Real Rigidity, Nominal Rigidity, and the Social Value of Information

American Economic Review 2016, 106(1): 200–227
http://dx.doi.org/10.1257/aer.20110865

Real Rigidity, Nominal Rigidity, and the Social Value of Information†

By George-Marios Angeletos, Luigi Iovino, and Jennifer La'O*

Does welfare improve when firms are better informed about the state of the economy and can thus better coordinate their production and pricing decisions? We address this question in an elementary business-cycle model that highlights how the dispersion of information can impede both kinds of decisions and, in this sense, be the source of both real and nominal rigidity. Within this context we develop a taxonomy for how the social value of information depends on the two rigidities, on the sources of the business cycle, and on the conduct of monetary policy. (*JEL* D21, D82, D83, E32, E52)

Economic agents have access to various sources of information about the state of the economy, some of which are private (such as the signals each firm extracts from its own market interactions) and some of which are public (such as macroeconomic statistics and central bank communications). By informing each agent about the activity of others, public information can ease coordination, whereas private information can hinder it. In this paper, we study the welfare consequences of this mechanism within the context of a micro-founded business-cycle model in which firms make their employment, production, and pricing choices under incomplete information about one another's choices and about the state of the economy.

Background.—We are not the first to study how information affects coordination and welfare. In an influential article, Morris and Shin (2002) used a "beauty contest" game—a linear-quadratic game in which actions are strategic complements—to formalize the coordinating role of public information and to study its welfare implications. In such a game, public signals have a disproportionate effect on equilibrium outcomes relative to what is warranted on the basis of their informational content regarding fundamentals alone. This is due to the fact that the players use such signals not only to predict fundamentals but also to coordinate their actions. In this

*Angeletos: Department of Economics, Massachusetts Institute of Technology, 77 Massachusetts Avenue E18-218, Cambridge, MA 02139, and NBER (e-mail: angelet@mit.edu); Iovino: Department of Economics, Bocconi University, via Roentgen 1, 20136 Milan, Italy, and IGIER (e-mail: luigi.iovino@unibocconi.it); La'O: Department of Economics, Columbia University, International Affairs Building, 420 West 118th Street, New York, NY 10027, and NBER (e-mail: jenlao@columbia.edu). This paper extends and subsumes an earlier draft that circulated as NBER Working Paper 17229, under the title "Cycles, Gaps, and Social Value of Information." The authors are grateful to five anonymous referees for extensive feedback. The authors also declare that they have no relevant or material financial interests that relate to the research described in this paper.
†Go to http://dx.doi.org/10.1257/aer.20110865 to visit the article page for additional materials and author disclosure statement(s).

regard, public signals can play a role akin to that of sunspots, possibly contributing to higher volatility and lower welfare.

Because strategic complementarity emerges naturally from the aggregate demand externalities that are embedded in macroeconomic models, Morris and Shin's analysis was used to inform the debate on the pros and cons of central bank transparency.[1] However, subsequent work raised questions about the validity of applying Morris and Shin's lessons to a macroeconomic context.

Using a richer game-theoretic framework, Angeletos and Pavan (2007) highlighted that Morris and Shin's welfare conclusion hinges on the assumption that coordination is socially harmful—an assumption that need not be valid in workhorse macroeconomic models. Reinforcing this observation, a line of applied work that includes Baeriswyl and Cornand (2010); Hellwig (2005); Lorenzoni (2010); Roca (2005); and Walsh (2007) found different welfare effects than those suggested by Morris and Shin in variants of the New-Keynesian model in which nominal rigidity originates from incomplete information rather than Calvo-like sticky prices.

This applied work, some of which we revisit in Section VI, has pushed the analysis of the question of interest from abstract games to workhorse macroeconomic models. This is a crucial step, as "anything goes" without the discipline of specific micro-foundations: different assumptions about the payoff structure of a game can justify any sort of welfare effect.

Yet, this work faces certain limitations. By equating the informational friction to a particular form of nominal rigidity, it abstracts from the bite that the former can have on productive efficiency *regardless* of nominal rigidity. Formally, it lets the informational friction impose a measurability constraint on nominal prices, but abstracts from any such constraint on real quantities. It is thus as if employment and production choices, in contrast to pricing choices, are made under complete information. Furthermore, this work intertwines the welfare effects of information with those of particular monetary policies, often confounding the informational incompleteness of the firms with frictions in the conduct of monetary policy—a point that we formalize and clarify in due course.

Our Contribution.—Seeking to overcome these limitations, in this paper we consider a framework in which firms make not only their pricing choices but also certain employment and production choices on the basis of dispersed noisy information about the underlying aggregate shocks. In this sense, we allow the incompleteness of information to be the source of both *real* and *nominal* rigidity, that is, to impose a measurability constraint on both quantities and prices.[2]

In addition, we dissect how the welfare effects of information depend on whether monetary policy coincides with or deviates from two important policy benchmarks. The first corresponds to a policy that replicates flexible prices, in the sense of implementing the same allocation as the one that would have obtained in the absence

[1] See, e.g., the follow-up articles in this journal by Svensson (2006); Morris, Shin, and Tong (2006); and James and Lawler (2011).

[2] Our notion of *real rigidity* differs from the one typically used in the New-Keynesian literature. In that literature, the term refers to the lack of response in a firm's desired relative price to aggregate disturbances due to technology, preferences, or market power. In our paper, the term instead refers to the lack of response in a firm's real quantity to aggregate disturbances due to incomplete information.

of the nominal rigidity. The second identifies the unconstrained optimal monetary policy, meaning the solution to the Ramsey problem in which the planner can set the nominal interest rate as an arbitrary function of the underlying state of nature.

This approach permits us to develop a certain taxonomy for how the answer to the question of interest depends on the conduct of monetary policy, on the nature of the underlying business-cycle disturbances, and on the aforementioned two types of rigidity.[3] In so doing, we also qualify some of the lessons that have appeared in the literature.

Isolating the Real Rigidity.—In the first part of the paper (Sections I–III), we study the case polar opposite to the one considered in prior work: we assume that firms choose employment on the basis of incomplete information, thus accommodating real rigidity, but we let prices adjust freely to the realized state. We thus abstract from nominal rigidity and thereby shut down the pivotal role that monetary policy plays once nominal rigidity is present. This part serves as a stepping stone toward the second part of the paper, which ultimately allows for both types of rigidity.

We first study how information affects two familiar welfare components: the volatility of the aggregate output gap and the inefficient cross-sectional dispersion in relative prices (and quantities). For each component separately, we show how the sign of these effects is governed by three sets of factors: (i) preference and technology parameters that pin down the coordination motives; (ii) whether the information is private or public; and (iii) the underlying business-cycle shocks.

We next show that, despite non-monotone and often conflicting effects on these two components, the sign of the overall welfare effect of *either* type of information is governed *solely* by the sources of the business cycle. When the business cycle is driven by non-distortionary forces such as technology shocks, welfare unambiguously increases with either private or public information. When instead the business cycle is driven by distortionary forces such as shocks to monopoly markups, welfare unambiguously decreases with either type of information.[4]

To some extent, this result is a priori intuitive: in the case of technology shocks, one may expect information to be welfare-improving because the firms' reaction to such shocks is socially desirable, while the converse is true in the case of markup shocks. However, this basic intuition can fall apart when information is dispersed: in Morris and Shin (2002), the equilibrium is first-best efficient when information is commonly shared, resembling what happens in our setting in the absence of monopoly distortions, yet welfare can decrease with the precision of public information when, and only when, information is dispersed. As we explain in due course, the sharpness of our results therefore hinges to the following property of our microfounded setting: the private value that the firms assign to the coordination of their

[3] In this paper, we focus on the distinct *normative* implications of the two types of rigidity. However, the two also have distinct *positive* implications. For example, when the rigidity is nominal, the response of macroeconomic outcomes to the underlying noise shock can take *any* sign, depending on the conduct of monetary policy. Furthermore, there is a Philips curve: any deviation of the level of real output from the complete-information point is *necessarily* associated with a commensurate movement in the price level. None of this is true when the rigidity is real.

[4] In the latter case, a countervailing effect is also at work unless one assumes, as is often done in the literature, that a noncontingent subsidy is used to eliminate the mean (or "steady-state") distortion in economic activity. We characterize this effect in Proposition 4 but, in line with the literature, abstract from it in the rest of our analysis.

choices coincides with the corresponding social value. In the absence of such a coincidence, the welfare effects of either type of information could have been reverted.

Adding Nominal Rigidity.—In the second part of the paper (Section IV), we study a more general framework, in which we let the informational friction impede not only the firms' employment and production choices but also their price-setting behavior. As noted before, we anchor our analysis to two benchmarks that help dissect the role of monetary policy. The first identifies policies that replicate flexible prices; the second identifies the unconstrained Ramsey optimum. As in the baseline New-Keynesian model, these two benchmarks coincide in the case of technology shocks, but not in the case of markup shocks. Importantly, the scenarios studied in the related prior work assume not only the absence of real rigidity but also specific deviations from these benchmarks.

Consider the first benchmark. When monetary policy replicates flexible prices, the question of interest admits the same answer as in our baseline model: welfare increases (respectively, decreases) with either type of information when the business cycle is driven by technology shocks (respectively, markup shocks). Furthermore, information matters at this benchmark *only* because of the real rigidity: in the absence of real rigidity, the flexible-price policy implements the complete-information outcome, irrespective of how noisy the firms' information might be.

Away from this benchmark, an additional effect emerges: information affects not only the bite of the real rigidity but also the firms' ability to forecast, and thus preempt, any action of the monetary authority that attempts to move the economy away from its flexible-price outcomes. The welfare contribution of this additional effect then depends on whether such deviations are socially desirable or not—a question that is directly answered by our policy benchmarks.

