

Discussion of "Noisy Business Cycles", by Angeletos and La'O *

Christian Hellwig[†]

UCLA and CEPR

July 21, 2009

Recent years have seen renewed interest in business cycle theories that emphasize the role of incomplete and possibly heterogeneous information in the propagation of aggregate shocks. Whereas the early literature (e.g. Lucas, 1972) placed its emphasis on models in which either information was homogeneous, or strategic interactions were limited, the new models emphasize the interactions between individuals with heterogeneous expectations, when equilibrium is shaped by uncertainty about fundamentals as well as uncertainty about aggregate behavior.¹ So far, this new approach has produced two important novel insights.

1. *In economies with decision complementarities, information heterogeneity amplifies adjustment delays (Woodford, 2003a).* With heterogeneous information, individuals update their decisions more slowly than their information, if they fear that new information is not shared by others, with whom their decisions are complementary. In contrast, the original models were unable to account for adjustment lags that exceeded the arrival of new information. Furthermore, information heterogeneity potentially explain the excessive weight that decisions often attach to public news, which play a disproportionate role in forecasting aggregate behavior.

2. *The welfare implications of heterogeneous information are tied to the presence of externalities (Morris and Shin, 2002; Hellwig, 2005; Angeletos and Pavan, 2007).* Increased persistence and excess sensitivity to public news need not be a symptom of inefficiency. (In-)efficiency instead results from externalities that are do not related to the strategic incentives. Besides payoff externalities, such as technical spill-overs or mark-up variation, and informational externalities when strategies affect the quality of information (such as information aggregated through market prices), the new

*I thank Marios Angeletos, Jennifer La'O, Mike Woodford and Venky Venkateswaran for helpful discussions, and the Alfred P. Sloan fellowship for financial support of my research. An extended version of this discussion with proofs of all formal claims is available from the author's website. The usual disclaimers apply.

[†]Email: chris@econ.ucla.edu

¹See AL for a detailed list of contributions. See Hellwig (2008) for a quick review of the recent and older literature.

heterogeneous information models have highlighted a trade-off between aggregate adjustment and cross-sectional heterogeneity resulting from the use of private information. Inefficient aggregate fluctuations or cross-sectional dispersion may result from a third type of externality in the relative weight of these two competing concerns.

The paper by Angeletos and La'O (2009, henceforth AL) is an excellent contribution to this emerging literature. It studies a real business cycle economy with preference complementarities, or “trade linkages” across products and sectors. In the model, firms commit to quantity decisions before they can observe the current aggregate conditions (i.e. productivity). Because of the trade linkages, each firm runs the risk of ending up with lower prices, revenues and profits, if it produces too much relative to the others. This entails a need to forecast the production level of others, a task that is rendered more difficult when firms don't all share the same information. The resulting uncertainty about aggregate activity leads to interesting departures from the canonical RBC model: it amplifies the effects of uncertainty about technology, mutes the short run response of output to technology shocks, and adds common expectation shocks as a potentially important source of short-run fluctuations. As a result, the model can account for evidence from identified VARs that appear to be inconsistent with the canonical RBC framework, namely the negative response of hours worked to productivity shocks, and the fact that productivity accounts only for a small fraction of short run fluctuations in output. Finally, the paper shows that, perhaps surprisingly these equilibrium adjustment dynamics do not reflect any inefficiencies. As long as there are no mark-up shocks and no endogenous signals, firms make socially optimal use of the available information in equilibrium.

In this discussion, I use a simplified version of the model in AL to review their business cycle implications and relate them to the general insights derived from the new heterogeneous information models.² I then discuss the potential of this class of models and the challenges they face in accounting for quantitative features of business cycles. Perhaps my main criticism about this paper is that it does not go far enough to subject its theoretical results to a quantitative evaluation.

1 Theoretical Results

A simple model: Consider a canonical new Keynesian economy, in which a continuum of firms indexed by $i \in [0, 1]$ each produce a differentiated commodity, using a differentiated labor input N_t^i

²The focus on business cycle implications should not divert attention from the efficiency result. The latter is important, since (i) it extends the efficiency properties of the canonical RBC model, and (ii) it is in stark contrast with models of nominal price adjustment or investment with complementarities and heterogeneous information, in which equilibrium information use is generally inefficient (Hellwig, 2005; Angeletos and Pavan, 2004).

as a unique factor of production. In a first stage of each period, intermediate producers commit to their output choices Q_t^i after observing their own TFP A_t^i , as well as additional noisy signals about aggregate TFP A_t . Markets open in a second stage.³ There are no frictions other than the firms' pre-commitment to output decisions and lack of information about aggregate TFP in stage 1.

