Business cycle fluctuations and monetary policy under heterogeneous information

Romain Baeriswyl, Pierrick Clerc, Camille Cornand

GATE LYON SAINT-ÉTIENNE
UMR 5824

Bâtiment MILC
35, rue Raulin
69007 Lyon - France

Maison de l'Université, Bâtiment B
10, rue Tréfilerie
42023 Saint-Etienne cedex 02 - France

https://www.gate.cnrs.fr
gate@gate.cnrs.fr

WP 2518 – July 2025

Business cycle fluctuations and monetary policy under heterogeneous information

Romain Baeriswyl, Pierrick Clerc, Camille Cornand

Abstract:

Modelling business cycle fluctuations and the non-neutrality of money is an ongoing challenge for macroeconomists. While the imperfect-information hypothesis developed by Phelps and Lucas in the 1970s had been abandoned in favour of the sticky-price hypothesis, a recent trend in the literature has put this hypothesis back on the agenda to explain business fluctuations and monetary non-neutrality. The success of this revival lies in the introduction of strategic uncertainty into a framework of heterogeneous information. The current paper presents this macroeconomic framework and provides a brief overview of recent advances in the literature.

Keywords:
heterogeneous information, strategic complementarities, higher-order beliefs, coordination, beauty-contest, monetary non-neutrality, business cycle, communication policy, monetary policy

JEL codes:
C72, D83, E12, E32, E52

Working Paper

Business cycle fluctuations and monetary policy under heterogeneous information*

Romain Baeriswyl[§] Pierrick Clerc[♯] Camille Cornand[†]

July 2025

Abstract

Modelling business cycle fluctuations and the non-neutrality of money is an ongoing challenge for macroeconomists. While the imperfect-information hypothesis developed by Phelps and Lucas in the 1970s had been abandoned in favour of the sticky-price hypothesis, a recent trend in the literature has put this hypothesis back on the agenda to explain business fluctuations and monetary non-neutrality. The success of this revival lies in the introduction of strategic uncertainty into a framework of heterogeneous information. The current paper presents this macroeconomic framework and provides a brief overview of recent advances in the literature.

JEL classification: C72, D83, E12, E32, E52.

Keywords: heterogeneous information, strategic complementarities, higher-order beliefs, coordination, beauty-contest, monetary non-neutrality, business cycle, communication policy, monetary policy.

*The views expressed in this paper are strictly those of the authors and do not necessarily reflect those of the Swiss National Bank.

[§]Swiss National Bank, Boersenstrasse 15, 8001 Zurich, Switzerland; email: romain.baeriswyl@snb.ch

[♯]HEC Liège, Rue Louvrex 14, 4000 Liège, Belgien; Email: Pierrick.Clerc@uliege.be

[†]CNRS, Université Lumière Lyon 2, Université Jean Monnet Saint-Etienne, emlyon business school, GATE, 69007, Lyon, France. Email: cornand@gate.cnrs.fr

1 Introduction

Modelling business cycle fluctuations and the non-neutrality of money is an ongoing challenge for macroeconomists. There is still a debate among them as to the most plausible frictional hypothesis to generate a model consistent with the empirical evidence. While the imperfect-information hypothesis developed by Phelps and Lucas in the 1970s had been abandoned in favour of the sticky-price hypothesis, a recent trend in the literature has put this hypothesis back on the agenda to explain business fluctuations and monetary non-neutrality. The success of this revival lies in the introduction of strategic uncertainty into a framework of heterogeneous information. Business cycle theories relying on this assumption state that economic agents face fundamental uncertainty about economic conditions as well as strategic uncertainty about how other market participants perceive these economic conditions. As a result, decisions tend to adjust slowly to changes in economic fundamentals, and small or temporary shocks can have large, lasting effects on macroeconomic aggregates, which is in line with empirical observations.

In another vein, recent advances in game theory under heterogeneous information have led to a renewed interest in incomplete information models. These models typically give a predominant role to higher-order beliefs in the decision-making process of agents, which leads to persistence in their reaction to new information. While the new information provides private agents with a view on exogenous variables, it only provides them with an ambiguous view on how other agents change their expectations about these exogenous variables and an even more ambiguous view on how other agents change their expectations about the expectations of others. The conjunction of strategic complementarities and heterogeneous information thus results in a powerful macroeconomic framework.

The current paper presents this macroeconomic framework and provides a brief overview of recent advances in the literature. Section 2 illustrates how higher-order beliefs lead to an overreaction to public information, a mechanism that was introduced by Morris and Shin (2002). It provides technical details to solve a simple class of models based on beauty-contest games. Section 3 relates to monetary non-neutrality under heterogeneous information and explains how Woodford (2003a) initiated the revival of the heterogeneous-information hypothesis by adding strategic complementarities to the model of Lucas (1972). It further discusses the contribution of Maćkowiak and Wiederholt (2009) and Pasten and Schoenle (2016) in rationalising the limited attention that firms pay to economic shocks. Section 4 explains how dispersed information impinges on the transmission of real shocks and business cycle fluctuations. For instance, Angeletos and La'O (2010) show that the nature of signals observed by economic agents is critical in the propagation of these shocks, and Angeletos and La'O (2013) highlight that non-fundamental shocks, such as sentiment shocks, can have significant economic effects through higher-order beliefs. Finally, section 5 deals with communication and stabilisation policies. It exposes the debate about the welfare effect of public information between Morris and Shin (2002), Svensson (2006) and Angeletos and Pavan (2007), before discussing how information can be disclosed in order to reduce the overreaction of market participants. The section discusses then two extensions, which have important implications on the optimal communication policy. The first is to consider that information is endogenous, either because the agents' source of information is an economic variable which itself depends on the agents' actions (e.g. Amador and Weill (2010)), or because the agents freely choose the attention they give to different information sources (e.g. Hellwig and Veldkamp (2009)). The second extension, following James and Lawler (2011), is to add a policy action to the framework and to solve jointly for the optimal communication and stabilisation policy.

Although brief, this overview of the literature reveals the potential that the conjunction of strategic complementarities and heterogeneous information offers to explain macroeconomic phenomena, and suggests a promising future for this field of economics.

2 A beauty-contest model

Strategic complementarities are widespread in the economy: often agents tend to align their action with that of others. For instance, an investor may decide to invest in a company based on what other investors think about that company; a depositor may trust her bank to the extent that other depositors trust it as well; or a firm may set its price by anticipating the price set by its competitors. Such games are often referred to as Keynesian beauty contests. Keynes (1936) compared the cognitive exercise of an investor in the financial markets to that of a participant in a newspaper beauty contest of the 1930s where participants were rewarded for selecting the most popular faces among all participants, rather than those they might personally find the most attractive. In these contests, Keynes stressed that "it is not a case of choosing those [faces] that, to the best of one's judgment, are really the prettiest, nor even those that average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees."

Bearing in mind the very general scope of the beauty-contest analogy, we illustrate the mechanism by referring to the price setting problem of firms in monopolistic competition. In a micro-founded economy[1] populated by a representative household and a continuum of monopolistically competitive firms, the demand side – expressed in logarithms – is summarised by the following equation:

$$q = p + y$$

where q is the level of aggregate nominal spending, p the price level, and y the level of aggregate real output (or output gap). The supply side is determined by the optimal price set by firm j

$$p_j = \mathbb{E}_j\left[(1-\xi)y + p\right] = \mathbb{E}_j\left[(1-\xi)q + \xi p\right] \tag{1}$$

where $\mathbb{E}_j$ represents the expectation operator of firm j conditional on its information and

[1]See e.g. Adam (2007) or Clerc and Dos Santos Ferreira (2022).

$\int_j p_j = p$. Each firm sets its price as a function of its own expectation of the nominal aggregate demand and its expectation about the price level. The parameter ξ determines to what extent the optimal price responds to the output gap and is increasing in the risk aversion of the household. As is generally the case in macroeconomic models, we assume that prices are strategic complements, so that $0 < \xi < 1$. The parameter ξ is the degree of strategic complementarities.

2.1 Information structure

In an environment characterised by strategic complementarities, a dual motive guides the economic agents in their action. An agent chooses his action to be close to the economic fundamental and, at the same time, close to the average action of other agents. Equation (1) says that each firm sets its price based on its expectation of nominal aggregate demand q and of the price level p. This dual motive gives rise to a dual role for information because information helps not only to forecast the economic fundamental, but also to forecast the action of others. Whereas the best forecast of economic fundamental relies on *accurate* information, the best forecast of others' action relies on *common* information.

Assume that nominal aggregate demand follows an improper prior distribution on the reals and that each firm j receives two signals on nominal demand. The first signal is centered on the true value of q and contains an error term that is normally distributed and that is identically but independently drawn for each firm:

$$q_j = q + \varepsilon_j, \qquad \text{with} \qquad \varepsilon_j \sim N\left(0, \sigma_\varepsilon^2\right) \tag{2}$$

The first signal is private as its error term is entirely specific to each firm. This private signal can be interpreted as an estimate of the demand that each firm makes from its own business activity.

The second signal also is centered on the true value of q and contains an error term that is normally distributed. However, all firms observe the same second signal:

$$D = q + \eta, \qquad \text{with} \qquad \eta \sim N\left(0, \sigma_\eta^2\right) \tag{3}$$

The second signal is public and common knowledge among firms. This public signal is typically disclosed by a public institution such as the central bank, and the error term η captures its perception error about the fundamental.

The distribution of error terms ε_j and η is common knowledge among firms.

2.2 Equilibrium

To derive the equilibrium behavior of firms, we need to compute the expectations of firm j about demand and the price level conditional on its information set. The expectation of demand $\mathbb{E}_j(q)$ writes

$$\mathbb{E}\left(q|q_j, D\right) = \frac{\sigma_\eta^2}{\sigma_\varepsilon^2 + \sigma_\eta^2} q_j + \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\eta^2} D = (1-\Omega) q_j + \Omega \cdot D \tag{4}$$

The best estimate of firm j about demand is an average of its signals, each weighted with Ω depending on their relative precision.