When the business cycle is driven by technology shocks, *any* deviation from flexible prices is welfare-deteriorating. Increasing the information that is available to firms may then improve welfare not only by alleviating the real rigidity but also by helping the firms forecast, and undo, the "mistakes" in monetary policy. In this sense, transparency is good.

When, instead, the business cycle is driven by markup shocks or other distortions, an appropriate deviation from flexible prices is desirable, for reasons once again familiar from the New-Keynesian framework: the optimal policy *exploits* the nominal rigidity in order to substitute for a missing tax instrument, namely, the state-contingent subsidy that would have offset the markup shock. More information in the hands of the private sector can then be detrimental for welfare, not only for the reasons highlighted in our baseline model, but also by reducing the effectiveness of monetary policy. In this sense, opacity becomes preferable.

To recap, the taxonomy we develop provides sharp answers to our question under two familiar policy benchmarks, but it also builds a roadmap for understanding the welfare effects of information away from them. As an application of this roadmap, in Section V we revisit the prior contributions of Baeriswyl and Cornand (2010); Hellwig (2005); Lorenzoni (2010); and Walsh (2007), shedding further light on the key mechanisms in these papers, qualifying some of the policy lessons, and facilitating a certain synthesis.

Layout.—Sections I–III study the baseline model. Section IV analyzes the role of monetary policy. Section V revisits the related literature. Section VI concludes. All the proofs are delegated to online Appendices A and B.

I. The Baseline Model

The baseline model builds on Angeletos and La'O (2010). The economy consists of a "mainland" and a continuum of "islands." Each island is inhabited by a continuum of workers and a continuum of monopolistic firms. Firms employ local workers through a competitive labor market and produce differentiated commodities, which they ultimately sell in a centralized market in the mainland. The latter is inhabited by a continuum of consumers, each of whom is tied to one worker and one firm from *every* island in the economy. Along with the fact that there will be no heterogeneity *within* islands, this guarantees that the economy admits a representative household: we can think of the latter as a "big family" that is comprised of all agents, collects all income, and consumes all output in the economy. Nevertheless, this geography introduces an informational friction: we assume that firms and workers observe the fundamentals on their own island, but face incomplete information about the underlying aggregate shocks and the choices that other agents (their "siblings") make on other islands. Finally, islands are indexed by $i \in I = [0,1]$; firms, workers, and commodities by $(i,j) \in I \times J = [0,1]^2$; and periods by $t \in \{0,1,2,\ldots\}$.

A. *Fundamentals*

The utility of the representative household is given by

$$\mathcal{U} = \sum_{t=0}^{\infty} \beta^t \Big[U(C_t) - \int_I \int_J \chi_{it} V(n_{ijt})\, dj di \Big],$$

where $U(C) = \frac{1}{1-\gamma} C^{1-\gamma}$, $V(n) = \frac{1}{1+\epsilon} n^{1+\epsilon}$, and $\gamma, \epsilon \geq 0$. Here, n_{ijt} is the labor input in firm j of island i (or the effort of the corresponding worker), χ_{it} is an island-specific shock to the disutility of labor, and C_t is aggregate consumption. The latter is given by the following nested CES structure:

$$C_t = \left[\int_I (c_{it})^{\frac{\rho-1}{\rho}} di \right]^{\frac{\rho}{\rho-1}}, \quad \text{with} \quad c_{it} = \left[\int_J (c_{ijt})^{\frac{\eta_{it}-1}{\eta_{it}}} dj \right]^{\frac{\eta_{it}}{\eta_{it}-1}} \forall i,$$

where c_{ijt} denotes the consumption of commodity j from island i, c_{it} represents a composite of all the goods of island i, and ρ and η identify the elasticities of substitution, respectively, across and within islands. In equilibrium, ρ ends up controlling the strength of aggregate demand externalities, while η controls the degree of monopoly power. We let $\eta \neq \rho$ so as to isolate the distinct roles of these two forces; we then let η_{it} be random so as to accommodate markup, or cost-push, shocks.

Recall that the representative household receives labor income and profits from all islands in the economy. Its budget constraint is thus given by the following:

$$\int_I \int_J p_{ijt} c_{ijt}\, djdj + B_{t+1} \;\leq\; \int_I \int_J \pi_{ijt}\, didj + \int_I (1 - \tau_{it})\, w_{it} n_{it}\, di + (1 + R_t) B_t + T_t .$$

Here, p_{ijt} is the period-t price of the commodity produced by firm j on island i, π_{ijt} is the period-t nominal profit of that firm, w_{it} is the period-t nominal wage on island i, R_t is the period-t nominal net rate of return on the riskless bond, and B_t is the amount of bonds held in period t.

The variables τ_{it} and T_t satisfy $T_t = \int_I \tau_{it} w_{it} n_{it}\, di$. One can thus interpret τ_{it} as an island-specific distortionary tax and T_t as the lump-sum transfers needed to balance the budget. Alternatively, we can consider a variant of our model with monopolistic labor markets as in Blanchard and Kiyotaki (1987), in which case τ_{it} could re-emerge as an island-specific markup between the wage and the marginal revenue product of labor. In line with much of the DSGE literature, we can thus introduce exogenous variation in $1 - \tau_{it}$ and interpret this variation as shocks to the "labor wedge."

Finally, the output of firm j on island i during period t is given by

$$y_{ijt} = A_{it} n_{ijt},$$

where A_{it} is the island-specific TFP, and the firm's realized profit is given by $\pi_{ijt} = p_{ijt} y_{ijt} - w_{it} n_{ijt}$.

B. *Information Structure*

Different authors have motivated informational frictions on the basis of either market segmentation (Lucas 1972; Lorenzoni 2009; Angeletos and La'O 2013) or some form of inattention (Sims 2003; Mankiw and Reis 2002; Woodford 2002; Mackowiak and Wiederholt 2009). In either case, the key friction is an agent-specific measurability constraint, reflecting the dispersed private information upon which certain economic decisions are conditioned. In this paper we wish to understand the welfare effects of relaxing this constraint, not its possible micro-foundations. Furthermore, we seek to isolate the information about aggregate, as opposed to idiosyncratic, shocks, because it is only the former that have nontrivial general-equilibrium effects.

With these points in mind, we assume that the firms and workers of any given island know the local fundamentals, but have incomplete information about the aggregate state of the economy. We then model the available information as a combination of private and public signals and proceed to characterize equilibrium welfare as a function of the precisions of these signals. The details, and a justification, are provided in Section III. For now, we note that the results of Section II use only the weaker assumption that the stochastic structure is Gaussian.

II. Equilibrium, Welfare, and Coordination

The equilibrium is defined in a familiar manner: prices clear markets and quantities are (privately) optimal given the available information. Following the same steps as in Angeletos and La'O (2010),[5] one can show that equilibrium output is pinned down by the following fixed-point relation:

$$\chi_{it} V'\left(\frac{y_{it}}{A_{it}}\right) = \frac{1}{\mathcal{M}_{it}} \mathbb{E}_{it}\left[U'(Y_t)\left(\frac{y_{it}}{Y_t}\right)^{-\frac{1}{\rho}}\right] A_{it}, \tag{1}$$

where $\mathcal{M}_{it} \equiv \frac{1}{1-\tau_{it}} \frac{\eta_{it}}{\eta_{it}-1}$ measures the overall wedge due to monopoly power, taxes, and/or labor-market distortions, $\mathbb{E}_{it}$ denotes the expectation conditional on the information that is available to island i, and Y_t denotes aggregate output (with $Y_t = C_t$, since there is no capital).

In the absence of informational frictions, condition (1) holds without the expectation operator; in its presence, equilibrium outcomes diverge from their complete-information counterparts insofar as aggregate output, Y_t, is not commonly known. Building on this observation, the following lemma helps reveal a formal connection between the *positive* properties of our model and those of the class of beauty-contest games studied by Morris and Shin (2002); Angeletos and Pavan (2007); and Bergemann and Morris (2013).

LEMMA 1: *The equilibrium level of output is pinned down by the following fixed-point relation:*

$$\log y_{it} = \phi_0 + \phi_a a_{it} + \phi_\mu \mu_{it} + \alpha \, \mathbb{E}_{it}[\log Y_t], \tag{2}$$

where ϕ_0, $\phi_a > 0$, *and* $\phi_\mu < 0$ *are scalars,* $a_{it} \equiv \log A_{it} - \frac{1}{1+\epsilon}\log \chi_{it}$ *and* $\mu_{it} \equiv \log \mathcal{M}_{it}$ *capture the local shocks, and*

$$\alpha \equiv \frac{1-\rho\gamma}{1+\rho\epsilon} < 1. \tag{3}$$

Condition (2), which is simply a log-linear transformation of condition (1), is formally identical to the best-response condition that characterizes the aforementioned class of beauty-contest games. In the context of these games, the scalar α identifies the degree of strategic complementarity and encapsulates the *private* value of coordination: it measures how much the players in the game (the firms in our model) care to align their actions (their production levels).

In an abstract game, this scalar can be a free variable. In our setting, it is pinned down by the underlying micro-foundations and it reflects the balance of two forces. On the one hand, an increase in aggregate income raises the demand faced by each firm, which stimulates firm profits, production, and employment; this effect captures the "aggregate demand externality." On the other hand, an increase in aggregate income discourages labor supply and raises real wages, which has the opposite effect on firm profits, production, and employment. In our view, the most plausible scenario is one in which the former effect dominates, so that $\alpha > 0$. To simplify the exposition, the comparative statics of volatility and dispersion in Proposition 2 focus

[5] The equilibrium characterization of the baseline model, and a variant of Lemma 1, can also be found in Angeletos and La'O (2010). Our contribution starts with the welfare decomposition in Lemma 2.

on this case. However, our key welfare results (Theorems 1, 2, 3, and 4) hold true regardless of the sign and the precise value of α. They only require $\alpha < 1$, a qualification that is henceforth omitted because it is necessarily satisfied for all admissible values of the underlying preference and technology parameters.