The economy is summarized by firm-specific production, demand, and labor supply functions, and an aggregate pricing kernel:

$$\text{Production function: } Q_t^i = A_t^i (N_t^i)^\theta \quad (1)$$

$$\text{Product demand: } Q_t^i = Q_t (P_t^i)^{-\rho} \quad (2)$$

$$\text{Wages: } W_t^i = Q_t^\gamma (N_t^i)^\varepsilon \quad (3)$$

$$\text{Pricing kernel: } \lambda_t = Q_t^{-\gamma}. \quad (4)$$

The variable $Q_t = \int (Q_t^i)^{1-1/\rho} di$ denotes aggregate output, and P_t^i denotes firm i 's relative price (so that $\int (P_t^i)^{1-\rho} di = 1$). The parameter $\theta \in (0, 1)$ represents the returns to scale, ρ the CES elasticity of demand, ε^{-1} the Frisch elasticity of labor supply, and γ is the household's degree of relative risk aversion, or more relevant for our purpose, the output elasticity of wages.

Before choosing quantities, firms perfectly observe their own TFP $\ln A_t^i \equiv a_t^i = a_t + \eta_t^i$, where $a_t \sim N(0, \kappa_a^{-1})$ and $\eta_t^i \sim N(0, \kappa_\eta^{-1})$ are aggregate and idiosyncratic TFP shocks, respectively.⁴ In addition, each firm observes an idiosyncratic signal x_t^i and an aggregate signal $z_t = a_t + v_t$ of the aggregate technology level, where $x_t^i \sim N(a_t, \kappa_x^{-1})$ and $v_t \sim N(0, \kappa_z^{-1})$. The firm's information set is $\mathcal{I}_t^i = (a_t^i, x_t^i, z_t)$, keeping in mind that the firm's own TFP a_t^i also serves as a signal of aggregate TFP. As discussed by AL, the noise in common signal v_t adds a second source of aggregate fluctuations, which is interpreted as a common expectation error. The firm's optimal quantity choices Q_t^i in stage 1 maximize $\mathbb{E}(\lambda_t (P_t^i Q_t^i - W_t^i N_t^i) | \mathcal{I}_t^i)$ subject to (1)-(4). An equilibrium consists of optimal quantity choices $Q_t^i(\mathcal{I}_t^i)$, taking as given the aggregate output dynamics implied by the other firms' quantity choices. Taking logs of the FOCs and normalizing constants that do not affect the adjustment to zero, the equilibrium rule for log-quantities $q(\mathcal{I}_t^i)$ solves

$$q(\mathcal{I}_t^i) = k \cdot a_t^i + \alpha \cdot \mathbb{E}(q_t | \mathcal{I}_t^i) \quad (5)$$

$$\text{where } \alpha = \frac{1 - \gamma\rho}{1 + \rho(\theta^{-1}(1 + \varepsilon) - 1)} \text{ and } k = \frac{\theta^{-1}(1 + \varepsilon)}{\rho^{-1} + \theta^{-1}(1 + \varepsilon) - 1}. \quad (6)$$

³One can interpret the first stage as a labor market in which firms commit to their hiring decisions, and the second stage as a product market, in which output is produced and sold to the household.

⁴Throughout I write $x = \log X$ for any variable X . Moreover, I assume that all aggregate shocks and signals are i.i.d. over time, and idiosyncratic shocks and signals are i.i.d. over time and across firms.

The firm's optimal quantity is therefore increasing in the firm's own productivity, and it depends on the firm's expectation of aggregate quantities through the complementarity/substitutability parameter α . The firms' quantity choices are strategic complements ($\alpha > 0$), whenever the firms' inverse demand elasticity exceeds the elasticity of wages to aggregate output, or when a firm's own prices are more sensitive to a change in aggregate output than its own labor costs. This captures the insight of AL that trade linkages across products or sectors (captured by a low ρ) generate complementarities in quantity choices. The absolute value of α is decreasing in the firms' marginal cost elasticity $\theta^{-1}(1 + \varepsilon) - 1$.

