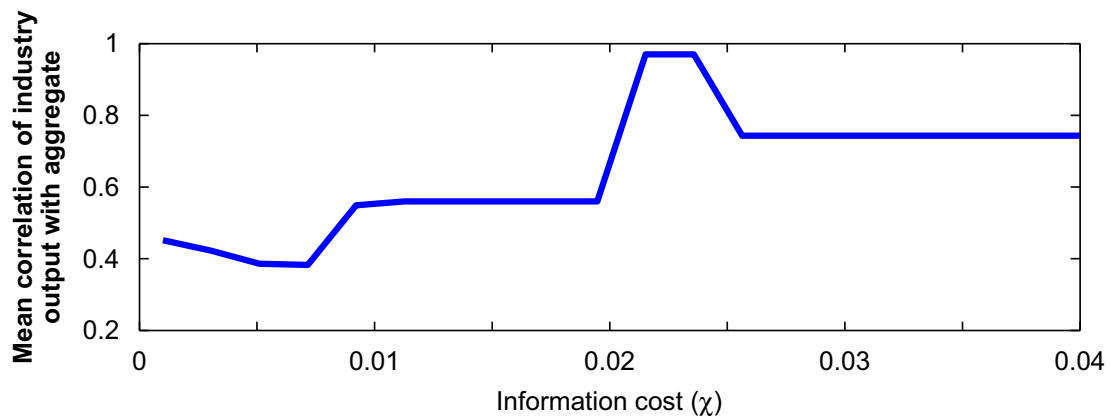
Fig. 2. A falling information cost makes comovement increase, then decrease. Average correlation of industry and aggregate output for the simulated labor market model with home production. Parameters are the same as in previous island and labor market simulations. The total amount paid for all information varies between 0.05% and 0.13% of aggregate GDP.

When information costs are very low, the economy resembles the full-information model of Section 4.2. Because each sector gets its sector-specific information, there are no aggregate shocks introduced by information. As the information cost rises, firms economize on information costs by purchasing cheaper, aggregate information. In turn, this leads all firms to focus on the aggregate shock, yielding substantial comovement. When information becomes very expensive, little information is acquired. Less information has less effect on agents' beliefs. As information about the aggregate shock becomes scarcer, comovement declines. In the extreme, if neither the market nor the home sector were to acquire any information, there would be no movement in beliefs, output correlation would be determined only by TFP correlation, and there would be zero excess comovement.

If the economy started somewhere to the left of the peak in comovement, then arbitrarily small decreases in information cost have the potential to dramatically decrease output comovement. The opposite prediction, that a decline in information costs would increase comovement, would require information to be so scarce that it paralyzed decision-making. While this is theoretically possible, it is not realistic. A model with such limited information would have no movement in inputs. Furthermore, free data about the aggregate economy, in practice, is not scarce. Therefore, decreases in information costs should cause comovement to fall over time.[8]

6. Conclusion

Industry comovement in business cycles may arise from firms' desires to economize on information costs. Learning aggregate, rather than firm-specific information allows firms to share the costs of aggregate information with other firms, or to free-ride off information other firms acquire. Both of these possibilities arise because of the non-rival nature of

---

[8]Davis et al. (2006) have recently documented the same decrease in firm comovement among publicly-traded firms as (Comin and Philippon, 2005). However, they find that the cross-sectional dispersion of growth rates for privately-traded firms has fallen as aggregate volatility has declined. This finding suggests that future work should focus on differences in the information that public and private firms acquire.

information. Because information used by one firm can also be used by another firm, whereas capital or labor cannot be shared in such a way, information is a natural candidate for a source of hidden complementarities.

In practice, many companies do not purchase information. Instead, governments and news agencies facilitate information cost-sharing. Governments use tax revenue to collect data and provide some statistics for free. The statistics they choose to provide are the high-social-value statistics relevant to many of their constituents—or in other words, aggregate signals. Likewise, the goal of news agencies is to maximize the size of their audience. To do this, they also provide high-demand aggregate information.