Lemma 1 permits one to characterize the positive properties of our baseline model as a direct translation of the positive properties of the aforementioned class of beauty-contest games. For example, one can readily show that a higher α maps to higher sensitivity of equilibrium production to noisy public news and therefore also to higher non-fundamental volatility; this mirrors a similar result in Morris and Shin (2002). Alternatively, following Bergemann and Morris (2013), one can show that the entire set of equilibrium allocations that obtain under arbitrary Gaussian information structures can be spanned with the two-dimensional signal structure we specify in the next section.

None of these facts, however, informs us about the *normative* properties of our model. To understand these properties, we develop a certain decomposition of the welfare losses that obtain in equilibrium relative to the first best. Let y_{it}^* and Y_t^* denote the first-best levels of, respectively, local and aggregate output; these are pinned down by the variant of condition (1) that sets $\mathcal{M}_{it} = 1$ and that lets Y_t be commonly known. Next, define the corresponding local and aggregate gaps by, respectively, $\log \tilde{y}_{it} \equiv \log y_{it} - \log y_{it}^*$ and $\log \tilde{Y}_t \equiv \log Y_t - \log Y_t^*$, and let

$$\Sigma \equiv \mathrm{Var}\big(\log \tilde{Y}_t\big) \quad \text{and} \quad \sigma \equiv \mathrm{Var}\big(\log \tilde{y}_{it} - \log \tilde{Y}_t\big),$$

measure, respectively, the volatility of the aggregate gap and the cross-sectional dispersion in local gaps.[6] Next, let $\bar{Y} \equiv \mathbb{E}[Y]$ denote the mean, or "steady-state," level of output. Finally, consider, as a reference point, the allocation that obtains when the mean wedge $\bar{\mu}$ is chosen so as to maximize welfare and let $\hat{Y}$ denote the associated mean level of output; as in the New-Keynesian policy literature, this reference point abstracts from distortions in the steady state and can be attained with the introduction of an appropriate noncontingent subsidy on employment or income. We can then reach the following characterization of equilibrium welfare.

LEMMA 2: *There exist functions* $v, w : \mathbb{R}_+ \to \mathbb{R}$, *which are invariant to the information structure*, *such that equilibrium welfare is given by*

$$\mathcal{W} = v(\Delta) w(\Lambda),$$

where

$$\Delta \equiv \frac{\bar{Y}}{\hat{Y}} \quad \text{and} \quad \Lambda \equiv \Sigma + \frac{1}{1-\alpha}\sigma. \tag{4}$$

Furthermore, $\mathcal{W}$ *attains its maximum* (*the first-best level*) *at* $\Delta = 1$ *and* $\Lambda = 0$, *is strictly concave in* Δ *and strictly decreasing in* Λ.

[6] In the literature, it is customary to recast σ as a measure of dispersion in relative prices. Such a transformation is valid in our setting but is not needed for our purposes.

To interpret this lemma, note that $v(\Delta)$ captures the welfare loss caused by any distortion in the mean level of economic activity, whereas $w(\Lambda)$ captures the loss due to volatility in the aggregate output gap and/or due to cross-sectional misallocation. The first loss disappears when $\Delta = 1$ $\left(\text{equivalently, } \bar{Y} = \hat{Y}\right)$, the second when $\Lambda = 0$ (equivalently, $\Sigma = \sigma = 0$).[7]

This lemma and a set of companion results we provide in Section IV extend the kind of welfare decompositions that are familiar in the New-Keynesian framework (Woodford 2003; Galí 2008) to the incomplete-information economies we are interested in. While these decompositions need not be surprising on their own right, and variants of them have appeared in all the related prior work, they serve two purposes. First, they help identify the different channels through which information can affect welfare. Second, they complete the mapping between the macroeconomic models of interest and the abstract games studied in Morris and Shin (2002); Angeletos and Pavan (2007); and Bergemann and Morris (2013), thus also clarifying whether there is any discrepancy between the private and the social value of coordination in the macroeconomic models of interest.

The first point will become evident as we proceed, especially once we add nominal rigidity. To understand the second point, consider any of the games studied in the above papers and momentarily recast Σ and σ as, respectively, the volatility and the dispersion of the gaps between the equilibrium and the first-best actions in that game. Following Angeletos and Pavan (2007), the combined welfare loss due to these gaps can be shown to be proportional to the following sum:

$$\Lambda = \Sigma + \frac{1}{1-\alpha^*}\sigma,$$

where α^* is a scalar that depends on the payoff structure of the game and that encapsulates the *social* value of coordination.[8] In general, this scalar may differ from the one that measures the degree of strategic complementarity, reflecting a divergence between private and social motives to coordinate. In the light of Lemmas 1 and 2, however, our economy maps to a game in which $\alpha = \alpha^*$, meaning that there is no such divergence.[9]

PROPERTY 1: *In our setting, the private value of coordination coincides with its social counterpart.*

This property underscores a crucial difference between our setting and that of Morris and Shin (2002). In their game, it is assumed that $\alpha^* = 0 < \alpha$; that is, the

[7] In addition, we normalize $v(1) = 1$. It follows $w(0)$ coincides with the first-best level of welfare.

[8] Formally, α^* is defined as the degree of strategy complementarity in a fictitious game whose equilibrium strategy coincides with the strategy that maximizes welfare in the economy under consideration; it therefore reflects how much agents *should* care to coordinate, as opposed to how much they actually do care in equilibrium.

[9] The mapping between our model and Angeletos and Pavan (2007) is complicated by the fact that the Δ term can vary with the available information in our setting, a kind of effect that is not accommodated by the linear-quadratic framework of Angeletos and Pavan (2007). This complication turns out to be inconsequential in the case of technology shocks, but not in the case of markup shocks. See Proposition 4 and the discussion surrounding this proposition. Also, for the case of technology shocks, the coincidence of α and α^* was first pointed out in Angeletos and La'O (2010) by comparing directly the equilibrium to the constrained efficient allocation. That paper, however, did not arrive at the precise mapping between the welfare effects of information in our setting and those in Angeletos and Pavan (2007), nor did it consider the extension with nominal rigidity we consider in Section IV.

private motives to coordinate are socially wasteful *by design*. The micro-foundations of our setting, instead, imply that $\alpha^* = \alpha > 0$; that is, the private incentives to coordinate are perfectly aligned with their social counterparts, as if by the magic of the invisible hand. As we explain in the next section, this alignment is also key to understanding why the *combined* effect of information on Λ turns out to be unambiguous, even though its *component* effects on volatility and dispersion are ambiguous in general and are often in conflict with one another.

III. The Effects of Information on Volatility, Dispersion, and Welfare

In this section, we characterize the comparative statics of the volatility measure Σ, the dispersion measure σ, and overall welfare $\mathcal{W}$ with respect to the information structure. We do so by distinguishing two polar cases. In the first, the underlying fundamental uncertainty is over technology or preferences. In the second, it is over monopoly power or labor wedges. The first case captures the scenario in which the business cycle would have been efficient had information been complete; in this case, Σ and σ are nonzero only due to the incompleteness of information. The second case captures the scenario in which the business cycle originates from distortions in product and labor markets; in this case, Σ and σ reflect the combination of the informational friction with such distortions.

A. *Efficient Fluctuations*

In this part, we fix $\mathcal{M}_{it} = \overline{\mathcal{M}}$ for all (i,t) and concentrate on the case of technology shocks. (The case of preference shocks is identical in terms of welfare.)

To facilitate sharp comparative statics, we specify the stochastic structure as follows. First, we let local productivity be $a_{it} \equiv \log A_{it} = \bar{a}_t + \xi_{it}$, where $\bar{a}_t$ is an aggregate shock and ξ_{it} is an idiosyncratic shock. The aggregate shock $\bar{a}_t$ is i.i.d. over time, drawn from $\mathcal{N}(0,\sigma_a^2)$, while the idiosyncratic shock ξ_{it} is i.i.d. across both t and i, independent of $\bar{a}_t$, and drawn from $\mathcal{N}(0,\sigma_\xi^2)$. Next, we summarize all of the private (local) information of island i regarding the underlying aggregate shock $\bar{a}_t$ in an island-specific signal x_{it} given by

$$x_{it} = \bar{a}_t + u_{it}, \tag{5}$$

where the noise term u_{it} is i.i.d. across i and t, orthogonal to $\bar{a}_t$, and drawn from $\mathcal{N}(0,\sigma_x^2)$.[10] Similarly, we summarize all of the public (aggregate) information in a public signal z_t given by

$$z_t = \bar{a}_t + \varepsilon_t, \tag{6}$$

[10] Note that local productivity is itself a private signal of aggregate productivity. The sufficient statistic x_{it} is meant to include this information. More precisely, $x_{it} \equiv (1-\omega)a_{it} + \omega x'_{it}$, where $x'_{it} = \bar{a} + u'_{it}$ is a signal that captures any private information *other* than the one contained in local productivity; u'_{it} is the noise in that signal, which is i.i.d. across i and t, orthogonal to $\bar{a}_t$ and ξ_{it}, and drawn from $\mathcal{N}(0,\sigma_x'^2)$; $\sigma_x^{-2} \equiv \sigma_\xi^{-2} + \sigma_x'^{-2}$; and $\omega \equiv \sigma_x'^{-2}/\sigma_x^{-2}$.

where the noise term ε_t is i.i.d. across t, orthogonal to all other shocks, and drawn from $\mathcal{N}(0, \sigma_z^2)$. Finally, to ease notation, we let $\kappa_a \equiv \sigma_a^{-2}$, $\kappa_\xi \equiv \sigma_\xi^{-2}$, $\kappa_x \equiv \sigma_x^{-2}$, and $\kappa_z \equiv \sigma_z^{-2}$.