These comparative statics should not come as a big surprise. Firms seek to equate their mark-up to the inverse demand elasticity. If $\gamma = \rho^{-1}$, the firms' marginal costs and revenues have the same elasticity w.r.t. aggregate output, so that the firms' mark-ups do not depend on aggregate output, implying neither complementarity nor substitutability. When $\gamma > \rho^{-1}$, marginal cost is more elastic than marginal revenue. An increase in expected aggregate output raises marginal costs by more than marginal revenues, so that firms will want to reduce quantities to increase their relative price and offset the mark-up change. With $\gamma < \rho^{-1}$, we find the opposite.

We can compare α with the pricing complementarity in New Keynesian models of nominal price adjustment. In Hellwig and Venkateswaran (2009), the pricing complementarities are given by⁵

$$\alpha_P = \frac{(\rho - \gamma^{-1})(\theta^{-1}(1 + \varepsilon) - 1)}{1 + \rho(\theta^{-1}(1 + \varepsilon) - 1)}. \quad (7)$$

Therefore, pricing decisions are complementary, whenever quantity decisions are substitutes (and vice versa), and the complementarity or substitutability becomes stronger as returns to scale are more strongly decreasing. Other comparative statics w.r.t. γ and ρ are also reversed. Thus the parameter values of standard New Keynesian models that give rise to strong pricing complementarities also lead to substitutability in quantity choices, while those parameter values that favor complementarities in quantity choices are also make pricing decisions strategic substitutes.⁶ In the pricing model, γ^{-1} measures the elasticity of aggregate demand with respect to aggregate prices. Any firm's demand is therefore increasing/constant/decreasing in aggregate prices, if and only if $\rho \gtrless 1/\gamma$, and with $\theta^{-1}(1 + \varepsilon) > 1$, increases in aggregate demand raise marginal costs more than marginal revenues, increasing the firm's optimal prices. When $\rho \geq 1/\gamma$, aggregate price increases thus raise output and optimal prices for individual firms, while the opposite occurs with $\rho \leq 1/\gamma$.

⁵Their model studies nominal adjustment with heterogeneous information. It is based on the same micro-foundations of preferences and technologies, but augmented with a demand for real balances which satisfies $M_t/P_t = Q_t^\gamma$. They also assume that $\log M$ follows a random walk and focus on the real effects of nominal shocks.

⁶See Woodford (2003b) for an extensive discussion of these parameter choices in New Keynesian pricing models.

Under complete information, the equilibrium satisfies $q_t = k/(1 - \alpha) \cdot a_t$. The model thus captures the standard RBC channel of propagation from technology to employment through the parameter $k/(1 - \alpha)$. The complementarity α decomposes this response into a direct response to firm-specific productivity increases and an indirect response to aggregate conditions coming from trade linkages. With $\alpha = 0$ or homogeneous information this indirect channel doesn't play a role for aggregate dynamics, but it becomes important once information is heterogeneous and $\alpha \neq 0$.

Main results: By guessing and verifying we characterize the equilibrium as

$$q(\mathcal{I}_t^i) = ka_t^i + \frac{k\alpha}{1 - \alpha} \frac{(1 - \alpha)(\kappa_\eta a_t^i + \kappa_x x_t^i) + \kappa_z z_t}{(1 - \alpha)(\kappa_\eta + \kappa_x) + \kappa_z + \kappa_a}. \quad (8)$$

The firm's expectations of aggregate TFP and aggregate output are given by

$$\mathbb{E}(a_t | \mathcal{I}_t^i) = \frac{\kappa_\eta a_t^i + \kappa_x x_t^i + \kappa_z z_t}{\kappa_\eta + \kappa_x + \kappa_z + \kappa_a} \text{ and } \mathbb{E}(q_t | \mathcal{I}_t^i) = \frac{k}{1 - \alpha} \cdot \frac{(1 - \alpha)(\kappa_\eta a_t^i + \kappa_x x_t^i) + \kappa_z z_t}{(1 - \alpha)(\kappa_\eta + \kappa_x) + \kappa_z + \kappa_a}. \quad (9)$$