The information that producers purchase does not have to be about productivity. Information about any factor that has a firm-specific and an aggregate component to it that is relevant to production decisions could cause comovement. Alternatives include demand, wage, or price information. Also, the theory could explain firm comovement within industries.

Real business-cycle models were designed to explain short-run and long-run growth with the same tools. In long-run growth, information's non-rival nature allowed it to sustain growth, where accumulating physical goods could not (Romer, 1990). Given that production of non-rival information unlocked long-run growth, why should not it play a key role in business cycles as well?

References

Arouba, B., 2006. Data revisions are not well-behaved. Working paper, University of Maryland.

Basu, S., Fernald, J., Kimball, M., 2006. Are technology improvements contractionary? American Economic Review.

Benhabib, J., Rogerson, R., Wright, R., 1991. Homework in macroeconomics: household production and aggregate fluctuations. Journal of Political Economy.

Boldrin, M., Christiano, L., Fisher, J., 2001. Habit persistence, asset returns, and the business cycle. American Economic Review 91 (1), 149–166.

Chen, Q., Goldstein, I., Jiang, W., 2006. Price informativeness and investment sensitivity to stock price. Review of Financial Studies.

Christiano, L., Fitzgerald, T., 1999. The business cycle: it's still a puzzle. Economic Perspectives, Federal Reserve Bank of Chicago, vol. 1998(4), pp. 56–83.

Cochrane, J., 1994. Shocks, Carnegie–Rochester Conference Series on Public Policy, vol. 41, pp. 295–364.

Comin, D., Philippon, T., 2005. The rise in firm-level volatility: causes and consequences. In: NBER Macroeconomics Annual.

Davis, S., Haltiwanger, J., Jarmin, R., Miranda, J., 2006. Volatility and dispersion in business growth rates: publicly traded versus privately held firms. NBER Working Paper #12354.

Dupor, W., 1999. Aggregation and irrelevance in multi-sector models. Journal of Monetary Economics 43, 391–409.

Durnev, A., Morck, R., Yeung, B., Zarowin, P., 2003. Does greater firm-specific return variation mean more or less informed stock pricing? Journal of Accounting Research 41 (5), 797–836.

Durnev, A., Morck, R., Yeung, B., 2004. Value enhancing capital budgeting and firm-specific stock return variation. Journal of Finance 59 (1), 65–105.

Gabaix, X., 2005. The granular origins of aggregate fluctuations. MIT Working Paper.

Hornstein, A., Praschnik, J., 1997. Intermediate inputs and sectoral comovement in the business cycle. Journal of Monetary Economics 40 (3), 573–595.

Horvath, M., 1998. Cyclicality and sectoral linkages: aggregate fluctuations from independent sectoral shocks. Review of Economic Dynamics 1, 781–808.

Horvath, M., 2000. Sectoral shocks and aggregate fluctuations. Journal of Monetary Economics 45, 69–106.

Jaimovich, N., Rebelo, S., 2006. Can news about the future drive the business cycle? Working paper.

Long, J.B., Plosser, C.I., 1983. Real business cycles. Journal of Political Economy 91 (1), 39–69.

Lucas, R., 1972. Expectations and the neutrality of money. Journal of Economic Theory 4 (2), 103–124.

Morris, S., Shin, H.S., 2002. Social value of public information. American Economic Review 92 (5), 1521–1534.

Piotroski, J., Roulstone, D., 2003. Analysts, institutional investors and insiders: what information do they provide? Working paper.

Rebelo, S., 2005. Real business cycle models: past, present, and future. Scandinavian Journal of Economics 107 (2), 217–238.

Romer, P., 1990. Endogenous technological change. Journal of Political Economy 98 (5), 71–102.

Shea, J., 2002. Complementarities and comovements. Journal of Money, Credit and Banking 34 (2), 412–433.

Veldkamp, L., 2006. Media frenzies in markets for financial information. American Economic Review 96 (3), 577–601.