The subsequent analysis focuses on the comparative statics of the equilibrium volatility, dispersion, and welfare with respect to the scalars κ_x and κ_z, which measure the precisions of, respectively, the available private and public information. When interpreting our results, however, it is worth keeping in mind the following point. As noted before, the results of Bergemann and Morris (2013) guarantee that the equilibrium allocation obtained by *any* Gaussian information structure can always be replicated with an information structure like the one specified above. This means that the adopted specification is without serious loss of generality and that the scalars κ_x and κ_z represent more generally a convenient parameterization of the information structure.

Prior work has often emphasized the different effects that each type of information can have on volatility and dispersion. We thus start by revisiting these effects in the context of our model.

PROPOSITION 1: (*i*) *An increase in* κ_z *necessarily reduces dispersion* σ, *whereas it reduces volatility* Σ *if and only if* κ_z *is high enough.* (*ii*) *Symmetrically*, *an increase in* κ_x *necessarily reduces* Σ, *whereas it reduces* σ *if and only if* κ_x *is high enough.*

To understand part (i), note that an increase in the precision of public information induces firms and workers to reduce their reliance on their private signals, which in turn reduces the contribution of idiosyncratic noise to cross-sectional dispersion. At the same time, because these agents increase their reliance on public signals, the contribution of public information to aggregate output gaps is ambiguous: the reduction in the level of the noise itself tends to reduce Σ, while the increased reaction of the agents tends to raise Σ. Which effect dominates depends on how large the noise is, which explains part (i). The intuition for part (ii) is symmetric.

Although each type of information can have a negative effect on either volatility or dispersion, the *combined* welfare effect is unambiguously positive: Λ necessarily decreases with either κ_x or κ_z. Along with the fact that Δ turns out to be invariant to the information structure,[11] this gives us the following result.

THEOREM 1: *Suppose the business cycle is driven by technology shocks. Welfare necessarily increases with the precision of either public or private information*, *for all* α. *Moreover*, *when* $\alpha > 0$, *the marginal welfare benefit of public information increases with* α, *that is*, $\frac{\partial^2 \mathcal{W}}{\partial \alpha \partial \kappa_z} > 0$.

As anticipated, this result owes its sharpness to the coincidence of the private and social values of coordination and can thus can be seen as a variant of Proposition 6 in Angeletos and Pavan (2007). If the scalar that governs the relative contribution of volatility and dispersion in Λ were lower from the one that governs the strategic complementarity $(\alpha^* < \alpha)$, then *public* information would have a non-monotone

[11] The intuition for this particular property is discussed in the context of Proposition 4 below.

welfare effect, in line with the result of Morris and Shin (2002); and if the converse were true ($\alpha^* > \alpha$), then it would be *private* information that would have a non-monotone welfare effect. It is thus Property 1 that explains why both types of information have a similar and unambiguously welfare effect in our setting.[12]

B. *Inefficient Fluctuations*

We now shift focus to the case of inefficient fluctuations, which we capture with shocks to monopoly markups (or, equivalently, to labor wedges). We thus fix $A_{it} = \chi_{it} = 1$ for all (i, t) and let the log of the local wedge be given by

$$\mu_{it} \equiv \log \mathcal{M}_{it} = \bar{\mu}_t + \xi_{it},$$

where $\bar{\mu}_t$ is an aggregate component and ξ_{it} is an idiosyncratic component. The former is i.i.d. across t, drawn from $\mathcal{N}\left(\bar{\mu}, \sigma_{\mu}^2\right)$; the latter is i.i.d. across both t and i, independent of $\bar{\mu}_t$, and drawn from $\mathcal{N}\left(0, \sigma_{\xi}^2\right)$. Finally, we let $\kappa_{\mu} \equiv \sigma_{\mu}^{-2}$ and model the information structure in the same way as in the previous section: the available signals are given by (5) and (6), replacing $\bar{a}_t$ with $\bar{\mu}_t$.

In the case of the technology shocks, equilibrium allocations could fluctuate away from the first best *only* because of the incompleteness of information. Here, by contrast, the *entire* variation in equilibrium allocations represents a deviation from the first best, no matter whether this variation originates in the noise or in the fundamentals themselves. The comparative statics of the resulting volatility and dispersion measures are described below.

PROPOSITION 2: *Suppose $\alpha > 0$. (i) Volatility Σ increases with either κ_x or κ_z. (ii) Dispersion σ decreases with κ_z, and is generally non-monotone in κ_x.*

In spite of the possible conflict between the component effects, Property 1 guarantees that the combined effect is once again unambiguous—but now of the opposite sign than in the case of technology shocks.

PROPOSITION 3: *The combined welfare loss due to volatility and dispersion, as captured by Λ, increases with either κ_x or κ_z, for all α.*

As before, it is useful to relate the above finding to Angeletos and Pavan (2007). Corollary 9 of that paper uses an abstract example in which $\alpha^* = \alpha = 0$ to illustrate the basic insight that information can be detrimental for welfare when it regards shocks that only move the complete-information equilibrium away from the first best. However, by leaving open the possibility that $\alpha^* \neq \alpha$ in workhorse macroeconomic models, and in fact conjecturing that $\alpha < \alpha^*$, that paper also left open the door for ambiguous welfare effects. Similarly to Theorem 1, the above result

[12] Theorem 1 can also be inferred from the result in Angeletos and La'O (2010) that the equilibrium is constrained efficient in the absence of markup shocks. This, however, does not extend to the rest of our results (Theorems 2, 3, and 4).

therefore owes its sharpness to Property 1, the coincidence of the private and social values of coordination.

Welfare depends not only on Λ, which we characterized above, but also on Δ. In the case of technology shocks, Δ was pinned down by the mean wedge $\bar{\mathcal{M}}$, and was invariant to the information structure. Here, instead, Δ varies with the level of noise.

PROPOSITION 4: *Δ increases with either κ_x or κ_z, for all α.*

This finding can be explained as follows. The uncertainty that firms face in predicting aggregate demand affects the mean level of economic activity, due to curvature at both the firm level (curvature of the profit function) and the aggregate level (imperfect substitutability across products). This effect is present irrespective of the nature of the underlying aggregate shocks. Its welfare consequences, however, hinge on the nature of the shocks. In the case of technology shocks, the equilibrium use of information is socially optimal and this effect does not represent a distortion, which explains why Δ does not vary with the information structure. In the case of markup shocks, instead, the planner would prefer the agents not to respond to the underlying uncertainty and the aforementioned effect is thus associated with an increase in Δ.

Recall that welfare is strictly convex in Δ, with a maximum attained at $\Delta = 1$. It follows that the aforementioned effect represents a welfare loss when $\Delta > 1$ and a welfare gain when $\Delta < 1$. In the former case, this effect therefore complements the one of Λ. In the latter case, instead, the two effects conflict with each other. Which one dominates then depends on the distance of Δ from the bliss point $\Delta = 1$. Finally, this point is itself attained if the planner has at his disposal a fiscal instrument that permits him to control the mean level of output, such as a noncontingent subsidy on employment, output, or sales. We thus reach the following result.

THEOREM 2: *Suppose the business cycle is driven by markup shocks. There exists a threshold $\hat{\Delta} \in (0,1)$ such that welfare decreases with the precision of either public or private information if and only if $\Delta > \hat{\Delta}$. Furthermore, the latter condition holds, with $\Delta = 1$, if a noncontingent subsidy is available and set optimally.*

It is interesting to note that the threshold $\hat{\Delta}$ is pinned down solely by preference and technology parameters and it is the same whether we consider the effect of private information or that of public information. This kind of symmetry between the two types of information is yet another symptom of Property 1: no matter which case we have considered, this property has guaranteed that the distinction between private and public information is inconsequential for the question of interest.

We conclude this section by noting that it is customary in the literature to shut down any "steady-state" distortion (that is, to set $\Delta = 1$) by assuming from scratch the presence of the noncontingent subsidy. Although the effect on Δ documented above may be of interest in its own right, in the sequel we also opt to abstract from it and, instead, extend the analysis in the direction of adding nominal rigidity and studying the role of monetary policy.

IV. Nominal Rigidity and Monetary Policy

In the preceding analysis we isolated the role of the informational friction as a source of real rigidity. We now extend the analysis to the more realistic scenario in which the informational friction is also a source of nominal rigidity: firms set their nominal prices on the basis of the kind of noisy private and public signals that were featured in our preceding analysis.

A. *Setup*

As usual, the introduction of nominal rigidity requires that we allow a margin of adjustment in quantities: at least one input must be free to adjust to realized demand, or else markets would fail to clear at the posted prices. Accordingly, we allow for two types of labor: one that is chosen on the basis of incomplete information, thus preserving the type of real rigidity that was at the core of our baseline model; and another that adjusts freely to the underlying state of nature, thus preserving market clearing in the presence of the nominal rigidity.

More specifically, the output of the typical firm in island i is now given by

$$y_{it} = A_{it} n_{it}^{\theta} \ell_{it}^{\eta},$$

where n_{it} is the labor input that is chosen on the basis of incomplete information (as in our baseline analysis), ℓ_{it} is the alternative input that adjusts to the realized state (so that markets can clear), and θ and η are positive scalars, with $\theta + \eta \leq 1$. One may think of n_{it} as bodies of employed workers whom the firm hires on the basis of incomplete information and of ℓ_{it} as labor utilization, overtime work, or other margins that adjust to realized demand. The precise interpretation of these inputs, however, is not essential. Rather, the essence is to accommodate the two types of rigidity we are interested in. A useful feature of the presumed specification is that it helps nest the scenario studied in prior work as the extreme case in which $\theta = 0$ (meaning that all output is free to adjust to the realized state).