The firms' forecasts of aggregate activity thus attach lower weights to private signals and higher weights to the prior and the public signal than their forecasts of productivity. This shift results because common information is relatively more useful in forecasting aggregate output: While public and private signals are used in proportion to their precision to forecast the response of output decisions to private signals, private signals play no role in forecasting the output response to public signals, which by its nature is common knowledge among all firms. Aggregate output is

$$q_t = \left(\frac{k}{1 - \alpha} - \phi_a \right) a_t + \phi_z v_t \quad (10)$$

$$\text{where } \phi_a = \frac{k\alpha}{1 - \alpha} \frac{\kappa_a}{(1 - \alpha)(\kappa_\eta + \kappa_x) + \kappa_z + \kappa_a} \text{ and } \phi_z = \frac{k\alpha}{1 - \alpha} \frac{\kappa_z}{(1 - \alpha)(\kappa_\eta + \kappa_x) + \kappa_z + \kappa_a}$$

Relative to full information, incomplete, heterogeneous information reduces the output response to technology shocks by ϕ_a . In addition, output responds to expectational noise by ϕ_z . The parameters ϕ_a and ϕ_z are increasing in α , and, for a given overall signal precision $\kappa_\eta + \kappa_x + \kappa_z$, raising $\kappa_\eta + \kappa_x$ and reducing κ_z amplifies the departures from the full information benchmark. Moreover, these two effects are mutually reinforcing, which underlies the idea that heterogeneity and complementarities jointly reinforce the effects of noise. The full information benchmark obtains only if $\alpha = 0$, in which case firms only respond to their own productivity, which is fully observed.

Let us further consider the response of employment and the labor wedge in this model:⁷

$$n_t = \theta^{-1}(q_t - a_t) = \theta^{-1} \left(\frac{k}{1 - \alpha} - \phi_a - 1 \right) a_t + \theta^{-1} \phi_z v_t \quad (11)$$

⁷See Chari, Kehoe, and McGrattan (2007) for a discussion on the role of these variables in accounting for business cycles in the data.

$$\tau_t^N = \log\left(\frac{MPL_t}{MRS_t}\right) = K\left(\frac{k}{1-\alpha}a_t - q_t\right) = K(\phi_a a_t - \phi_z v_t), \quad (12)$$

where $K = \gamma + \theta^{-1}(1 + \varepsilon) - 1 > 0$. AL discuss how this characterization leads to interesting departures from the canonical RBC model:

Proposition 1 (Main Results of AL): (i) *With $\alpha > 0$, information heterogeneity mutes the response of output to technology shocks. Technology shocks may even induce negative co-movement of technology and employment, but they generate positive co-movement with the labor wedge.*

(ii) *Informational shocks provide a source of positive co-movement between output and hours, and negative co-movement with the labor wedge. Depending on parameters, informational shocks can account for any fraction of the variance decompositions of these variables.*

These shifts summarize the main business cycle implications of heterogeneous information mentioned in the beginning. The shift towards the common prior explains why heterogeneous beliefs and complementarities amplify adjustment delays. The shift towards public signals explains the excess sensitivity to public information and common expectation shocks.

Thus, as emphasized by Proposition 4 in AL, output may not respond to TFP, even if firms have very precise information: if α is close to 1, $\kappa_\eta + \kappa_x$ is large, but $(1 - \alpha)(\kappa_\eta + \kappa_x)$ and κ_z are small compared to κ_a , ϕ_a may be close to $k/(1 - \alpha)$ and the response of output to TFP close to zero. The labor wedge on the other hand is constant at zero under full information and becomes more pro-cyclical the larger is the departure from full information: the under-adjustment to TFP under incomplete information raises the marginal product relative to the marginal rate of substitution.

Finally, if public signals receive significant weight in decisions, signal noise generates significant positive co-movement in labor and output without any change in technology, inducing countercyclical fluctuations in the labor wedge: false good news about technology raise output, lowering the marginal product, and raising the marginal rate of substitution, and hence the labor wedge.

2 Comments

Comments about the Theory. AL suggest that these results may help account for empirical results that appear to be inconsistent with a canonical RBC model. VAR estimates suggest that hours respond negatively on impact to technology shocks, and technology shocks only account for a small fraction of short-run fluctuations. Chari, Kehoe and McGrattan (2007) argue that countercyclical movements in the labor wedge are important in accounting for business cycles. The

logic of Bayesian updating however imposes restrictions on the ability of informational noise to account for these features and serve as an independent source of fluctuations more generally.