Deriving the conditional expectation of firm j about the price level is less straightforward because one has to form expectations about the behavior of other firms. There are two ways to derive the equilibrium, the *recursive method* and the *method of indeterminate coefficients*. Though equivalent, we present both methods for pedagogical purposes.

2.2.1 The recursive method

The recursive method consists in computing the weighted average of first- and higher-order expectations of firm j about demand. The pricing rule (1) can be rewritten by recursively substituting the average price by higher-order expectations on demand:

$$\begin{aligned}
p_j &= \mathbb{E}_j\big[(1-\xi)q + \xi \underbrace{\bar{\mathbb{E}}\left[(1-\xi)q + \xi p\right]}_{p}\big] \\
&= \mathbb{E}_j\left[(1-\xi)q + \xi\bar{\mathbb{E}}\left[(1-\xi)q + \xi\bar{\mathbb{E}}\left[(1-\xi)q + \xi\bar{\mathbb{E}}\left[(1-\xi)q + \ldots\right]\right]\right]\right] \\
&= (1-\xi)\sum_{k=0}^{\infty} \xi^k \mathbb{E}_j\left[\bar{\mathbb{E}}^k(q)\right] \qquad (5)
\end{aligned}$$

where k is the degree of higher-order iteration, $\bar{\mathbb{E}}$ is the average expectation operator in the population such that $\bar{\mathbb{E}}(.) = \int_j \mathbb{E}_j(.)dj$ and we use the following notation for higher-order beliefs: $\bar{\mathbb{E}}^0(x) = x$ is the variable x, $\bar{\mathbb{E}}^1(x) = \bar{\mathbb{E}}(x)$ is the average expectation of x, $\bar{\mathbb{E}}^2(x) = \bar{\mathbb{E}}\bar{\mathbb{E}}^1(x) = \bar{\mathbb{E}}\bar{\mathbb{E}}(x)$ is the average expectation of the average expectation of x, and so on.

Based on the first-order expectation of demand (4), expectations of higher orders are

$$\begin{aligned}
\mathbb{E}_j\bar{\mathbb{E}}(q) &= (1-\Omega)[(1-\Omega)q_j + \Omega D] + \Omega D = (1-\Omega)^2 q_j + [1-(1-\Omega)^2]D \\
\mathbb{E}_j\bar{\mathbb{E}}^{(2)}(q) &= (1-\Omega)^2[(1-\Omega)q_j + \Omega D] + [1-(1-\Omega)^2]D = (1-\Omega)^3 q_j + [1-(1-\Omega)^3]D \\
\mathbb{E}_j\bar{\mathbb{E}}^{(k-1)}(q) &= (1-\Omega)^k q_j + [1-(1-\Omega)^k]D \qquad (6)
\end{aligned}$$

When information is heterogeneous across agents, an agent assigns more weight to common information in his expectation of others' expectation about fundamentals than in his own expectation about fundamentals: the law of iterated expectations fails and the weight assigned to common information increases with the degree of strategic complementarities. The weight assigned to the public signal, D, increases in expectations of higher orders because this signal better helps to predict others' expectation as its error term is common to all firms. These higher-order expectations receive themselves an increasing weight in price setting with higher degrees of strategic complementarities. Plugging higher-order expectations into (5), the pricing rule yields

$$
\begin{aligned}
p_j &= (1-\xi)\sum_{k=0}^{\infty}\xi^k\left[(1-\Omega)^{k+1}q_j+[1-(1-\Omega)^{k+1}]D\right]=\frac{(1-\xi)(1-\Omega)}{(1-\xi)+\xi\Omega}q_j+\frac{\Omega}{(1-\xi)+\xi\Omega}D \\
&= \frac{(1-\xi)\sigma_\eta^2}{\sigma_\varepsilon^2+(1-\xi)\sigma_\eta^2}q_j+\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2+(1-\xi)\sigma_\eta^2}D=(1-\omega)q_j+\omega D \qquad (7)
\end{aligned}
$$

The firm assigns a larger weight on the public signal D in its price setting than in its first-order expectation about the fundamental: $\omega > \Omega$. The firm is said to *overreact* to common information in the sense that, because of the coordination motive, it gives more weight to it than its accuracy would actually justify. This overreaction increases with the degree of strategic complementarities. Rational though from the agent's point of view, this overreaction can cause welfare losses, as emphasised by Morris and Shin (2002) and explained in section 5.

2.2.2 The method of indeterminate coefficients

The method of indeterminate coefficients consists in stating that the equilibrium price is a linear combination of the private and public signals with indeterminate weights, $p_j = (1-\omega)q_j + \omega D$, and solving for this guess in the pricing rule (1):

$$
\begin{aligned}
p_j &= \mathbb{E}_j\left[(1-\xi)q+\xi\left[(1-\omega)q+\omega D\right]\right] \\
&= \underbrace{\left[(1-\xi\omega)(1-\Omega)\right]}_{1-\omega}q_j+\underbrace{\left[\Omega+\xi\omega(1-\Omega)\right]}_{\omega}D \\
&= \frac{(1-\xi)(1-\Omega)}{(1-\xi)+\xi\Omega}q_j+\frac{\Omega}{(1-\xi)+\xi\Omega}D=\frac{(1-\xi)\sigma_\eta^2}{\sigma_\varepsilon^2+(1-\xi)\sigma_\eta^2}q_j+\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2+(1-\xi)\sigma_\eta^2}D
\end{aligned}
$$

The method of indeterminate coefficients yields the same solution as the recursive method in (7), indicating that the linear equilibrium is the unique equilibrium. While the method of indeterminate coefficients is shorter to derive, the recursive method explicitly exposes the increasing weight of public information in higher-order expectations, which helps to understand the role of common information in strategic complementarities.

With this simple illustration in mind, we can now look at how the introduction of strategic complementarities in macroeconomic models has led to a revival of the imperfect-information hypothesis as the source of the monetary non-neutrality.

3 Monetary non-neutrality

The 'islands' model proposed by Lucas (1972) was the first model attempting to provide an explanation of monetary non-neutrality based on heterogeneous information. This model, however, had a major drawback: it was not able to generate persistent effects of monetary shocks on real variables. Lucas's model indeed assumed that firms became fully aware of aggregate demand shocks one period after their occurrence. As a result, all nominal magnitudes were fully adjusted after a single period, so that the effects on real variables (however large) completely vanished at that point. By contrast, a well-known fact about monetary shocks – extensively documented in the VAR literature – is that they have highly persistent real effects.[2] Lucas's model was therefore unable to account for a crucial dimension of monetary non-neutrality.

Michael Woodford, in a paper prepared for a conference held in 2001 in honor of Edmund Phelps (and published as Woodford 2003a), proposed to reconsider what he called the 'Phelps-Lucas hypothesis' – according to which the real effects from nominal shocks stem from dispersed information. Woodford argued that two important departures from Lucas's model would make this hypothesis consistent with persistent monetary non-neutralities. The first departure deals with the structure of information. In Lucas's model, firms become aware of the current level of aggregate demand as soon as this level becomes public information (which is assumed to happen after one period). This means that firms are (implicitly) supposed to be able to fully absorb and process all available information. Woodford, instead, follows a line of argumentation initially developed by Sims (1998, 2003), according to which firms would have limited capacity for processing information. In Woodford's model, even though information about the current state of aggregate demand is immediately published, firms cannot pay more than limited attention to this information. Hence, it takes time for firms to process available information, and to integrate it into their information set.

The second departure deals with the structure of markets. In Lucas's model, firms operate in perfectly competitive markets. Woodford's model instead assumes a monopolistic competition framework featuring strategic complementarities in price setting. Individual firms thus play a 'beauty-contest game' with their competitors, in which their 'action' consists in setting the price of the item they sell. In such an environment, firms find it optimal to set their price by reacting not only to their own expectations about current aggregate demand, but also (since they are uncertain about the expectations of other firms) to their 'higher-order expectations' – i.e., their expectations about other firms' expectations regarding current aggregate demand, their expectations about other firms' expectations about other firms' expectations, and so on. Therefore, besides their incapacity to know the current level of aggregate demand (due to their limited capacity for processing information), firms do not know either what other firms know about this level. In other words, Woodford's model involves what game theory calls imperfect common knowledge. The fact that higher-order expectations matter for price setting is a central aspect of Woodford's strategy to explain persistence of real effects of monetary shocks, because higher-order expectations display even more inertia in response to aggregate demand shocks than firms' own expectations about these shocks.[3]

[2] Christiano et al. (2005), for instance, find that aggregate output returns to its initial level no sooner than three years after the occurrence of an expansionary monetary shock.

[3] The importance of higher-order expectations has been initially raised by Phelps (1983) and Townsend (1983), but not for the study of the non-neutrality of money.

3.1 Imperfect common knowledge

Following Clerc and Dos Santos Ferreira (2022), we can formalise the main insights of Woodford's model as follows. Expressing all variables in logarithms and introducing a time subscript t, the demand side is summarised by the following equation:

$$q_t = p_t + y_t = q_{t-1} + \mu + \varepsilon_{q,t} \tag{8}$$

where q_t denotes the level of aggregate nominal spending (in period t), p_t the price level, y_t the level of aggregate output, μ the drift in aggregate nominal spending, and $\varepsilon_{q,t}$ the shocks on aggregate demand. These shocks follow a Gaussian white noise process with variance σ_q^2. We assume that monetary shocks are the single source of aggregate demand disturbances, so that $\varepsilon_{q,t}$ actually denotes monetary shocks.