We next let the per-period utility of the representative household be given by the following sum:

$$\frac{1}{1-\gamma} C_t^{1-\gamma} - \frac{1}{1+\epsilon_n} \int_I n_{it}^{1+\epsilon_n} di - \frac{1}{1+\epsilon_\ell} \int_I \ell_{it}^{1+\epsilon_\ell} di,$$

where ϵ_n and ϵ_ℓ are positive scalars that parameterize the Frisch elasticities of the two types of labor. To simplify, and without serious loss, we let $\epsilon_n = \epsilon_\ell = \epsilon$.

Note that curvature in both the utility function and in the production function in the two types of inputs ensures that the alternative input ℓ_{it} cannot be a perfect substitute for n_{it}. That is, while ℓ_{it} adjusts to the realized state, it cannot fully undo the effect of the predetermined labor input n_{it}.

Consider now the specification of monetary policy. In general, this opens the door to delicate modeling issues. What is the information upon which the monetary authority acts? Does this contain only signals of the exogenous shocks or also signals of endogenous economic outcomes? What are the objectives, targets, or policy rules that guide the policymaker? How one chooses to answer these questions is bound to affect the welfare properties of the model. In what follows, we develop a taxonomy that seeks to dissect the role of different monetary policies, without however getting into the granular details of how policy is conducted.

We assume that the policy instrument is the nominal interest rate and allow the latter to follow an arbitrary stochastic process. We only require that this process is log-normal in order to maintain the Gaussian structure of the equilibrium. Following the tradition of the Ramsey literature, we then adopt an approach that permits us to span directly the set of all the allocations that can obtain in equilibrium under such an arbitrary monetary policy. The benefit of this approach is its flexibility; the cost is that it suppresses the question of what exactly it takes for the policymaker to be able to implement a particular allocation.

To economize on space, the characterization of the set of allocations that can be implemented with *arbitrary* monetary policies is delegated to online Appendix B (see Section B1 and Lemmas 6–8 therein). To facilitate the subsequent analysis, we nevertheless need a "topography" of this set, that is, a way to index the different points in it. We provide such a topography in the next lemma.

LEMMA 3: (*i*) *In any equilibrium, nominal GDP satisfies*

$$\log M_t = \lambda_s \bar{s}_t + \lambda_z z_t + m_t, \tag{7}$$

where λ_s *and* λ_z *are scalars,* $\bar{s}_t$ *stands for either the technology or markup shock, and* m_t *is a random variable that is drawn from* $\mathcal{N}(0, \sigma_m^2)$, *for some* $\sigma_m \geq 0$, *and is orthogonal to both* $\bar{s}_t$ *and* z_t.

(*ii*) *Suppose that the interest rate satisfies*

$$\log(1 + R_t) = \rho_s \bar{s}_t + \rho_z z_t + r_t, \tag{8}$$

where ρ_s *and* ρ_z *are scalars and* r_t *is a random variable that is drawn from* $\mathcal{N}(0, \sigma_r^2)$ *for some* $\sigma_r \geq 0$, *and is orthogonal to both* $\bar{s}_t$ *and* z_t. *For any triplet* $(\lambda_s, \lambda_z, \sigma_m)$, *there exists a monetary policy as in* (8) *such that* (7) *holds in the equilibrium induced by this policy.*

(*iii*) *A policy as in* (8) *can replicate the equilibrium allocation induced by any other policy.*

Part (i) follows from regressing the equilibrium value of nominal GDP on the fundamental and the public signal, and letting (λ_s, λ_z) be the projection coefficients and m_t the residual. This part is therefore trivial, but it is useful for our purposes because in conjunction with the rest of the lemma it permits us to index different equilibria with different values for the triplet $(\lambda_s, \lambda_z, \sigma_m)$. Parts (ii) and (iii) then provide us with a class of monetary policies that can implement *any* value for this triplet and that span the entire set of the allocations that obtain under arbitrary monetary policies.

Although it is possible to interpret condition (8) as a policy rule, one can also arrive to it from a different specification of how policy is conducted. For instance, suppose that the monetary authority adheres to the following Taylor rule:

$$\log(1+R_t) \;=\; r_z z_t + r_y\left(\log Y_t + \epsilon_t^y\right) + r_p\left(\log P_t + \epsilon_t^p\right) + \tilde{r}_t,$$

where $\left(r_z, r_y, r_p\right)$ are policy coefficients, ϵ_t^y and ϵ_t^p are measurement errors in the monetary authority's observation of real output and the price level, and $\tilde{r}_t$ is a monetary shock. Once one solves for equilibrium output and prices, the above reduces to condition (8), with $\left(\rho_{\bar{s}}, \rho_z\right)$ being functions of the policy coefficients $\left(r_z, r_y, r_p\right)$ and r_t being a mixture of the monetary shock $\tilde{r}_t$ and the measurement errors $\left(\epsilon_t^y, \epsilon_t^p\right)$. In a nutshell, condition (8) can always be recast as a representation of the equilibrium implemented by any given policy rule.

Furthermore, although condition (8) requires that the interest rate react to the current technology or markup shock, such a contemporaneous reaction is not strictly needed for the policymaker to implement a particular response in macroeconomic activity to the shock. Rather, it suffices that monetary policy reacts at some point in the future: Lemma 7 in online Appendix B establishes that the entire set of implementable allocations remains the same whether monetary policy responds within the same period or with a lag. This follows directly from iterating the Euler equation and noting that current consumption depends on the entire path of future nominal interest rates. As a result, it makes no difference whether the desired movements in aggregate demand are implemented by moving the current interest rate or by committing to move future rates.[13]

These points underscore that conditions (7) and (8) are equivalent representations of all the equilibrium allocations that can obtain under arbitrary monetary policies. We have found (7) to be most convenient for our purposes, for reasons that will become evident in the statement of the formal results in this section, as well as in the discussion of the related literature in the next section.[14]

To close the model, we must specify the information upon which firms can condition their production and pricing decisions. As in the baseline model, we assume that this is summarized by a pair of signals about the underlying fundamental: the public signal z_t and the private signal x_t. We proceed to investigate the comparative statics of welfare with respect to the corresponding precisions, κ_x and κ_z. Note that this rules out the possibility that the firms also have information about the shock m_t, which can be interpreted as a monetary shock. This alternative kind of information is the subject matter of Hellwig (2005) and is briefly discussed at a later point.

B. *A Familiar Benchmark*

Consider, as a reference point, the hypothetical scenario in which the nominal rigidity is removed, by which we mean the case in which p_{it} is free to adjust to the realized state. This scenario is henceforth referred to as "flexible prices" and the equilibrium allocation that obtains under it as the "flexible-price allocation." The

[13] This property is not specific to our model; it is a standard feature of the New-Keynesian framework and underlies the argument about "forward guidance" at the zero lower bound. Investigating the realism or the robustness of this property is however beyond the scope of this paper.

[14] Some papers, such as Woodford (2002) and Hellwig (2005), treat M_t as an exogenous random process. Others, such as Baeriswyl and Conrand (2010), assume that the policy instrument is M_t rather than the interest rate. Our approach can accommodate these possibilities, but is not limited to them.

next lemma identifies a set of monetary policies that implement this allocation when the nominal rigidity is present.

LEMMA 4: *There exists a pair* (λ_s^*, ρ^*) *such that a monetary policy replicates flexible prices if* (7) *holds with* $\lambda_s = \lambda_s^*$ *and* $\sigma_m = 0$ *or, equivalently, if* (8) *holds with* $\rho_s = \rho^*$ *and* $\sigma_r = 0$.

This lemma is a special case of a more general result in Angeletos and La'O (2014): just as in the baseline New-Keynesian framework there are monetary policies that can undo the nominal rigidity induced by Calvo-like sticky prices, in the class of incomplete-information models studied here (and in related papers) there are monetary policies that can undo the nominal rigidity induced by informational frictions. These policies presume that the policymaker can observe the aggregate state perfectly, although perhaps with a time lag, and that she has perfect control over aggregate demand. These policies are therefore not particularly realistic. They nevertheless represent a useful benchmark, separating the informational friction of the market from any friction on the policymaker's side, and facilitating sharp welfare conclusions.[15]

When these policies are in place, information matters only through the real rigidity. This suggests a connection to our baseline analysis, which we formalize next. Let $q_{it} \equiv A_{it} n_{it}^{\theta}$ denote the component of output that is determined on the basis of incomplete information. Next, define the corresponding aggregate as

$$Q_t \equiv \left[\int_I (q_{it})^{\frac{\hat{\rho}-1}{\hat{\rho}}} di\right]^{\frac{\hat{\rho}}{\hat{\rho}-1}},$$

and finally let

$$\hat{\alpha} \equiv \frac{1 - \hat{\rho}\,\hat{\gamma}}{1 + \hat{\rho}\,\hat{\epsilon}}, \tag{9}$$

where $\hat{\epsilon} \equiv \frac{1+\epsilon-\theta}{\theta}$, $\hat{\gamma} \equiv 1 - \frac{(1-\gamma)(1+\epsilon)}{1+\epsilon-\eta(1-\gamma)}$, and $\hat{\rho} \equiv \frac{\rho(1+\epsilon-\eta)+\eta}{1+\epsilon+\eta(1-\rho)}$ are transformations of the underlying preference and technology parameters. Similarly to the baseline model, we can verify that $\hat{\alpha} < 1$ for all preference and technology parameters, and can obtain the following characterization of the flexible-price allocation.