Remark 1: *The unconditional covariance between the labor wedge and output is always positive:*

$$\text{cov}(\tau_t^N, q_t) = \alpha K \left(\frac{k}{1-\alpha} \right)^2 \frac{(1-\alpha)(\kappa_\eta + \kappa_x + \kappa_z + \kappa_a)}{((1-\alpha)(\kappa_\eta + \kappa_x) + \kappa_z + \kappa_a)^2} > 0.$$

The model thus fails to account for the unconditionally counter-cyclical fluctuations in the labor wedge that, according to Chari-Kehoe and McGrattan are a key feature of business cycle fluctuations. Notice that the pro-cyclicality of labor wedges results even if most of the fluctuations in the labor wedge are driven by noise shocks.

There is a more general sense in which incomplete information about productivity appears to push the RBC model in the ‘wrong’ direction, with respect to the data: incomplete information reduces the importance of TFP for short-run output fluctuations, but it does so at the expense of a pro-cyclical labor wedge: The under-reaction of output to TFP translates into pro-cyclical responses in the labor wedge to TFP. Expectational noise adds some counter-cyclicality to labor wedges, along with more pro-cyclical employment, but this is not sufficient to offset the pro-cyclical TFP effects. The model is thus not able to simultaneously account for counter-cyclical labor wedges and for a reduced short-run role of TFP, suggesting perhaps other fundamental or preference shocks as a more fruitful avenue to pursue, separately or in conjunction with information shocks.

This tension between the importance of TFP and the pro-cyclicality of the labor wedge is a direct consequence of Bayesian updating. Bayesian updating links the strength of negative co-movement in the labor wedge from informational shocks to the positive co-movement from productivity shocks: firms respond to the public signals and hence the informational noise, only to the extent that they believe that these signals may also reflect the true fundamentals. For the same reason, the model also imposes restrictions on expectational noise as an independent source of fluctuations: if $\kappa_z \approx 0$, public signal noise is large, but irrelevant because firms don’t respond to public signals ($\phi_z \approx 0$). If instead, κ_z is large, firms respond to public information, but the noise in the information is small, and hence again irrelevant. Public signal noise has the biggest effects at intermediate levels, and their overall size is also tightly linked to the size of TFP shocks.

Remark 2: *The negative co-movement of employment and TFP requires k to be small:*

$$\text{cov}(a_t, n_t) < 0 \iff 1 > \frac{k}{1-\alpha} \left(1 - \alpha \frac{\kappa_a}{(1-\alpha)(\kappa_\eta + \kappa_x) + \kappa_z + \kappa_a} \right) > k.$$

The negative responses of hours to technology thus require firms to reduce hours in response to an increase in their technology ($k < 1$), even if they attribute this technological change to

idiosyncratic factors. A negative co-movement of hours and TFP in the aggregate thus also implies negative co-movement of hours and technology in the cross-section of firms.

Remark 3: *Are the results due to informational heterogeneity or informational noise?*

The paper's main results above all still hold when $\kappa_\eta = \kappa_x = 0$, i.e. when all information is common. Thus, they are not a consequence of information heterogeneity per se. What is more, the absolute values of ϕ_a and ϕ_z are decreasing in κ_η and κ_x . For a given size of aggregate shocks and a given informativeness of public signals, adding informative private signals or reducing the variance of idiosyncratic technology shocks therefore reduces the departure from the full information benchmark. In other words, for a given magnitude of aggregate shocks and common information, the effects of informational noise are largest, when there are no additional private signals, and hence no informational heterogeneity.

This suggests noise or incomplete information rather than information heterogeneity as the main driving force. It also suggests that it is a matter of interpretation whether information heterogeneity amplifies or mutes the quantitative importance of noise: for a given quality of *public* information, heterogeneity in the form of additional, or more precise, private signals mutes the effects of noise, while for a given quality of *overall* information, heterogeneity in the form of a shift from public to private information increases the effects of noise. This discussion of contrasting comparative statics results therefore remains essentially meaningless in the absence of a clear empirical counterpart.

Early incarnations of RBC models explored incomplete information as a potential source of amplification, but they soon chose to abstract from it, in part because they found that it did not generate a substantial departure from the simpler full information benchmark. The renewed interest in incomplete information is to a large extent based on the premise that adding heterogeneity increases the effects of informational noise. However, if the common information models were already deemed too close to the full information benchmark, and adding heterogeneous information brings them even closer, this raises the question whether this new class of models will be any more successful in accounting for important features of the business cycles.