The supply side is composed of a continuum of firms in monopolistic competition. If there were no information frictions, firm j would set its period t price, $p_{j,t}^*$, according to:

$$p_{j,t}^* = \xi p_t + (1 - \xi)(q_t - y^n) \tag{9}$$

where $\xi < 1$ denotes a reduced-form parameter measuring the degree of strategic interactions in price setting, and y^n the level of potential output (assumed to be constant).[4] There is strategic complementarity in price setting when $\xi > 0$ (since the optimal price of a given firm positively depends on the prices set by the other firms), while there is strategic substitutability in price setting when $\xi < 0$.

All information about the aggregate state of the economy is publicly available to every firm. However, because of finite information-processing capacity, firms observe aggregate variables only imperfectly. The price actually set by firm j in period t, $p_{j,t}$, is thus given by:

$$p_{j,t} = \xi \mathbb{E}_{j,t} p_t + (1 - \xi)(\mathbb{E}_{j,t} q_t - y^n) \tag{10}$$

where $\mathbb{E}_{j,t}$ denotes the expectations operator conditional upon the information available to firm j in period t. Firms can solely observe a noisy private signal about current aggregate demand. The signal observed by firm j in period t, denoted by $s_{q,j,t}$, is represented as follows:

$$s_{q,j,t} = q_t + e_{q,j,t} \tag{11}$$

with $e_{q,j,t}$ following a Gaussian white noise process with variance $\tilde{\sigma}_q^2$, independent from $\varepsilon_{q,j,t}$, and independently distributed across firms. The precision of the signal is therefore given by $1/\tilde{\sigma}_q^2$: the higher $1/\tilde{\sigma}_q^2$, the more informative is the signal.

Since firms receive different signals (and ignore the signals received by others), information about current aggregate demand is dispersed.

Firms never know exactly the true value of aggregate demand, even after many periods. As a result, they can only use the history of their private signals up to period t in order

[4] Equation (9) is actually the combination of firm j's log-linear approximation of the first-order condition, namely $p_{j,t}^* = p_t + (1 - \xi)(y_t - y^n)$, and $q_t = p_t + y_t$.

to form expectations about q_t. In the case of firm j, we have:

$$\mathbb{E}_{j,t}q_t = \mathbb{E}[q_t|s_{q,j,t}, s_{q,j,t-1}, s_{q,j,t-2}, s_{q,j,t-3}, \ldots] \tag{12}$$

However, equation (10) implies that the price set by a given firm also depends on this firm's expectation of the average level of prices charged by the other firms. The combination of dispersed information and strategic interactions in price setting thus induces firms to play beauty-contest games with their competitors.

Since all firms set prices in a similar way, each firm has to form an expectation of the average expectation among the other firms about q_t, an expectation of the average expectation of that average expectation, and so on. Hence, on top of its own expectation of q_t, every firm has to form higher-order expectations about q_t.

Using the same notation as in (6), equation (10) can then be rewritten as:

$$p_{j,t} = (1-\xi)\left[\sum_{k=0}^{\infty}\xi^k \mathbb{E}_{j,t}^{(k+1)}q_t\right] - (1-\xi)y^n \tag{13}$$

The importance of higher-order expectations for pricing decisions critically depends on the value of ξ. When $\xi = 0$, there are no strategic interactions in price setting since (according to equation (10)) the price set by a given firm is independent from the prices set by the other firms. In this case, equation (13) makes it clear that pricing decisions only depend on firms' own expectations of q_t (i.e., on their first-order expectations of q_t), $\mathbb{E}_{j,t}q_t$. By contrast, when $\xi > 0$, there are strategic complementarities in price setting. In this case – which is the one usually retained in the pricing literature – equation (13) shows that higher-order expectations not only matter for pricing decisions, but also increasingly influence those decisions.[5] Indeed, the stronger strategic complementarities (i.e., the closer ξ is to 1), the higher the order of expectations firms have to consider, and the larger the weights firms have to attach to expectations of higher orders. Intuitively, this stems from the fact that the stronger the strategic elements between firms, the more important become the expectations held by the other firms in the decisions taken by a given firm.

Moreover, it can be shown that expectations of higher order react less strongly to private signals than expectations of lower order. For instance, when a positive monetary shock occurs, the resulting rise in the private signal observed by each individual firm induces a rise in third-order expectations which is more sluggish than the rise in second-order expectations, itself more sluggish than the rise in first-order expectations.

Finally, applying optimal filtering techniques, averaging the various equations over j,

[5]Several sources of strategic complementarities in price setting have been proposed, among which stand out decreasing returns to scale, firm-specific input markets, input-output linkages across sectors, and real wage rigidities at the aggregate level. See Woodford (2003b, Chapter 3) and Leahy (2011) for a thorough discussion on these sources.

and making technical assumptions to ensure the uniqueness of equilibrium, yields:

$$y_t - y^n = \kappa(y_{t-1} + \varepsilon_{q,t}) \quad \text{with } \kappa = 1 - 1/2 - (1-\xi)\frac{\sigma_q^2}{\tilde{\sigma}_q^2} + [((1-\xi)\frac{\sigma_q^2}{\tilde{\sigma}_q^2})^2 + 4(1-\xi)\frac{\sigma_q^2}{\tilde{\sigma}_q^2}]^{1/2} \tag{14}$$

From equation (14), we can see that y_t depends on both $\varepsilon_{q,t}$ and y_{t-1}. As a result (and contrary to what happens in Lucas's model, where it could be shown that y_t depends solely on $\varepsilon_{q,t}$), aggregate output responds not only to current realisations of monetary shocks, but also – and crucially – to past realisations of these shocks. This means that in Woodford's model, monetary shocks have persistent real effects.

The quantitative importance of this persistence is driven by κ. This coefficient is made of two components, each encapsulating a key mechanism at work. The first component is the precision of the signal, $1/\tilde{\sigma}_q^2$, illustrating the role played by the informativeness of the signal observed: when this signal is poorly informative (i.e., when $1/\tilde{\sigma}_q^2$ is low), it takes a very long time for firms to learn that a monetary shock has occurred (even if the information on this occurrence is available from the start), and to adjust their price accordingly.[6] Hence, aggregate output displays persistent deviations from its potential level. The second component is the parameter ξ, illustrating the role played by strategic complementarities in price setting. When these complementarities are strong (i.e., when ξ is close to 1), firms need to form expectations of very high orders, and to attach large weights to them in their pricing decisions. At the same time, expectations of higher orders react more sluggishly to information than expectations of lower orders. Strong strategic complementarities therefore slow down the response of prices to monetary shocks, leading to persistent deviations of aggregate output from potential.

Furthermore, and quite importantly, κ can be made arbitrarily large for values of ξ sufficiently close to 1, whatever the value of $1/\tilde{\sigma}_q^2$. This implies that when strategic complementarities in price setting are very strong, monetary shocks can have highly persistent effects on aggregate output even if firms are individually very well informed about the occurrence of these shocks. When prices are very strong strategic complements, indeed, the price set by any given firm is almost entirely determined by its higher-order expectations (and especially by those of very high orders), and only marginally by its first-order expectations. Individual firms thus find it optimal to adjust their price only sluggishly (chiefly reflecting the dynamics of expectations of very high orders), even if they are well aware that a monetary shock has occurred (i.e., even if first-order expectations adjust fairly quickly). Hence, provided a sufficiently high strategic uncertainty, highly persistent monetary non-neutralities can be obtained for a relatively low fundamental uncertainty.

3.2 Rational inattention

We have just seen that in Woodford's model, a poorly informative private signal – embodied in a low value of the signal precision $1/\tilde{\sigma}_q^2$ – is a central mechanism through which

[6] Firms, however, will never be certain that such a shock has occurred since they never know exactly the true value of aggregate demand.

monetary shocks produce persistent real effects. A poorly informative signal is meant to represent the limited attention firms pay to the current state of aggregate demand when they have limited capacity for processing information. In this model, however, $1/\tilde{\sigma}_q^2$ is a free parameter whose value can be arbitrarily selected by the modeler.[7] This means that retaining a low value for $1/\tilde{\sigma}_q^2$ amounts, in fact, to imposing that firms pay only limited attention to the state of aggregate demand. Therefore, an important source of persistent monetary non-neutralities in Woodford's model turns out to be nothing more than postulated.

The main contribution of Maćkowiak and Wiederholt (2009) is to explain the persistent real effects of monetary shocks by the rational decision of firms to pay limited attention to the state of aggregate demand. In the framework set out in that paper, firms allocate their attention to track both idiosyncratic and aggregate demand shocks. Maćkowiak and Wiederholt show that the allocation chosen by firms depends on the relative impact of these shocks on the prices firms would set absent information frictions. In this context, the signal precision $1/\tilde{\sigma}_q^2$ is no longer a free parameter: it is a variable chosen by firms as a function of the structural parameters driving the volatility of shocks and the responses of firms' frictionless optimal prices. For realistic values of these parameters, the resulting $1/\tilde{\sigma}_q^2$ is found to be very low, which means that it is actually optimal for firms to pay very limited attention to information about aggregate demand shocks. This rational inattention with respect to the state of aggregate demand implies that firms adjust very sluggishly their prices in response to monetary shocks, inducing large and persistent fluctuations in aggregate output.

The framework proposed by Maćkowiak and Wiederholt is basically the same as Woodford's. There are, however, three important differences. First, the price that would be set absent information frictions also depends on idiosyncratic (cost or demand) conditions. Firm j's frictionless optimal price in period t becomes:

$$p_{j,t}^* = \xi p_t + (1-\xi)(q_t - y^n) + \varepsilon_{z,j,t} \tag{15}$$

where $\varepsilon_{z,j,t}$ denotes an idiosyncratic state variable reflecting conditions specific to firm j. Firm-specific conditions follow a Gaussian white noise process with variance σ_z^2, independent from $\varepsilon_{q,t}$, and independently distributed across firms.