PROPOSITION 5: *Consider the flexible-price allocation and let* q_{it} *and* Q_t *be the quantities defined above.*

(*i*) *There exist scalars* $\hat{\phi}_a > 0$, $\hat{\phi}_\mu < 0$, *and* $\hat{\phi}_{\bar{\mu}}$ *such that, for any information structure, the equilibrium* q_{it} *is determined by the solution to the following:*

$$\log q_{it} = \hat{\phi}_a a_{it} + \hat{\phi}_\mu \mu_{it} + \hat{\phi}_{\bar{\mu}} \mathbb{E}_{it}[\bar{\mu}_t] + \hat{\alpha}\,\mathbb{E}_{it}[\log Q_t]. \tag{10}$$

[15] Our analysis also abstracts from any interference the nominal rigidity may have with the response to sectoral or idiosyncratic shocks, or from other types of relative-price distortions that monetary policy may be unable to correct even when the policymaker observes perfectly the state.

(*ii*) *There exists a decreasing function w such that, for any information structure, equilibrium welfare is given by* $\mathcal{W} = w(\Lambda)$, *where*

$$\Lambda = \Sigma + \frac{1}{1-\hat{\alpha}}\sigma + \omega, \tag{11}$$

where Σ *and* σ *are defined in the same way as in the baseline model, modulo replacing output y with the component q defined above and the scalar* α *with the scalar* $\hat{\alpha}$, *and where* ω *is a scalar that does not depend on either* κ_x *or* κ_z *and that vanishes in the absence of markup shocks.*

This result extends Lemmas 1 and 2 from the baseline model to the flexible-price allocation of the extended model, with conditions (10) and (11) being the analogues of, respectively, conditions (2) and (4) of the baseline model.

If we compare condition (10) to condition (2), we see three differences. First, q_{it} and Q_t have taken the place of, respectively, y_{it} and Y_t. This is because it is only q_{it} (or n_{it}), not y_{it}, that is restricted to depend on incomplete information. Second, the hatted scalars $\hat{\alpha}$, $\hat{\rho}$, etc. have taken the place of the corresponding un-hatted scalars in the baseline model. This reflects the more general specification of preferences and technologies. Finally, a new term has emerged: in addition to the firm's own markup, the expected aggregate markup enters the firm's best-response condition. This is because the realized aggregate markup affects the realized aggregate output for any given Q_t, implying in turn that a firm's optimal choice of q_{it} depends directly on its expectation of $\bar{\mu}_t$.

Comparing now condition (11) to condition (4), we see that a new term shows up also in the definition of Λ. Even if we hold constant the entire cross-sectional profile of q_{it}, the realized markup distorts the firms' choice of ℓ_{it}. This explains why Λ contains not only the terms Σ and σ but also the new term ω in condition (11), which is proportional to the volatility of the markup shock.

Notwithstanding these differences in the micro-foundations of conditions (10) and (11), the scalar $\hat{\alpha}$ that captures the strategic complementarity in (10) continues to determine the relative welfare costs of volatility and dispersion in condition (11). The following variant of Property 1 therefore holds in the extended model.

PROPERTY 2: *The private and the social value of coordination coincide as long as monetary policy replicates flexible prices.*

This observation leads once again to a sharp answer to the question of interest. Insofar as monetary policy replicates flexible prices, ℓ_{it} choice adjusts to the realized supply and demand conditions *as if* information were complete. The ω term in (11) is therefore invariant to κ_x and κ_z, and can be ignored for our purposes. Consider now the case of technology shocks. In this case, q_{it} solves *exactly* the same fixed-point relation as y_{it} did in the baseline model, guaranteeing that the mapping from the information structure to the equilibrium values of Σ, σ, and Λ are also the same, modulo the replacement of the scalars α and ϕ_a with their hatted counterparts. When instead the business cycle is driven by markup shocks, the presence of the aggregate markup in condition (10) breaks the equivalence between the two models. Nevertheless, because this term is just a different facet of the distortionary effects of markup shocks, it does not interfere with the essence of the lessons of the baseline model. We thus arrive to the following extension of Theorems 1 and 2.

THEOREM 3: *Suppose monetary policy replicates flexible prices. Welfare increases with both types of information in the case of technology shocks, and decreases with them in the case of markup shocks.*

C. *Away from the Benchmark*

We now turn to policies that deviate from replicating flexible prices. From Lemma 3, any such deviation contains at most three components: one perfectly correlated with the fundamental; one perfectly correlated with the public signal; and a residual, which may be interpreted as a monetary shock. The second component has no welfare consequences, because the dependence of policy on z_t is commonly predictable when firms set prices and can thus have no real effect. We thus reach the following decomposition of welfare for *any* monetary policy.

PROPOSITION 6: *There exists a decreasing function w such that, for any policy and any information structure, the equilibrium level of welfare is given by*

$$\mathcal{W} = w(\Lambda + \mathcal{K} + \mathcal{T}),$$

where the following are true:

(i) *Λ is the welfare loss at the flexible-price allocation.*

(ii) *$\mathcal{K}$ is the welfare effect of $\lambda_s \neq \lambda_s^*$ and can be expressed as $\mathcal{K} = \mathcal{K}(\lambda_s)$, where*

$$(12)\quad \mathcal{K}(\lambda_s) \equiv \begin{cases} \Theta\,(\lambda_s - \lambda_s^*)^2\,\sigma_a^2 & \textit{with technology shocks} \\ \left[-2\,\Theta_1\,(\lambda_s - \lambda_s^*) + \Theta_2(\lambda_s - \lambda_s^*)^2\right]\sigma_\mu^2 & \textit{with markup shocks} \end{cases}$$

and where Θ, Θ_1, and Θ_2 are scalars that depend on (κ_x, κ_z) and that vanish as $\kappa_x \to \infty$ or $\kappa_z \to \infty$.

(iii) *$\mathcal{T}$ is the welfare loss caused by $\sigma_m \neq 0$.*

This result complements our earlier welfare decompositions and formalizes the sense in which the two forms of rigidity map into two channels through which information affects welfare: the role of the real rigidity is captured by Λ; the additional effect of the nominal rigidity is captured by the sum $\mathcal{K} + \mathcal{T}$. This sum is nonzero only insofar as monetary policy deviates from the benchmark of replicating flexible prices. By contrast, Λ is necessarily positive.[16]

As already mentioned, $\mathcal{T}$ captures the welfare consequences of m_t, or equivalently of r_t. The latter represents a deviation that is orthogonal to both the underlying fundamental and the public signal—a deviation that can be interpreted as a

[16] With the exemption, of course, of the limit case in which $\theta = 0$: without real rigidity, $\Lambda = 0$.

monetary shock or a policy "mistake." Whatever the interpretation, $\mathcal{T}$ is necessarily nonnegative and independent of the available information about $\bar{s}_t$. By contrast, $\mathcal{K}$ depends on that information precisely because it captures the deviations that are correlated with $\bar{s}_t$. Furthermore, both the sign and the comparative statics of this term with respect to the available information depend on the nature of the underlying business-cycle forces.

From part (ii), we see that the minimum of $\mathcal{K}$ is zero and it is attained at $\lambda_s = \lambda_s^*$ when the business cycle is driven by technology shocks, whereas it is positive and it is attained at $\lambda_s \neq \lambda_s^*$ when the business cycle is driven by markup shocks. This verifies that a familiar policy lesson extends from the New-Keynesian setting to the present framework.

PROPERTY 3: *A monetary policy that replicates flexible prices is optimal in the case of technology shocks, but not in the case of markup shocks.*

Notwithstanding the similarity to the New-Keynesian framework, the following difference is worth mentioning in the case of technology shocks: unless the real rigidity is shut down ($\theta = 0$), replicating flexible prices does *not* implement the first-best allocation, nor is it synonymous to targeting price stability.[17]

Turning to the case of markup shocks, certain deviations from the flexible-price allocation are welfare-improving because they substitute for a missing tax instrument, namely the state-contingent subsidy that would have offset the markup shock. The only key difference from the New-Keynesian framework then is that, since the nominal rigidity originates in an informational friction rather than Calvo-like sticky prices, the ability of the monetary authority to counter the markup shock hinges on its ability to respond to information that is not available to the firms when the latter set their prices. This suggests that more precise information in the firms' hands may contribute toward lower welfare not only by reducing the "base" level of welfare that obtains in the flexible-price allocation but also by limiting the ability of the monetary authority to counteract the markup shock.[18] As it turns out, however, this intuition is only partially correct.

LEMMA 5: *Suppose the business cycle is driven by markup shocks. The optimal policy corresponds to $\lambda_s = \lambda_s^{**}$ and $\sigma_m = 0$, where $\lambda_s^{**} \equiv \arg\min_{\lambda_s} \mathcal{K}(\lambda_s)$. Let $K(\kappa_x, \kappa_z)$ denote this minimum, as a function of the precisions of information.*

(*i*) *Suppose $\theta = 0$. Then, $K(\kappa_x, \kappa_z)$ is increasing in both κ_x and κ_z.*

(*ii*) *Suppose $\theta > 0$. There are values of the preferences and technology parameters for which $K(\kappa_x, \kappa_z)$ is non-monotone in either κ_x* or *κ_z.*

(*iii*) *Suppose $\theta > 0$ and let $\bar{K}(\kappa, \varrho)$ be the function defined by $K(\kappa_x, \kappa_z)$ along the locus of (κ_x, κ_z) such that $\kappa_x + \kappa_z = \kappa$ and $\frac{\kappa_x}{\kappa_\mu + \kappa_z} = \varrho$. Then, $\bar{K}(\kappa, \varrho)$ increases in κ, but is non-monotone in ϱ.*

[17] For more details on these points, see Angeletos and La'O (2014).