The authors rightly point out that this observation may not be robust to more general specifications of the information structure that do not restrict attention to purely private and purely public signals, and allow for more general correlation patterns of expectations across agents. Corroborating this conjecture or providing proof to the contrary will be an important step in the future development in this literature, because it will contribute to a better understanding of the separate effects of noise and heterogeneity. This separation is central to determining whether heterogeneous information models have any hope of providing quantitative implications that are substantially

different from their common information ancestors.

Quantitative Challenges. My comments about the theory all point to the need for a tighter connection between theoretical results and quantitative implications. Let me now briefly discuss the challenges that arise if one is to subject this class of models to a quantitative evaluation.

The central challenge in any quantitative evaluation is how one picks the structural parameters underlying trade linkages, as well as the parameters that determine informational incompleteness and heterogeneity. One natural source of discipline comes from the model’s cross-sectional implications for output, prices, profits, and productivity. The authors question the usefulness of this type of information, in part because a ‘firm’ and an ‘island’ have different interpretations in the model, and enrichments of the model allow for cross-sectional movements that are driven by parameters unrelated to aggregate fluctuations. As an alternative, they propose to infer informational parameters and complementarities using data on survey expectations about aggregate activity.

In contrast to the authors, I view the model’s parsimony, which derives rich cross-sectional as well as aggregate implications from only a minimal set of parameters as one of its main virtues. To illustrate the power of the cross-sectional moments to discipline parameter choices, consider what parameters can account for important trade linkages (high α) and explain large firm-specific output fluctuations (high k). The definitions of these parameters reveal a tension between these targets, as a higher ρ or $\theta^{-1}(1 + \varepsilon)$ increases k , but reduces α .

Table 1	AL, $\gamma = 0.2$		HR, $\gamma = 0.2$		AL, $\gamma = 0.05$		HR, $\gamma = 0.05$	
ρ	α	k	α	k	α	k	α	k
10	-0.06	1.56	-1	10	0.03	1.32	0.5	10
5	0	1.47	0	5	0.09	1.27	0.75	5
2	0.15	1.25	0.6	2	0.23	1.15	0.9	2
1	0.32	1	0.8	1	0.38	1	0.95	1
0.5	0.51	0.71	0.9	0.5	0.56	0.79	0.98	0.5
0.2	0.74	0.38	0.96	0.2	0.76	0.48	0.99	0.2
0.1	0.85	0.22	0.98	0.1	0.87	0.29	1	0.1

The authors fix an income elasticity of wages of $\gamma = 0.2$, a labor share of $\theta = 0.6$, and a Frisch elasticity of $\varepsilon^{-1} = 2$, and they alter ρ to vary the degree of complementarity α . Table 1 compares the implications of these values for α and k with alternative parameter specifications (“Hansen-Rogerson”, HR), in which there are no decreasing returns ($\theta = 1$) and linear disutility of labor ($\varepsilon = 0$). We also vary the income elasticity of wages from $\gamma = 0.2$ to $\gamma = 0.05$. As this table shows,

AL's parameter choices generate high complementarities only at the expense of a small value of k . Hansen-Rogerson preferences make α and k more sensitive to changes in ρ , but on their own do not resolve this issue. Instead, high values of α and k can jointly be sustained with Hansen-Rogerson preferences and very low income effects of wages. Interestingly, this does not require implausible values of firm-level demand elasticities, with values of ρ between 2 and 5.

This exercise raises the issue of the right level of aggregation, or the right empirical counterpart of an 'island' in the AL model. Since the islands define informational separation between localized markets, it may be natural to interpret them as separate firms, as the firm represents a natural informational boundary to outsiders who do not have immediate access to the same proprietary information about technology and products, or accounting information about their demand, sales and costs. This view was implicitly adopted when the within and across island elasticities were set equal in this discussion. HR preferences with small wealth effects are consistent with this interpretation, which also implies positive co-movement between output, productivity and employment in the cross section.