Second, the flow of information that firms can effectively process in a given period is now explicitly introduced. This flow, denoted by κ, reflects firms' limited capacity for processing information (the lower firms' capacity, the lower κ). Moreover, firms can allocate this information flow – i.e., their attention – to track both idiosyncratic and aggregate demand shocks. The flow of information regarding idiosyncratic shocks is denoted by κ_z, while the flow of information regarding aggregate demand shocks is denoted by κ_q. Firms choose the values of κ_q and κ_z so as to maximise their expected profits, under the

[7] By contrast, the parameter ξ – which encapsulates the second source of persistence in Woodford's framework – is a reduced-form parameter determined by the structural parameters of the model.

constraint:[8]

$$\kappa = \kappa_q + \kappa_z \tag{16}$$

As in Woodford's model, information processing limitations prevent firms from observing the precise values of aggregate variables. In particular, firms solely observe a noisy private signal about current aggregate demand. Firm j's signal about aggregate demand in period t is still denoted by $s_{q,j,t}$, and still given by equation (11). In addition, Maćkowiak and Wiederholt assume that information processing limitations also prevent firms from observing the precise values of idiosyncratic variables: firms solely observe a separate noisy private signal about their current idiosyncratic conditions. This signal is modeled in the same fashion as the signal about aggregate demand. Hence, firm j's signal about its own conditions in period t, denoted by $s_{z,j,t}$, is:

$$s_{z,j,t} = \varepsilon_{z,j,t} + e_{z,j,t} \tag{17}$$

with $e_{z,j,t}$ following a Gaussian white noise process with variance $\tilde{\sigma}_z^2$, independent from $\varepsilon_{z,j,t}$, and independently distributed across firms. The precision of this signal is therefore $1/\tilde{\sigma}_z^2$.

Third, the information structure is now endogenous. Firms, indeed, can choose the precision of the signals received. This latter depends on firms' attention: the more firms pay attention to a particular type of shocks, the less noisy the associated signal. For instance, the more firms pay attention to aggregate demand shocks (i.e., the higher κ_q), the more precise is the signal $s_{q,j,t}$ (i.e., the higher is $1/\tilde{\sigma}_q^2$).

While more complicated to derive analytically, the resulting dynamics of aggregate output are actually very similar to those induced by Woodford's model, and can thus be approximated by equation (14). This means that, exactly as in that model, monetary shocks generate real effects whose persistence is determined by two mechanisms – respectively encapsulated by ξ and $1/\tilde{\sigma}_q^2$). In Maćkowiak and Wiederholt's model, however, $1/\tilde{\sigma}_q^2$) is no longer a free parameter: it is a variable whose value results from the amount of attention firms find it optimal to pay to aggregate demand shocks. Hence, the persistence produced by this model crucially depends on how firms allocate their attention between idiosyncratic and aggregate demand shocks.

Firms' allocation of attention is driven by the relative impact of idiosyncratic and aggregate demand shocks on frictionless optimal prices. This relative impact is, in turn, essentially determined by two factors. The first factor is the relative volatility of aggregate-demand and idiosyncratic shocks, i.e. σ_q^2/σ_z^2. This is intuitive: if, for instance, aggregate demand shocks are more volatile than idiosyncratic shocks (i.e., if $\sigma_q^2/\sigma_z^2 > 1$), aggregate demand shocks will represent a relatively more important source of fluctuations for frictionless optimal prices, creating an incentive for firms to track these shocks with more accuracy. The second factor is the degree of strategic interactions in price setting, i.e.,

[8] In their baseline model, Maćkowiak and Wiederholt assume that firms choose the optimal values of κ_q and κ_z once and for all, in period $t = 0$.

ξ. Let us assume that $\sigma_q^2/\sigma_z^2 > 1$, so that firms allocate most of their attention to track aggregate demand shocks. In this case, firms quickly adjust their prices in response to these shocks, since they are quite aware of their occurrence. As a result, the price level displays large fluctuations in response to aggregate demand shocks. If there are strong strategic complementarities (i.e., if ξ is close to 1), it can be seen from equation (15) that there will be an important feedback effect for firm j: the large movements in the price level imply that firm j's frictionless optimal price will react even more to aggregate demand shocks. Firm j's incentive to allocate most of its attention to this type of shocks is therefore reinforced. Hence, strategic complementarities in price setting result in strategic complementarities in attention allocation: when ξ is close to 1, it is optimal for firms to allocate more attention to the shocks for which the other firms allocate most of their attention.

Maćkowiak and Wiederholt argue that, empirically, aggregate demand shocks are much less volatile than idiosyncratic shocks (whose large size is necessary to replicate the large average size of price changes observed in the micro-price data). This means that the actual value of σ_q^2/σ_z^2 is much lower than one, suggesting that firms would have a strong incentive to allocate most of their attention to track idiosyncratic shocks. Maćkowiak and Wiederholt also recall that, in the literature on price setting, ξ is usually set between 0.8 and 0.9. This large degree of strategic complementarities in price setting would entail a large degree of strategic complementarities in attention allocation, reinforcing firms' incentive to allocate attention mostly to idiosyncratic shocks. Under their baseline calibration, Maćkowiak and Wiederholt find that firms choose to allocate 96% of their attention to idiosyncratic shocks and, consequently, only 4% to aggregate demand shocks (i.e., $\kappa_q^* = 0.04\kappa$). Since the signal precision $1/\sigma_q^2$ positively depends on the attention paid to aggregate demand shocks, and since the fluctuations in aggregate output generated by monetary shocks are all the more persistent as $1/\tilde{\sigma}_q^2$ is low, the very limited attention firms rationally choose to pay to aggregate demand shocks induces highly persistent monetary non-neutralities.

3.3 Rational inattention with multi-product firms

In Maćkowiak and Wiederholt's model, firms produce and set the price of only one good. This assumption is standard in the literature investigating the real effects of monetary shocks – and especially in DSGE models. However, the menu cost literature has recently argued (in the wake of Midrigan 2011) that this assumption would actually have important implications for monetary non-neutrality. In particular, considering multi-product firms would allow to weaken the selection effect which is inherent in state-dependent pricing models, and which explains why these models usually deliver much less monetary non-neutrality than their time-dependent counterparts. Alvarez and Lippi (2014) notably show that when the number of different goods priced by each firm tends to infinity, the selection effect vanishes and the amount of monetary non-neutrality generated by state-dependent models is as high as that generated by the popular model of Taylor (1980).

Pasten and Schoenle (2016) embed multi-product firms into Maćkowiak and Wiederholt's model. They assume that firms have to allocate their limited attention to track three types of disturbances: aggregate demand shocks, firm-specific shocks, and good-specific shocks (whose existence is documented by the authors). At first sight, it seems that a firm pricing a larger number of goods needs to spread thin its fixed amount of attention over a larger number of shocks. The attention allocated to track aggregate demand shocks is therefore mechanically reduced as the number of goods priced increases. As we have seen in the previous subsection, a lower amount of attention allocated to aggregate demand shocks turns into a lower signal precision $1/\tilde{\sigma}_q^2$, implying larger and more persistent real effects of monetary shocks. However, Pasten and Schoenle show that this spread thin effect is dominated by the existence of economies of scope in information processing. These economies of scope stem from the fact that tracking aggregate-demand and firm-specific shocks, while requiring the same attention as tracking good-specific shocks, provides information that can be used to price all goods. The benefits of tracking aggregate-demand and firm-specific shocks thus scale up with the number of different goods priced, while the benefits of tracking good-specific shocks do not. Firms accordingly find it optimal to allocate more attention to aggregate demand shocks when they price more goods. The signal precision $1/\tilde{\sigma}_q^2$ therefore increases with the number of goods priced, which implies – in sharp contrast to what happens in menu cost models – that the amount of monetary non-neutrality decreases with this number.[9]

Pasten and Schoenle find that the amount of monetary non-neutrality declines quite substantially with the number of goods priced. When firms price two goods, for instance, this amount is cut by three with respect to the single-good (i.e., Maćkowiak and Wiederholt's) case. When firms price eight goods or more, monetary shocks have almost no effect on aggregate output. This decline reflects the increasing attention allocated to aggregate demand shocks as the number of goods priced increases: while firms allocate 3% of their total attention to track aggregate demand shocks in the single-good case, this figure rises to 11% and 24% in the two-good and eight-good cases, respectively.

Finally, it is worth emphasising that strategic complementarities in price setting play an important role in Pasten and Schoenle's results. We have seen in the previous subsection that in rational inattention models, strategic complementarities in price setting induce strategic complementarities in attention allocation. These latter are particularly strong in the multi-product firms model. To illustrate this strength, Pasten and Schoenle develop an extension of their baseline model where different firms price different numbers of goods. It turns out that in this economy with heterogeneous firms, the firms which price a single good allocate much more attention to aggregate demand shocks than they do in the baseline economy (with homogenous firms). Hence, the interaction with firms allocating more attention to aggregate demand shocks leads to more attention allocated to these shocks

[9] Pasten and Schoenle nevertheless make it clear that the allocation of attention, as well as the resulting amount of monetary non-neutrality, is independent from the number of goods firms price when those goods are not subject to shocks (or when good-specific shocks have no impact on firms' profits). Hence, what ultimately matters is the number of shocks that do require paying attention to, not the number of goods per se.

by firms that would have otherwise chosen to remain largely uninformed about them.