[18] The latter possibility is also highlighted in Baeriswyl and Cornand (2010), albeit in a model that adds a certain friction in the conduct of monetary policy. As explained in Section V, this friction is the key to understanding why that paper does not reach the kind of unambiguous result we obtain in part (i) of Lemma 5 below.

Part (i) verifies the aforementioned intuition in the special case in which the real rigidity is absent; as already mentioned, this is the case considered in prior work. Part (ii) establishes that the intuition can be overturned once the real rigidity is present. Part (iii) concludes by providing a qualified variant of the intuition that holds true irrespective of the real rigidity: any additional information at the hands of the firms necessarily reduces the welfare contribution of the optimal monetary policy if that comes without a change to the degree to which information is common (meaning an increase in the overall precision, κ, without a change in the relative precision, ϱ).

We do not fully comprehend the non-monotonicities documented in part (ii). At this point, what we know for sure, thanks to parts (i) and (iii), is only that these non-monotonicities derive exclusively from the interaction of the real rigidity with the degree to which information is correlated across the firms.[19] We thus reach the following summary.

THEOREM 4: *Suppose that monetary policy is optimal.*

(*i*) *When the business cycle is driven by technology shocks, more information improves welfare by, and only by, improving the efficiency of the underlying flexible-price allocation.*

(*ii*) *When instead the business cycle is driven by markup shocks, more information contributes to lower welfare both by exacerbating the inefficiency of the underlying flexible-price fluctuations and by reducing the ability of the monetary authority to combat these fluctuations. Nevertheless, when and only when the real rigidity is present, an ambiguous effect can obtain with changes in the composition of information.*

This result refers to the solution of a Ramsey problem where the planner is free to select an arbitrary Gaussian process for the interest rate or, equivalently, to induce any triplet $(\lambda_s, \lambda_z, \sigma_m)$ he wishes in condition (7). We now discuss what happens when the monetary policy falls short of this "unconstrained" optimum. This could be because the monetary authority has imperfect control of aggregate demand, because it observes the underlying shocks and/or the economic outcomes with noise, because its objectives diverge from the welfare criterion in the model, or because of any other reason that is left outside our model.

Not surprisingly, not much can be said if one puts no structure whatsoever on the deviation from optimality. To understand the logic, let us concentrate on the case of technology shocks. When the nominal rigidity is shut down, the informational friction represents a real distortion that moves the equilibrium away from the first best. Increasing the precision of the available information necessarily reduces the welfare cost of this distortion. When this is the only distortion, more information is unambiguously welfare-improving. But if an additional distortion is present due to the combination of nominal rigidity and suboptimal monetary policy, a second-best result applies: reducing information may increase welfare by having one distortion offset the other.

[19] In part (iii), the ratio ϱ is defined as the ratio of the precision of the private information to that of the public information. However, following Angeletos and Pavan (2007) and Bergemann and Morris (2013), this ratio can be interpreted more generally as a measure of the extent to which information is correlated across the agents.

The opposite scenario, however, is also possible and seems relevant for the following reason. To the extent that she is guided by standard New-Keynesian lessons, the policymaker may fail to incorporate how the informational friction affects the nature of the optimal allocation and the corresponding policy targets. In so doing, the policymaker may inadvertently introduce distortions in addition to those induced by the informational friction. But when the latter vanishes, the policymaker's "mistake" also vanishes. Under this scenario, more precise information may help increase welfare not only by attenuating the real rigidity but also by alleviating the policy suboptimality. We illustrate this logic in Section B3 of online Appendix B, with a numerical example that examines the welfare implications of policies that target either price-level or output-gap stabilization.

D. *Monetary Shocks*

Consider now $\mathcal{T}$, the welfare term corresponding to deviations from the flexible-price benchmark that are orthogonal to the fundamental and the public signal. As already noted, these deviations can be interpreted as monetary shocks. So far, we have assumed that the firms have no information about them. But now suppose the contrary. How does this distinct type of information matter for welfare?

This question is the subject matter of Hellwig (2005). The answer is as follows. Consider first the case of public information. In equilibrium, any *commonly* predictable variation in m_t can have no real effect. Furthermore, any residual variation in m_t necessarily contributes to welfare losses, because m_t is orthogonal to the underlying preferences and technologies. It follows that more precise public information about m_t necessarily improves welfare (it reduces $\mathcal{T}$).

Consider next the case of private information. As with public information, private information dampens the aggregate real effect of any given monetary shock. But unlike public information, private information does so in an imperfect manner, because the lack of common knowledge hinders the coordination of the firms' pricing decisions. At the same time, an increase in the precision of private information can exacerbate the cross-sectional misallocation of resources. It follows that more precise private information about monetary shocks can have a non-monotone welfare effect. We refer the reader to Hellwig (2005) for a more detailed analysis of this particular effect.

E. *Remarks*

We conclude this section with four remarks regarding the possible endogeneity of the information structure and applicability of our results.

Remark 1: Although our analysis has treated the information structure of the firms as an exogenous object, this does not necessarily limit the usefulness of our results. Suppose, for example, that some of the available public information obtains from the release of macroeconomic indicators or from policy actions. Alternatively, suppose that the available private information is the product of costly information acquisition, i.e., the informational friction is a symptom of inattention.[20] Under these scenarios, the precisions κ_x and κ_z become endogenous to the behavior of the firms, as well as to that of the monetary authority. How this endogeneity shapes the mapping from "deeper" parameters to the precisions κ_x and κ_z, or how it impacts the nature of the optimal monetary policy, is beyond the scope of this paper. We refer the reader to Amador and Weill (2010); Baeriswyl and Cornand (2010); and Paciello and Wiederholt (2014) for certain explorations in this direction. But no matter how κ_x and κ_z are determined in the first place, the anatomy of the welfare effects of information that our paper has offered remains valid in the following regard: the mapping from the precisions (κ_x, κ_z) to equilibrium welfare is invariant to the micro-foundations of the former.

Remark 2: The reinterpretation of the informational friction in terms of limited attention also explains why the policymaker may be unable to eliminate the friction even if she happens to know the state and can reveal it to the firms before the latter make their production and pricing decision. But even if the informational friction cannot itself be eliminated, whether it ultimately has a bite on real allocations and welfare still hinges on whether the rigidity is real or nominal: when the rigidity is only nominal, a monetary policy that replicates flexible prices, and *only* this policy, guarantees that the friction ceases to have a bite on welfare. This underscores, once again, the distinct normative implications of the two rigidities and the pivotal role of monetary policy vis-à-vis the nominal one.

Remark 3: The two policy benchmarks we characterized in Theorems 3 and 4 require that the policymaker observe the state of the economy. But as already mentioned, they do not presume that policymaker has an informational advantage over the firms at any time. It suffices that the policymaker commits to act in the future, after the state of the economy has become public information, provided of course that the market expects this to happen. This underscores the role of the policymaker in "managing expectations."

Remark 4: Our analysis has orthogonalized the information structure in three dimensions: one corresponding to preference and technologies; one corresponding to monopoly markups and other real distortions; and one corresponding to pure monetary shocks. When translating our results to certain applied contexts, however, it may be natural to consider signals that confound two or more of these distinct types of information. For instance, to the extent that business cycles are driven by a mixture of technology, markup, and monetary shocks, macroeconomic statistics will serve as a mixed signal of all these shocks, and their combined welfare effect will itself be a mixture of the effects we have documented. A similar point applies

[20] One can contemplate various micro-foundations of how agents allocate attention. See, e.g., Pavan (2014) for a flexible approach. But as long as one maintains a Gaussian specification (which is the golden standard in the related literature), the results of Bergemann and Morris (2013) imply that the resulting equilibrium allocations can always be replicated with an information structure like the one we assume.

to central bank communications insofar as they may contain only an overall assessment of the state of the economy.

V. Related Literature

In this section we seek a certain synthesis of the related literature, revisiting the prior works of Hellwig (2005); Walsh (2007); Lorenzoni (2010); and Baeriswyl and Cornand (2010) under the lens of our analysis.

As already noted, these papers rule out real rigidity, a scenario nested in our framework by letting $\theta = 0$. In this limit case, a monetary policy that replicates flexible prices implements the complete-information allocation. This is because the absence of real rigidity guarantees that the informational friction ceases to have a bite on real allocations and welfare once the "right" monetary policy is in place. Of course, such a policy may be unrealistic or undesirable. But the point we wish to make, for pedagogical reasons, is that the welfare effects reported in these prior works hinge entirely on deviations from the benchmark of replicating flexible prices: had monetary policy replicated flexible prices, the complete-information outcomes would have obtained.

Consider Hellwig (2005). Because that paper models M_t as an exogenous random walk and rules out shocks in preferences and technologies, all the volatility in M_t represents a monetary shock of the type described in Lemma 3. As noted in the previous section, the exercise conducted in Hellwig (2005) therefore boils down to studying the comparative statics of $\mathcal{T}$ with respect to the information that firms have about this shock, and the key unexpected finding is that private information has a non-monotone effect on relative-price dispersion and thereby on welfare as well. Clearly, these non-monotonicities reflect a discrepancy between the private and social value of coordination. But whereas the prior work appears to suggest that this discrepancy is innate to the Dixit-Stiglitz preference specification of modern macroeconomic models (see especially the discussion in Section 6.3 of Angeletos and Pavan 2007), our analysis clarifies that this discrepancy stems from the particular monetary policy assumed in Hellwig (2005).