An alternative view interprets each island as a separate sector. This view is more consistent with AL's parameter values, with lower values of ρ and consequently k , and higher demand elasticities within sectors. Low values of k also imply that labor moves out of sectors with higher productivity and into sectors with lower productivity. The sectorial interpretation raises the question of what causes the informational heterogeneity between sectors but not between firms within a sector, and why the former should receive much more weight than the latter.

I conclude by highlighting the sensitivity of theoretical and quantitative results to informational assumptions, and in particular the role of pre-commitment to quantity choices and the non-availability of information from market transactions. For this purpose, consider a small variation of the main model, in which firms are able to condition stage 1 choices on their stage 2 cash flows, or equivalently the prices and wages in stage 2. To make things interesting, assume that there are firm-specific shocks Z_t^i and B_t^i to demand and labor supply $Q_t^i = B_t^i Q_t (P_t^i)^{-\rho}$ and $W_t^i = Z_t^i Q_t^\gamma (N_t^i)^\varepsilon$; otherwise, the observation of prices and wages would trivially reveal aggregate output, with full adjustment as an immediate consequence. The next proposition shows that even with these firm-specific shocks, the equilibrium exhibits complete adjustment, irrespective of the degree of informational incompleteness or heterogeneity.

Proposition 2 *Suppose that $\mathcal{I}_t^i = (a_t^i, x_t^i, z_t; p_t^i, w_t^i)$. Then, the equilibrium exhibits full adjustment. Moreover noisy signals of aggregate technology (x_t^i, z_t) are irrelevant for adjustment decisions.*

The logic behind this proposition is the following: to sustain the full information outcome, it suffices that firms have enough information to fully infer their optimal full information quantities. This requires much less than full information about idiosyncratic and aggregate state variables. Remarkably, markets provide firms with just that information, because they enable firms to infer their marginal costs and revenues from the prices and wages, eliminating the need to form expectations about TFP or aggregate activity. Equilibrium decisions thus mimic the full information benchmark, even though the firms' actual information about fundamentals may be very limited.

The result also holds as an approximation, if firms observe current market conditions with noise. The departure from the full information equilibrium, and the role of signals about aggregate conditions becomes more important, the more uncertain firms are about their own market environment, and hence their own optimal actions.

Ultimately, this result reflects an old insight going back to Hayek (1945): markets parsimoniously convey all relevant information to those making the decisions. Firms don't need to know whether prices and costs change for idiosyncratic or aggregate reasons; they only need to be able to respond to these changes in their own decisions.

The result has several important implications. First, it highlights an important necessary condition for the effects of informational incompleteness and heterogeneity to affect equilibrium outcomes: departures from the full information outcome arise only if firms face uncertainty about their own optimal decisions; if they can accurately infer the latter, incomplete information about aggregate conditions are largely irrelevant for adjustment dynamics. This also suggests that survey expectations about aggregate conditions alone are unlikely to be sufficient for disciplining the quantitative effects of information heterogeneity, if they are not coupled with additional information on how uncertain firms are about their own market environment.

Second, it serves as a benchmark for discussing what factors determine the quantitative importance of information heterogeneity. It seems natural to assume that firms have access to and would want to use the information conveyed by their own market transactions, especially since this information parsimoniously and accurately conveys everything the firms need to know. Even if firms are not able to use this information instantaneously, they presumably will do so eventually. This delay is captured by the pre-commitment assumption. But then, the relevant issue is how well information from past market transactions is able to forecast current or future market behavior. This depends on how persistent the underlying idiosyncratic and aggregate shocks are, how fast firms learn about them, and how fast they are able to respond to new information. The first two are properties of the underlying stochastic processes, while the last is determined by the model's

notion of a period length.

Hellwig and Venkateswaran (2009) discuss these issues in a model of nominal price adjustment. They calibrate the afore-mentioned model of nominal price adjustment to match micro data on price and quantity fluctuations. If pre-commitment to pricing choices vanishes and firms can respond to market information at very high frequency, prices adjust rapidly to changes in aggregate conditions, even though the firms are not aware of these aggregate changes. This quantitative result follows from the observations that firms face large idiosyncratic shocks in response to which they adjust prices. These shocks are reflected in market data along with aggregate shocks, and hence confounded with the latter. When an aggregate shock occurs, firms adjust their nominal prices thinking that they have been hit by an idiosyncratic shock. Even if firms revisit pricing decisions once a week or once a month, prices absorb a significant fraction of nominal shocks within only a few weeks.