4 Business cycle fluctuations

The recent literature involving informational imperfections has also investigated how dispersed information would impinge on the transmission of fundamental shocks and extrinsic disturbances such as sentiment shocks. An important conclusion emerging from this work is that the exogenous or endogenous nature of the private signals observed by economic agents seems to be critical in the propagation of fundamental shocks: exogenous signals tend to dampen business cycle fluctuations, while endogenous signals tend to amplify them. Furthermore, sentiment shocks generate demand-driven fluctuations through self-fulfilling variations in expected demand.

4.1 The propagation of fundamental shocks

In order to illustrate these ideas, let us consider a standard Real Business Cycle model, in which we introduce the three following assumptions (taken from Angeletos and La'O 2010).

First, capital is assumed away in order to simplify the analysis.

Second, the economy is split into a large number of islands, indexed by $j \in (0, 1)$, and a mainland. Each island is populated by a self-employed worker (hereafter farmer), who produces a differentiated, island-specific consumption good. This introduces strategic interactions across islands: the optimal production on each island will depend on other islands' production. The goods produced on the different islands are subsequently traded on the mainland.

Third, each period is divided into two stages: production takes place in stage 1, while trading and consumption take place in stage 2. In stage 1, the farmer of any given island is uncertain about the productivity and output of other islands. Information is therefore dispersed at the time when production decisions are made. In stage 2, all information that was previously dispersed becomes publicly known, and markets open. Quantities are predetermined by the production decisions made during stage 1, but prices adjust so as to clear markets.

Expressing all variables in logarithms, production on island j is given by:

$$y_{j,t} = a_{j,t} + n_{j,t}$$

where $y_{j,t}$ denotes the level of output produced in period t, $n_{j,t}$ the level of employment, and $a_{j,t}$ the level of the exogenous local productivity.

Local productivity is the sum of an aggregate and an idiosyncratic component: $a_{j,t} = a_t + \varepsilon_{j,t}$, where a_t is the aggregate technology shock, and $\varepsilon_{j,t}$ an idiosyncratic one. The former follows a random walk: $a_t = a_{t-1} + \varepsilon_{a,t}$, with $\varepsilon_{a,t}$ following a Gaussian white noise process with variance σ_a^2. The latter follows a Gaussian white noise process with variance

σ_j^2, independent from $\varepsilon_{a,t}$, and independently distributed across islands.

The local productivity of island j is perfectly observed by the farmer located on that island. However, she is unable to distinguish between the aggregate and island-specific components at the time when production decisions are made. We assume that in stage 1 of each period, the only information the farmer has about the current-period aggregate shock, a_t, is summarised in a private signal of the form:

$$s_{a,j,t} = a_t + e_{a,j,t} \tag{18}$$

with $e_{a,j,t}$ following a Gaussian white noise process with variance $\tilde{\sigma}_a^2$, independent from $\varepsilon_{a,t}$, and independently distributed across islands.

Angeletos and La'O (2010) show that the general equilibrium of the resulting model reduces to the solution of the following fixed-point equation:

$$y_{j,t} = \alpha \mathbb{E}_{i,t} y_t + (1 - \alpha)\chi a_{j,t} \tag{19}$$

where $\mathbb{E}_{i,t}$ denotes the expectations operator conditional upon the information available to island j in period t, y_t the level of aggregate output, and $\alpha < 1$ and $\chi > 0$ are reduced-form parameters that depend on the underlying preferences and technology.

The parameter α measures the degree of strategic interactions across islands. As stressed by Angeletos and Lian (2016), two opposing mechanisms determine the level of α: a demand-side effect that contributes towards strategic complementarity ($\alpha > 0$), and a supply-side effect contributing towards strategic substitutability ($\alpha < 0$). Indeed, when a farmer expects aggregate output to go up, the demand for the local good and its relative price are expected to go up as well. This effect motivates the farmer to work and produce more. This is the demand-side effect, which induces local output to increase with expectations of aggregate output. The opposing supply-side effect stems from an income effect: when a farmer expects aggregate output to go up, income is also expected to go up, which tends to discourage labor supply and production. Whether α is positive or negative therefore depends on which of the two effects dominates. In what follows, we assume that the demand-side effect dominates, so that there are strategic complementarities across islands ($\alpha > 0$). Had information been complete rather than dispersed, we would have had $\mathbb{E}_{j,t} y_t = y_t$ for all j, and, consequently:

$$y_t = \chi a_t$$

Hence, as predicted by the standard RBC model, aggregate output would have been: i) exclusively pinned down by the technology shock; ii) completely independent from the degree of strategic interactions α.

As equation (19) makes it clear, the combination of dispersed information and strategic complementarities across islands implies that individual farmers have to play beauty-contest games with the other farmers. They thus need to form higher-order expectations.

Using the same notation as in (6), Angeletos and La'O (2010) show that the aggregate level of output can be expressed as follows:

$$y_t = \chi(1-\alpha)\sum_{k=0}^{\infty}\alpha^k \mathbb{E}_{j,t}^{(k+1)} a_t \tag{20}$$

The reliance on higher-order expectations dampens the responses of aggregate output to aggregate technology shocks. Indeed, since expectations of higher order react less strongly to private signals than expectations of lower order, the fluctuations of aggregate output are milder than in the complete information case. Moreover, the strength of this dampening effect increases with the strength of strategic complementarities: the stronger strategic complementarities across islands (i.e., the closer α is to 1), the higher the order of expectations farmers have to consider, and the larger the weights they have to attach to expectations of higher orders.

This dampening effect can be more easily seen by noting that equation (20) can be rewritten as:

$$y_t = \chi a_{t-1} + \Phi\varepsilon_{a,t}$$

where Φ denotes a reduced-form parameter whose value satisfies $0 < \Phi < \chi$. Furthermore, Φ converges to 0 from above as α converges to 1 from below. Therefore, the impact of an aggregate technology shock on aggregate output, driven by Φ, is lower than its counterpart in the complete information case, which is driven by χ. In addition, this impact declines with the degree of strategic complementarities, whereas it is independent from this degree in the complete information case.

Chahrour and Gaballo (2021) argue that the existence of such a dampening effect ultimately stems from the fact that the signal received by each individual farmer is purely exogenous. They consider a more complicated RBC economy, comprising an housing sector. More specifically, households are located on different islands and consume both a traded consumption good and local housing. The consumption good is produced using labor from all islands, while local housing is produced using land, local labor, and a traded productive factor (the commodity good) whose supply is fixed. Local house prices can move either because of an island-specific shock affecting the future productivity of local labor, or because of a current aggregate shock to housing production. Given that most fluctuations in those prices are driven by local labor productivity, observing high house prices induces households to become optimistic about their own income prospects. However, an increase in local house prices can also result from a negative aggregate shock to housing production, since this kind of shock generates a rise in house prices across islands. When such a shock occurs, the observed increase in local house prices is misinterpreted by households as good news about future wages. They thus raise their demand for both consumption and housing. The resulting increase in aggregate demand further increases house prices (as well as the price of the commodity factor), reinforcing the initial price increase.

Hence, what started as a small change in housing supply leads, in equilibrium, to an economy-wide increase in house prices and to a boom in consumption and housing. Chahrour and Gaballo point out that the logic of their model extends to other sorts of fundamental shocks, and to learning from the observation of any local price. In each case, the crucial feature is a price-optimism feedback channel: higher local prices begets economic optimism, which begets even higher local prices, and so on. Economic fluctuations are therefore amplified, rather than dampened, and this amplification originates from the fact that agents receive information through the observation of signals taking the form of endogenous variables.

4.2 Sentiments as a source of disturbance

The literature investigating the role of heterogeneous information in business cycle fluctuations has not solely focused on fundamental shocks. Indeed, beauty-contest games – and the associated higher-order uncertainty – open the door to the emergence of shocks affecting the expectations of each agent regarding the expectations and actions of all other agents. These extrinsic disturbances have been commonly referred to as sentiment shocks after Angeletos and La'O (2013) and Benhabib et al. (2015). Sentiment shocks are assumed to directly affect higher-order expectations: each agent, while being well aware that no fundamental shock has occurred, suddenly believes that the other agents expect such a shock to have actually taken place. Sentiment shocks therefore take the form of waves of optimism or pessimism and, as such, can be associated with forces akin to animal spirits.

Let us consider the RBC economy introduced in the previous subsection, with two modifications.

First, we drop the idiosyncratic component of local productivity, so that this latter is only made of the aggregate component: $a_{j,t} = a_t$. Individual farmers are assumed to be unable to observe local productivity when they make their production decisions in stage 1 of each period. They can only observe the signal given by equation (18).

Second, we assume that each farmer believes that the private signals observed by the other farmers are biased: the prior of the farmer located on island j is that the noise of her signal follows a Gaussian white noise process with variance $\tilde{\sigma}_a^2$, and that the noise of the signals received by the other farmers follows a Gaussian process with mean μ_t and variance $\tilde{\sigma}_a^2$, where μ_t is a random variable that represents the perceived bias in one another's signals. We assume that μ_t follows an AR(1) process:

$$\mu_t = \rho\mu_{t-1} + \nu_t$$

where $0 \leq \rho < 1$ and ν_t follows a Gaussian white noise process with variance σ_μ^2. Innovations in μ_t are called sentiment shocks.