Consider next Walsh (2007). That paper allows monetary policy to react systematically to shocks in preferences and markups (which Walsh interprets as, respectively, "demand shocks" and "cost-push shocks"). A deviation from the unconstrained optimum we studied in Section IV, however, obtains because of two types of policy frictions: the restriction that monetary policy can respond only to contemporaneous and noisy signals of the state of the economy, and the assumption that policy objectives differ from the model's ex ante utility. Unable to obtain analytic results, Walsh employs numerical simulations and arrives at a somewhat inconclusive answer to the question of interest, namely the welfare effects of central bank transparency. If instead one abstracts from the aforementioned policy frictions, a particularly sharp answer becomes available on the basis of our results: maximal transparency is desirable in the face of benign forces such as technology shocks, and maximal opacity is desirable in the face of distortionary forces such as markup shocks.

Consider next Lorenzoni (2010). In that paper, there are no markup shocks, the policymaker observes perfectly the state of the economy (with a lag), and monetary policy is set so as to maximize welfare. On the basis of our results, one may have expected monetary policy to replicate flexible prices and, in conjunction with the right noncontingent subsidy, to implement the first best. However, this is not the case because of the presence of an additional friction, a certain segmentation in consumer markets: each firm is matched with a random subset of consumers in the economy, and each consumer gets to see only a random subset of the prices in the economy. In the presence of this friction, a policy that implements the first-best response to the underlying aggregate shocks is still feasible, yet a distortion remains because prices cannot adjust to the idiosyncratic shocks induced by the random matching between the firms and the consumers. It is this distortion that drives the distinct welfare results reported in that paper.

Finally, consider Baeriswyl and Cornand (2010). Under the lens of our analysis, this paper makes two key assumptions. First, it requires that nominal GDP, $M_t \equiv P_t Y_t$, satisfy

$$\log M_t = \lambda_a (\bar{a}_t + \epsilon_{a,t}) + \lambda_\mu (\bar{\mu}_t + \epsilon_{\mu,t}), \tag{13}$$

where $\bar{a}_t$ and $\bar{\mu}_t$ are the underlying technology and markup shocks, $\epsilon_{a,t}$ and $\epsilon_{\mu,t}$ are exogenous noises, and λ_a and λ_μ are scalars under the control of the monetary authority. Second, it allows each firm to observe a noisy private signal of $\log M_t$. Baeriswyl and Cornand (2010) interpret M_t as the policy instrument, condition (13) as a policy rule, the noises $(\epsilon_{a,t}, \epsilon_{\mu,t})$ as measurement errors in the policymaker's *contemporaneous* observation of the underlying shocks, and the firms' signal of M_t as a signal of the policy action. The key contribution of that paper is then to study how the signaling role of monetary policy interacts with its stabilization role.

Our results qualify that paper's analysis in the following regard. Interpreting (13) as the policy rule overlooks the ability of the monetary authority to control current outcomes by committing to move interest rates in the future. Such commitment would not only improve the stabilization role of monetary policy by utilizing additional information that may arrive in the future but would also mute the signaling effect of current policy actions. Ruling out this possibility is therefore a key unstated assumption, although perhaps a realistic one, behind the core result of that paper regarding the trade-off between the stabilization and the signaling roles of monetary policy.

Putting aside this point and the precise interpretation of condition (13), this condition represents a restriction on the set of implementable allocations. This restriction drives the optimal policy in that paper away from the unconstrained optimum we characterized in Section IV. This fact in turn is the key to understanding why welfare in that paper depends on the firms' information about the technology shock despite the absence of real rigidity, as well as why welfare is a non-monotone function of the firms' information about the markup shock, in contrast to the monotone effect we obtained in part (i) of Lemma 5. Finally, because condition (13) is nested in Lemma 3 of our paper by letting $m_t = \lambda_a \epsilon_{a,t} + \lambda_\mu \epsilon_{\mu,t}$, the welfare effects of the signal that a firm receives about M_t can be understood under the lens of our analysis as the mixture of three kinds of information: information about the technology shock; information about the markup shock; and information about the policy "mistake" caused by measurement error.[21]

Let us close this section by noting the obvious: none of the preceding discussion is meant to downplay the contribution of the prior works. The mechanisms studied therein seem both intriguing and relevant. We nevertheless hope that our discussion has shed additional light on the inner workings of these mechanisms and on the assumptions that underlie them, thus also illustrating more generally how our paper can facilitate a useful anatomy of the welfare effects of information in baseline macroeconomic models.

VI. Conclusion

By assuming away incomplete information and strategic uncertainty, standard macroeconomic models presume that firms can perfectly coordinate their production and pricing decisions. By contrast, in this paper we allow an informational friction to inhibit this coordination and we study how this shapes the social value of information within an elementary business-cycle model. The key lessons can be summarized as follows:

- The welfare effects of information can be decomposed into two channels: the real rigidity that emerges as firms make production choices on the basis of incomplete information and the nominal rigidity that emerges as firms also set prices on the basis of such information.
- The first channel is present irrespective of the conduct of monetary policy. It also has sharp comparative statics: more information is welfare-improving through this channel if the business cycle is driven by "benign" forces such as technology shocks and welfare-deteriorating if it is driven by distortionary forces such as markup shocks.
- By contrast, the second channel hinges on the conduct of monetary policy. As in the New-Keynesian framework, there is a policy that neutralizes the nominal rigidity. At the flexible-price benchmark, the welfare effects of information are shaped solely by the real-rigidity channel. Away from it, they hinge on whether the provision of more information dampens or amplifies the deviation of monetary policy and on whether that deviation was desirable to begin with.
- When the business cycle is driven by technology shocks, a monetary policy that replicates flexible prices is optimal. When, instead, the business cycle is driven by markup shocks, a deviation from this benchmark is desirable. More information then tends to decrease welfare not only because it exacerbates the inefficiency of the underlying flexible-price fluctuations but also because it curtails the monetary authority's ability to combat these fluctuations.

We view the sharpness of these lessons and their close connection to familiar normative properties of RBC and New-Keynesian models as the main strengths of our contribution. This sharpness, however, comes at a cost. By narrowing the

[21] The welfare effects of the first two types of information are those mentioned above. The third one has a non-monotone effect for the reason first explained in Hellwig (2005).

analysis within the context of an elementary model, we preclude any quantitative assessment. By treating the information structure as exogenous, we bypass the question of either how information gets collected or what policy instruments can affect it. Finally, while we allow the informational friction to inhibit the coordination of production and pricing decisions of firms, we assume away any such friction in, say, the consumption and saving choices of households or the trades of financial investors. The bite of incomplete information on the social efficiency of the latter kind of economic decisions, and the implications of this for the business cycle, is an important open research question.

REFERENCES

Amador, Manuel, and Pierre-Olivier Weill. 2010. "Learning from Prices: Public Communication and Welfare." *Journal of Political Economy* 118 (5): 866–907.

Angeletos, George-Marios, and Jennifer La'O. 2010. "Noisy Business Cycles." In *NBER Macroeconomics Annual 2009*, edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, 319–78. Chicago: University of Chicago Press.

Angeletos, George-Marios, and Jennifer La'O. 2013. "Sentiments." *Econometrica* 81 (2): 739–79.

Angeletos, George-Marios, and Jennifer La'O. 2014. "Optimal Monetary Policy with Informational Frictions." Unpublished.

Angeletos, George-Marios, and Alessandro Pavan. 2007. "Efficient Use of Information and Social Value of Information." *Econometrica* 75 (4): 1103–42.

Baeriswyl, Romain, and Camille Cornand. 2010. "The Signaling Role of Policy Actions." *Journal of Monetary Economics* 57 (6): 682–95.

Bergemann, Dirk, and Stephen Morris. 2013. "Robust Predictions in Games with Incomplete Information." *Econometrica* 81 (4): 1251–1308.

Blanchard, Olivier Jean, and Nobuhiro Kiyotaki. 1987. "Monopolistic Competition and the Effects of Aggregate Demand." *American Economic Review* 77 (4): 647–66.

Galí, Jordi. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework.* Princeton: Princeton University Press.

Hellwig, Christian. 2005. "Heterogeneous Information and the Welfare Effects of Public Information Disclosures." Unpublished.

James, Jonathan G., and Phillip Lawler. 2011. "Optimal Policy Intervention and the Social Value of Public Information." *American Economic Review* 101 (4): 1561–74.

Lorenzoni, Guido. 2009. "A Theory of Demand Shocks." *American Economic Review* 99 (5): 2050–84.

Lorenzoni, Guido. 2010. "Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information." *Review of Economic Studies* 77 (1): 305–38.

Lucas, Robert E., Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (2): 103–24.

Mackowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99 (3): 769–803.

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328.

Morris, Stephen, and Hyun Song Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (5): 1521–34.

Morris, Stephen, Hyun Song Shin, and Hui Tong. 2006. "Social Value of Public Information: Morris and Shin (2002) Is Actually Pro-transparency, Not Con: Reply." *American Economic Review* 96 (1): 453–55.

Paciello, Luigi, and Mirko Wiederholt. 2014. "Exogenous Information, Endogenous Information, and Optimal Monetary Policy." *Review of Economic Studies* 81 (1): 356–88.

Pavan, Alessandro. 2014. "Attention, Coordination, and Bounded Recall." Unpublished.

Roca, Mauro. 2005. "Transparency and Monetary Policy with Imperfect Common Knowledge." Unpublished.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.

Svensson, Lars E. O. 2006. "Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-transparency, Not Con." *American Economic Review* 96 (1): 448–52.

Walsh, Carl E. 2007. "Optimal Economic Transparency." *International Journal of Central Banking* 3 (1): 5–36.

Woodford, Michael. 2002. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, edited by Philippe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford, 25–58. Princeton: Princeton University Press.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.