These high-frequency assumptions may be less appealing for a model of quantity choices or hiring and production plans, which often require firms to look ahead at future market conditions, or follow specific seasonal cycles.⁸ Nevertheless, the model cannot escape the broader insight that its quantitative implications are going to be sensitive to the speed with which firms are able to adjust their production and hiring plans to information from market transactions (such as the arrival of new orders of production, or the purchases and prices of new raw materials or other inputs), and the degree to which this information is useful for forecasting future market conditions. Whether this delay is viewed as a technological lead time in production, or as resulting from the firms' inability to process and direct the relevant information to the right people, its source is tied to the internal organization of the firms or frictions in the decision process and not to the external lack of information. Inevitably a quantitative evaluation of the model requires a better understanding of these issues and therefore of the firms' internal decision processes.

In conclusion, informational incompleteness or heterogeneity can affect aggregate outcomes only if firms are not able to respond to information contained in market outcomes in real time. The pre-commitment assumption that captures this requirement introduces an additional adjustment friction, the effects of which are difficult to separate from pure lack of information. Any quantitative evaluation needs to take a stand on the importance of this additional friction - otherwise, the model's aggregate implications remain unidentified.

⁸On the other hand, Zipf's Law for firms suggests that firm-specific real shocks are highly persistent, which improves the inference drawn from past market conditions for future choices, even into the distant future.

3 Conclusion

The paper by AL presents a very elegant model of real business cycles with heterogeneous beliefs and informational noise. The model's appealing aggregate and cross-sectional implications result from trade linkages across firms or sectors and the need to forecast aggregate activity.

This discussion has suggested several open questions and areas that appear promising avenues for future developments. Central among those is the development of a consistent strategy for confronting the theoretical and quantitative implications of this class of models with their counterparts in the data. Among other things, this will require further development of truly dynamic business cycle models with heterogeneous information, along with a more dynamic view as to what information firms have access to and how this information enters their decision-making process.

On a few issues, my interpretation of the results differs quite significantly from the authors. Most importantly, I view the results as more sensitive to informational assumptions, especially on the quantitative side, and I do not think that information from economic forecasts alone will be sufficient to discipline the relevant parameters of the model. My emphasis on these differences should not be viewed as criticism, but rather as a contribution to the broader debate, in the hope that this will stimulate further research on these issues. These differences should also not obscure my agreement with the authors that heterogeneous information seems like a very fruitful and important avenue for gaining a better understanding of business cycle dynamics.

References

- [1] Angeletos, George-Marios, and Jennifer La'O (2009), "Noisy Business Cycles", 2009 NBER Macroeconomics Annual (this volume)
- [2] Angeletos, George-Marios, and Alessandro Pavan (2004), "Transparency of Information and Coordination in Economies with Investment Complementarities," *American Economic Review, Papers and Proceedings*, 94, 91-98.
- [3] Angeletos, George-Marios, and Alessandro Pavan (2007), "Efficiency Use of Information and Social Value of Information," *Econometrica*, 75(4), 1103-1142.
- [4] Chari, V. V., Patrick Kehoe and Ellen McGrattan (2007), "Business Cycle Accounting," *Econometrica*, 75 (3), 781-836.

- [5] Hayek, Friedrich (1945), "The Use of Knowledge in Society," *American Economic Review*, 35 (4), 519-530.
- [6] Hellwig, Christian (2005), "Heterogeneous Information and the Welfare Effects of Public Information Disclosures," discussion paper, U.C.L.A.
- [7] Hellwig, Christian (2008), "Monetary Business Cycles (Imperfect Information)," in: *The New Palgrave Dictionary of Economics*, 2nd edition, Eds. Steven N. Durlauf and Lawrence E. Blume, London: Palgrave-Macmillan.
- [8] Hellwig, Christian and Venky Venkateswaran (2009), "Setting the Right Prices for the Wrong Reasons," *Journal of Monetary Economics*, forthcoming.
- [9] R. Lucas, Jr. (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4, 103-124.
- [10] S. Morris and H. S. Shin (2002), "The Social Value of Public Information," *American Economic Review* 92, 1521-1534.
- [11] Woodford, Michael (2003a), "Imperfect Common Knowledge and the Effects of Monetary Policy," in: *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Eds. P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, Princeton: Princeton University Press.
- [12] Woodford, Michael (2003b), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.