It can be shown that the general equilibrium of the resulting model reduces to the solution of the following fixed-point equation:

$$y_{j,t} = \alpha \mathbb{E}_{j,t} y_t + (1-\alpha)\phi \mathbb{E}_{j,t} a_t \tag{21}$$

where $\phi > 0$ is a reduced-form parameter that depends on the underlying preferences and technology. Let us now assume a positive innovation in μ_t. Such a positive sentiment shock means that, while they are observing a signal indicating no shock on aggregate productivity, individual farmers suddenly believe that the other farmers are observing wrong signals indicating a positive shock on aggregate productivity. As a result, they keep their own expectations of aggregate productivity unchanged (i.e., $\mathbb{E}_{j,t} a_t$ is not affected), but adjust their higher-order expectations upwards. This upward adjustment leads individual farmers to become optimistic about the level of aggregate output (i.e., to raise $\mathbb{E}_{j,t} y_t$). According to equation (21), the optimal response of each farmer is to raise her own level of output – since an increase in the expected level of aggregate activity means an increase in the expected demand for her product. Sentiment shocks thus generate demand-driven fluctuations through self-fulfilling variations in expected demand: a positive shock induces individual farmers to expect a higher level of aggregate demand; each farmer accordingly raises production and spending, pushing up aggregate demand and validating the initial expectation. However, given that $\rho < 1$, this wave of optimism is only short lived. Consequently, farmers find it optimal to consume only a fraction of the increase in their income, and to save the rest. This additional savings would have induced an increase in investment had capital not been assumed away.

5 Communication and stabilisation policies

Sections 3 and 4 have shown how the conjunction of strategic complementarities and heterogeneous information allows macroeconomic models to generate monetary non-neutrality and realistic business fluctuations. This framework also allows to address economic policy issues. When imperfect information causes money non-neutrality and leads to business fluctuations, disclosure of information is likely to have a significant impact on the economy. An extensive literature has discussed the welfare effects of disclosing more information to market participants. Furthermore, as information disclosure shapes the behaviour of economic agents, communication and stabilisation policies need to be assessed in a comprehensive framework.

5.1 Welfare effect of public disclosure

The simple beauty-contest model derived in section 2 shows that – due to strategic complementarities – firms overreact to public information in order to take action close to that of others. Yet, how does overreaction to public information affect welfare? This question fuelled a lively debate in the literature. It turns out that the answer depends on the analytical framework, particularly on the welfare function. It also depends on the type of fundamental about which imperfect information relates. The analytical framework developed section 2 refers to an economy where the fundamental is the aggregate nominal demand and where imperfectly known variations of it are equivalent to demand shocks. Under complete information, demand shocks don't yield inefficiencies. Indeed, to the extent that a demand shock is perfect common knowledge among agents, no welfare loss follows in this framework. Welfare losses occur when information about aggregate demand is incomplete and agents make inefficient use of their information. We summarise below the main lines of argument regarding the welfare effect of public information about demand shocks and then discuss how information can be disclosed to control the overreaction of agents.

The conclusions raised below about optimal communication policy do not hold when shocks yield inefficiencies under complete information, as is the case of supply or cost-push shocks. Because these shocks create inefficiencies even when they are perfect common knowledge, more information is typically socially costly since it strengthens firms' responses to inefficient shocks (e.g. Angeletos and Pavan (2007), Angeletos et al. (2016)).

5.1.1 Is Morris and Shin (2002) pro or contra transparency?

Morris and Shin (2002) show that, due to the overreaction mechanism, the disclosure of public information can be detrimental to welfare. In their model, the welfare loss takes the form

$$L_{MS} = \int_j (p_j - q)^2 dj = \int_j (p_j - p)^2 \, dj + (q - p)^2 \tag{22}$$

which yields, using the information structure of section 2 and the equilibrium pricing (7), the unconditional expected social loss

$$\mathbb{E}(L_{MS}) = (1-\omega)^2 \sigma_\varepsilon^2 + \omega^2 \sigma_\eta^2$$

Let us assume that the central bank, or another public agency, has the choice between disclosing the public signal with precision σ_η^2 (transparency) and withholding this signal (opacity). What is the optimal communication strategy? When the central bank withholds the public signal, the price of each firm equals its private signal (i.e., $\omega = 0$) and the corresponding expected loss is determined by the price dispersion across agents σ_ε^2. It turns out that withholding the public signal is preferable to disclosing it when the following condition hold: $2\xi - 1 > \sigma_\varepsilon^2/\sigma_\eta^2$, indicating that the disclosure of public information can be detrimental to welfare under strategic complementarities. This important result undermined the general presumption that more information generally improves the efficiency of market outcomes by showing that more information can decrease welfare even if welfare would be maximal under complete information.

However, the presumed anti-transparency result of Morris and Shin (2002) has been challenged by two main criticisms. First, Svensson (2006) observes that withholding public information improves welfare in the model of Morris and Shin only under unrealistic parameter values. Transparency is indeed detrimental to welfare when the information of the central bank on the fundamental is less accurate than the private information of firms, that is when $\sigma_\varepsilon^2 < \sigma_\eta^2$. This seems rather unrealistic because a central bank usually has a more accurate view of economic fundamentals than any individual firm, as shown by Romer and Romer (2000). Svensson concludes that the model of Morris and Shin speaks actually in favour of rather than against transparency.

Second, Woodford (2005) and Hellwig (2005) show that the anti-transparency result of Morris and Shin rests on the specificity of their utility and welfare functions, and does not hold in a micro-founded framework. Because utility of the representative household decreases significantly with the price dispersion across goods in these models, public information improves household's utility more through its coordinating role on prices than it worsens utility through its distortive effect on the average price. In micro-founded models, the second-order approximation of the loss of the representative household increases with the dispersion of prices across firms and the distortion of the price level from demand:

$$L_{MF} = \frac{\theta}{1-\xi} \int_j (p_j - p)^2 \, dj + (q-p)^2 \tag{23}$$

where $\theta > 1$ is the parameter of price elasticity of demand in the Dixit-Stiglitz composite good. This micro-founded loss function assigned a much higher weight to price dispersion than the loss function (22) of Morris and Shin. As a result, the unconditional expected loss with the public signal, $\mathbb{E}(L_{MF}) = \theta(1-\omega)^2\sigma_\varepsilon^2/(1-\xi) + \omega^2\sigma_\eta^2$, is always smaller than the expected loss without the signal, $\theta\sigma_\varepsilon^2/(1-\xi)$, regardless of the variance of the error terms. Transparency improves thus welfare in micro-founded models.

Angeletos and Pavan (2007) generalise the analysis and show that the welfare effect of public information depends on the relation between the *equilibrium* degree of coordination in firms' action and the *efficient* degree of coordination from a social welfare perspective. The equilibrium degree of coordination corresponds to the degree of strategic complementarities which drives firms' coordination in equilibrium, i.e. ξ in (1). The efficient degree of coordination is the degree of coordination of firms that would maximise welfare. When the equilibrium degree of coordination coincides with the efficient degree, transparency improves welfare because firms make an efficient use in equilibrium of their private and public information. By contrast, when the equilibrium degree of coordination exceeds the efficient degree, firms react too strongly to public information compared to what is socially desirable and the disclosure of public information may be detrimental to welfare. This is the case in the model of Morris and Shin: differentiating the loss (22) with respect to ξ shows that the efficient degree of coordination is 0, which is smaller than the equilibrium degree ξ. Inversely, when the equilibrium degree of coordination is lower than the efficient degree, firms react too weakly to public information compared to what is socially desirable and the disclosure of public information improves coordination and welfare. This is the case in the micro-founded model: differentiating the loss (23) with respect to ξ shows that the efficient degree of coordination is $1 - (1-\xi)/\theta$, which is larger than the equilibrium degree ξ.

The welfare effect of public information depends thus on the gap between the equilibrium degree of coordination of firms and its efficient degree. The central bank has no influence on the equilibrium degree of coordination of firms and therefore cannot reduce its gap to the efficient degree to maximise welfare. But it can adjust its communication policy to control the overreaction to public information and reduce the welfare loss.

5.1.2 Reducing overreaction

How can the central bank reduce overreaction to its disclosures? The central bank can either disclose public information to a fraction of market participants only (partial publicity) or disclose public information with some ambiguity to all market participants (partial transparency). Both strategies are equivalent in reducing overreaction and improving welfare.

In the partial-publicity strategy, proposed by Cornand and Heinemann (2008), each firm j observes a private signal q_j and only a fraction of firms observes a public signal D. As previously, these signals deviate from the fundamental q by some error terms normally distributed. The private signal is defined as in (2). However, only the fraction Q of firms observes the common signal D defined as in (3). This signal is thus semi-public. Q is the degree of publicity.

Choosing a communication channel which does not reach all market participants reduces overreaction to the disclosure because uninformed participants cannot react to it, whereas informed participants react less strongly because they know that some of their peers are uninformed. The equilibrium action of informed firms who get both the private and the semi-public signal is a linear combination of their signals and is given by

$$p_{j,Q} = \frac{(1-\xi Q)\,\sigma_\eta^2}{\sigma_\varepsilon^2 + (1-\xi Q)\,\sigma_\eta^2} q_j + \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + (1-\xi Q)\,\sigma_\eta^2} D \tag{24}$$

If $Q = 1$, all firms observe the public signal D and the equilibrium price (24) is equivalent to (7) in the previous section. However, when the fraction of informed firms falls, informed firms assigned a smaller weight on the semi-public signal D because they anticipate in their higher-order expectations that the firms which do not observe the semi-public signal set their price equal to their private signal q_j.

What is the optimal share of firms Q to which the semi-public signal should be disclosed? The loss functions (22) and (23) discussed above can be rewritten as

$$L = \lambda \int_j (p_j - p)^2 \, dj + (q - p)^2 \tag{25}$$

with $\lambda = 1$ and $\lambda = \theta/(1-\xi)$, respectively. Differentiating the unconditional expected loss with respect to Q gives the optimal degree of publicity that maximises welfare: $Q^* = \min\left[\max\left[0, \frac{\lambda(\sigma_\varepsilon^2+\sigma_\eta^2)}{(2-2\lambda+3\xi\lambda)\sigma_\eta^2}\right], 1\right]$. Interestingly, when $\lambda = 1$ as in Morris and Shin (2002), Q^* can be lower than 1 even if the central bank has more accurate information on the fundamental than firms, that is even if $\sigma_\eta^2 < \sigma_\varepsilon^2$. This means that the critique of Svensson is not robust to the introduction of partial publicity as reducing publicity may improve welfare even if the central bank has more accurate information than firms.

The second strategy for reducing overreaction consists of disclosing public information to all market participants, but with ambiguity.[10] The degree of transparency is determined by the idiosyncratic noise of the public signal disclosed to all market participants. Each firm j observes a private signal q_j and a semi-public signal with an idiosyncratic noise D_j. As previously, these signals deviate from the fundamental q by some error terms normally distributed. The private signal is defined as in (2). The semi-public signal is defined as $D_j = D + \phi_j$ with D defined as in (3) and $\phi_j \sim N(0, \sigma_\phi^2)$. The signal D_j is semi-public as it contains an error term η that is common to all firms and an error term ϕ_j that is private to each firm. The idiosyncratic error term captures the ambiguity of the public disclosure and the fact that each firm may interpret it differently. Central bankers are indeed known to speak with ambiguity, making their speeches equivocal and subject to interpretation.[11] With this information setup, the equilibrium price of firm j is

$$p_j = \frac{(1-\xi)\sigma_\eta^2 + \sigma_\phi^2}{\sigma_\varepsilon^2 + (1-\xi)\sigma_\eta^2 + \sigma_\phi^2} q_j + \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + (1-\xi)\sigma_\eta^2 + \sigma_\phi^2} D_j$$

Communicating with ambiguity reduces overreaction to public information because ambiguity mitigates its focal role as it entails uncertainty about how other market participants interpret the disclosure. If $\sigma_\phi^2 = 0$, the central bank is perfectly transparent, all firms interpret its signal identically, and the equilibrium price is equivalent to (7) in the previous section.

Differentiating the unconditional expected loss (25) with respect to σ_ϕ^2 gives the optimal degree of transparency that maximises welfare: $\sigma_\phi^{2*} = \max\left[0, \frac{2-3\lambda+3\xi\lambda}{\lambda}\sigma_\eta^2 - \sigma_\varepsilon^2\right]$. Similarly to the strategy of partial publicity, the critique of Svensson is not robust to the introduction of partial transparency as σ_ϕ^{2*} can be higher than 0 when $\lambda = 1$ even if the central bank has more accurate information than firms.

Baeriswyl and Cornand (2014) show that partial publicity and partial transparency are equivalent in reducing overreaction to public information and in improving welfare. Moreover, a given equilibrium weight in price setting yields the same dispersion of prices across firms $\int_j (p_j - p)^2 dj$ and the same distortion of the average price from the fundamental $(y-p)^2$ in both partial-publicity and partial-transparency strategies. They also run a laboratory experiment to test this equivalence result and show that partial publicity and partial transparency succeed in reducing overreaction to public information, though not as much as theory predicts.

[10] Heinemann and Illing (2002) are the first to introduce this concept in a speculative attack game.

[11] To illustrate the ambiguity of central bankers' speeches, let us recall the bon-mot of Alan Greenspan in front of the American Congress: "Since I have become a central banker, I have learned to mumble with great incoherence. If I seem unduly clear to you, you must have misunderstood what I said" (1987). More recently (in December 2002), Mike Moskow, president of the Chicago Fed commenting about the way the Fed communicates says that it is: "a language in which it is possible to speak, without ever saying anything".

5.2 Endogenous information

An important extension of the analytical framework described above is to relax the assumption that agents' information is exogenously given, to make the acquisition of information or its content dependent on the actions of agents. In a nutshell, strategic complementarities in actions carry over to endogenous information, which strengthens the effect of complementarities on welfare. We review some of these extensions and their implications for welfare.

Learning from prices As already underlined by Hayek (1945), market prices are not only exchange ratios between goods but also aggregators of information: prices play an informational role by aggregating agents' knowledge and belief about the state of the economy. The literature has analysed firms' learning from market prices on the one hand and central bank's learning from market outcomes on the other. Relying on an island model à la Lucas (1972), Amador and Weill (2010) study the welfare effects of public information disclosure about productivity and/or monetary shocks in a microfounded model when firms learn from prices. The economy is subject to an aggregate monetary shock, an aggregate productivity shock and idiosyncratic real demand shocks. Economic agents cannot tell whether a high nominal price in their own sector is due to a high real demand shock, or to a high monetary shock. However, they learn (partially) about the monetary shock by observing the price level. The informativeness of the latter is endogenous: agents use the price level as a signal to improve their forecast of the monetary shock. Because the public forecast is the same for everyone, all there is to learn about monetary shocks comes from other agents' private forecasts. The informativeness of the price level is determined by the weight agents assign to their private forecasts. If agents put more weight on their private forecasts, inducing that more private information feeds into prices, then the price level becomes more informative, improving in turn the quality of agents' forecasts. In this context, releasing more precise public information about the monetary shock, the productivity shock, or both will induce agents to put more weight on these signals and thus less on private information, reducing the endogenous informational content of prices. This negative effect is amplified by strategic complementarities, which increase in equilibrium agents' uncertainty about the monetary shock and their real demands, and reduces welfare.

While learning from prices is exerted only via a public channel in Amador and Weill (2010), Amador and Weill (2012) envisage the possibility that agents also learn via a private channel. They consider that dispersed private knowledge about own local markets (that also informs agents about aggregate economic conditions) diffuses in the population over time as households and firms learn from each other valuable bits of information that are not yet known to everyone. While the latter dissemination operates via a private channel, on the opposite, the information aggregated in prices or in the macroeconomic figures published by agencies are typically public. Amador and Weill show that as long as some of the learning is due to a private channel, public information may increase the overall level of information in the short run, but this comes at the cost of lowering the overall level of information in the long run. Therefore, when agents also learn from private channels, a release of public information can reduce welfare.

In both cases, by inducing private agents to rely less on their private information, the disclosure of public information reduces the informational efficiency of the price system, giving rise to a detrimental welfare effect.

Another branch of the literature highlights that policy makers themselves – and not only private agents – do rely on market outcomes to make their decision: central banks themselves assess the state of the economy by observing market outcomes, such as prices or the level of economic activity. Rather than directly observing fundamental shocks, the central bank collects information about economic fundamentals from observing the economic outcome. Indeed, economic shocks cannot be observed independently of the behavior of economic agents. At the same time, the central bank shapes the economy through the conduct of monetary policy. This gives a dual role to the central bank as both observer and shaper of the economy, posing a dilemma documented by Amato and Shin (2006) and Morris and Shin (2005): the more a central bank influences economic outcomes, the less reliable those outcomes are as indicators of the fundamental of the economy. Baeriswyl (2011) shows that taking this mirror effect into account worsens the potential detrimental effect of overreaction to public information highlighted by Morris and Shin (2002) and calls for a lower degree of central bank transparency than in the case where central bank's information is exogenously given.

Information acquisition Whereas the above papers focus on the endogeneity of information content, there is a literature that considers firms' strategic information acquisition choice. The main idea developed by Hellwig and Veldkamp (2009) and Myatt and Wallace (2012) is that when firms face the choice of acquiring different information sources (at some cost), they give more attention to the most public (and least private) signals. Indeed, strategic complementarities in actions give rise to strategic complementarities in information acquisition. The more agents' actions are complementary, the more the information acquired becomes public. Intuitively, public signals act as effective focal points for players' coordination. As the desire for coordination strengthens, agents pay more attention to such signals.

These theoretical predictions do not seem to be fully borne out in reality. For example, while theory predicts forward guidance by central banks about their intended future policy to have a significant impact on financial markets, their reaction seems surprisingly weak. Shaping market expectations seems more challenging in practice than in theory. One branch of the literature has set out to explain the possible causes of this so-called forward guidance puzzle, putting forward, like Angeletos and Lian (2018), the lack of common knowledge among market participants, or, like Garcia-Schmidt and Woodford (2019), the limited level of reasoning as possible causes of the agents' weak reaction. Both hypotheses are supported in the laboratory experiment by Baeriswyl et al. (2021). This study shows that participants in the lab pay less attention to public disclosure than theoretically predicted, which prevents the emergence of common knowledge, and operate only a limited level of reasoning in choosing their action, both contributing to a weaker effect of public disclosure.

5.3 Stabilisation policy

The models presented in the previous sections examined the effect of information disclosure on agents' actions and welfare, but ignored the fact that information disclosure by a policy maker goes often hand in hand with the implementation of a policy action. In reality, central banks play a direct role in shaping macroeconomic outcomes through active policy interventions. James and Lawler (2011) make an important contribution by analysing how the inclusion of policy action affects the welfare effect of information disclosure.

Suppose that the central bank stabilises the economy affected by demand shocks q owing to its monetary instrument I. The nominal aggregate demand is composed of the demand shock and the monetary instrument. Firms set their price as a function of their first and higher order beliefs on the demand shock and the monetary instrument:

$$p_j = \mathbb{E}_j \left[(1-\xi)(q+I) + \xi p\right] \tag{26}$$

The central bank sets its monetary instrument as a function of its information on the demand shock, $I = \rho \cdot D$. This policy function is common knowledge among firms.

As in the set-up above, each firm j receives a private signal as defined in (2). The central bank also discloses its own information about the fundamental with some degree of transparency, D_j, as defined in section 5.1.2.

The equilibrium action of agent j yields

$$p_j = \frac{(1-\xi)\sigma_\eta^2 + (1+\rho)\sigma_\phi^2}{\sigma_\varepsilon^2 + (1-\xi)\sigma_\eta^2 + \sigma_\phi^2} q_j + \frac{(1+\rho)\sigma_\varepsilon^2 + \rho(1-\xi)\sigma_\eta^2}{\sigma_\varepsilon^2 + (1-\xi)\sigma_\eta^2 + \sigma_\phi^2} D_j = \omega_q \cdot q_j + \omega_D \cdot D_j \tag{27}$$

with $\omega_q + \omega_D = 1 + \rho$. Given the equilibrium action of agents, the policy action coefficient ρ that maximises the loss function

$$L = \lambda \int_j (p_j - p)^2 \, dj + (q + I - p)^2 = \lambda(\omega_q^2 \sigma_\varepsilon^2 + \omega_D^2 \sigma_\eta^2) + (\rho - \omega_D)^2 \sigma_\eta^2 \tag{28}$$

is given by

$$\rho^* = -\frac{(\lambda\sigma_\varepsilon^2 + (2\lambda - 2\lambda\xi - 1)\sigma_\eta^2 + \lambda\sigma_\phi^2)\sigma_\varepsilon^2}{(\lambda\sigma_\varepsilon^2 + \sigma_\eta^2)\sigma_\phi^2 + \lambda(\sigma_\varepsilon^2 + (1-\xi)\sigma_\eta^2)^2} \tag{29}$$

The central bank counteracts the demand shock with its instrument and its response increases in absolute value with the accuracy of its information ($\partial\rho/\partial\sigma_\eta^2 > 0$) to reach -1 when its information is perfect. Firms, of course, anticipate the stabilisation policy of the central bank when setting their price.

How should then the central bank communicate when it implements an optimal instrument? Plugging the equilibrium action (27) and optimal instrument (29) into the loss function (28) and differentiating it with respect to σ_ϕ^2 shows that unconditional expected loss decreases with the opacity of the central bank: $\partial \mathbb{E}(L|_{\rho^*})/\partial \sigma_\phi^2 < 0$. In other words, the central bank that stabilises the economy with its monetary instrument should remain completely opaque about its information on the demand shock. Moreover, not only is zero disclosure desirable as part of an optimal stabilisation policy, but it also achieves the first best possible outcome, i.e. the outcome that would be achieved if firms made efficient use of their information.

The analysis by Morris and Shin (2002) points out that strategic complementarities lead firms to make socially inefficient use of their information, as public information is more predictive of the behaviour of others than private information. In the James and Lawler (2011) framework where the central bank stabilises the economy with its instrument, the central bank uses its information efficiently as it aims to minimise the social loss rather than optimise the beauty contest of firms. The central bank's information is thus put to better use by implementing the monetary instrument than by disclosing it to firms that would overreact to it. Opacity is therefore beneficial when the central bank implements a policy action: a higher welfare can be reached when the central bank stabilises shocks without communicating since it avoids firms' overreaction while ensuring all information is exploited optimally. Nevertheless, the question arises as to whether it is realistic to assume that the central bank's action can remain unknown to market participants. Baeriswyl and Cornand (2010) or Melosi (2017) state that the implementation of monetary policy is necessarily observed by market participants and, therefore, signals central bank's information. The central bank must therefore take this signaling channel into account when making its monetary policy decision.

As in the case of pure communication policy, the analytical framework with stabilisation policy has been extended to endogenous information structures.[12] Paciello and Wiederholt (2014) study the combined effect of communication and stabilisation policies when the amount of attention that firms devote to aggregate conditions is endogenous. There are two types of aggregate shocks: shocks that cause efficient fluctuations under perfect information (as we have assumed so far) and shocks that cause inefficient fluctuations under perfect information (e.g., markup shocks). Owing to its policy, the central bank can discourage firms to pay attention to markup shocks. The optimal monetary policy in response to the first type of shocks is to fully stabilise the profit-maximising price, implying that firms do not have to change prices in response to the shock. In response to the second type of shock, the central bank would like to avoid the allocation under perfect information. The optimal monetary policy is therefore to reduce inefficient price dispersion by stabilising the profit-maximising price in response to markup shocks, such that firms pay no attention to the latter. In addition, the monetary policy that stabilises profit-maximising price in response to markup shocks has a direct negative effect on con-

[12] See e.g. Morris and Shin (2018) and Baeriswyl et al. (2020) for the study of the optimal stabilisation and communication policies when the central bank observes the market outcome.

sumption, but this effect is weakened by the fact that price setters pay less attention to markup shocks since this implies that the price level increases less after a positive markup shock. Therefore, a monetary policy that reduces the response of profit-maximising price to markup shocks reduces both inefficient price dispersion and consumption variance.

References

Adam, K. (2007). Optimal monetary policy with imperfect common knowledge. *Journal of Monetary Economics*, 54, 267-301.

Alvarez, F., Lippi, F. (2014). Price setting with menu cost for multiproduct firms. *Econometrica*, 82(1), 89-135.

Angeletos, G.-M., Iovino, L., La'O, J. (2016). Real rigidity, nominal rigidity, and the social value of information. *American Economic Review*, 106(1), 200-227.

Angeletos, G.-M., La'O, J. (2010). Noisy Business Cycles. *NBER Macroeconomics Annual 2009*, 24, 319-378.

Angeletos, G.-M., La'O, J. (2013). Sentiments. *Econometrica*, 81(2), 739-779.

Angeletos, G.-M., Lian, C. (2016). Incomplete Information in Macroeconomics. *Handbook of Macroeconomics*, Vol.2, Chapter 14, 1065-1240.

Angeletos, G.-M., Lian, C. (2018). Forward Guidance without Common Knowledge. *American Economic Review*, 108, 2477-2512.

Angeletos, G.-M., Pavan, A. (2007). Efficient Use of Information and Social Value of Information. *Econometrica*, 75, 1103-1142.

Amador, M., Weill, P.-O. (2010). Learning from Prices: Public Communication and Welfare. *Journal of Political Economy*, 118, 866-907.

Amador, M., Weill, P.-O. (2012). Learning from private and public observations of others' actions. *Journal of Economic Theory*, 147, 910-940.

Amato, J.D., Shin, H.S. (2006). Imperfect common knowledge and the information value of prices. *Economic Theory*, 27, 213–241.

Baeriswyl, R., Cornand, C., 2010. The Signaling Role of Policy Actions. *Journal of Monetary Economics*, 57, 682-695.

Baeriswyl, R. (2011). Endogenous Central Bank Information and the Optimal Degree of Transparency. *International Journal of Central Banking*, 7(2), 85-111.

Baeriswyl, R., Cornand, C. (2014). Reducing overreaction to central banks disclosure: theory and experiment. *Journal of the European Economic Association*, 12, 1087-1126.

Baeriswyl, R., Cornand, C., Ziliotto, B. (2020). Observing and shaping the market: the dilemma of central banks. *Journal of Money, Credit and Banking*, 52, 1973-2005.

Baeriswyl, R., Boun My, K., Cornand, C. (2021). Double overreaction in beauty-contests with information acquisition: theory and experiment. *Journal of Monetary Economics*, 118(C): 432-445.

Benhabib, J., Wang, P., Wen, Y. (2015). Sentiments and Aggregate Demand Fluctuations. *Econometrica*, 83(2), 549-585.

Chahrour, R., Gaballo, G. (2021). Learning from House Prices: Amplification and Business Fluctuations. *The Review of Economic Studies*, 88(4), 1720-1759.

Christiano, L. J., Eichenbaum, M., Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, 113(1), 1-45.

Clerc, P., Dos Santos Ferreira, R. (2022). Dispersed information and the non-neutrality of money: fifty years after Lucas 1972. *Journal of Economic Methodology*, 29(1), 86-104.

Cornand, C., Heinemann, F. (2008). Optimal Degree of Public Information Dissemination. *Economic Journal*, 118, 718-742.

Garcia-Schmidt, M., Woodford, M. (2019). Are low interest rates deflationary? A paradox of perfect-foresight analysis. *American Economic Review*, 109(1), 86-120.

Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review*, 35(4), 519-530.

Heinemann F., Illing G. (2002). Speculative Attacks: Unique Equilibrium and Transparency. *Journal of International Economics*, 58, 429-450.

Hellwig, C. (2005). Heterogeneous information and the welfare effects of public information disclosures. *Mimeo.*

Hellwig, C., Veldkamp, L. (2009). Knowing what others know: Coordination motives in information acquisition. *Review of Economic Studies*, 76, 223–251.

James, J.G., Lawler, P. (2011). Optimal Policy Intervention and the Social Value of Public Information. *American Economic Review*, 101, 1561–1574.

Keynes, J.M. (1936). *The General Theory of Employment, Interest, and Money*, Mc Millan, Cambridge University Press, 1973. Chap. XIII : "The State of Long-Term Expectation".

Leahy, J. (2011). A survey of New Keynesian theories of aggregate supply and their relation to Industrial Organization. *Journal of Money, Credit and Banking*, 43(Supplement), 87-110.

Lucas Jr, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, 4, 103-124.

Maćkowiak, B., Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3), 769-803.

Melosi, L., 2017. Signaling effect of monetary policy. *The Review of Economic Studies*, 84, 853–884.

Midrigan, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, 79(4), 1139-1180.

Morris, S., Shin, H.S. (2002). Social Value of Public Information. *American Economic Review*, 92, 1521-1534.

Morris, S., Shin, H.S. (2005). Central Bank Transparency and the Signal Value of Prices. *Brookings Papers on Economic Activity*, 2, 1-66.

Morris, S., Shin, H.S. (2018). Central Bank Forward Guidance and the Signal Value of Market Prices. *AEA Papers and Proceedings*, 108, 572-577.

Myatt D.P., Wallace, C. (2012). Endogenous information acquisition in coordination games. *Review of Economic Studies*, 79, 340-374.

Paciello, L., Wiederholt, M. (2014). Exogenous Information, Endogenous Information, and Optimal Monetary Policy. *Review of Economic Studies*, 81, 356-388.

Pasten, E., Schoenle, R. (2016). Rational inattention, multi-product firms and the neutrality of money. *Journal of Monetary Economics*, 80(1), 1-16.

Phelps, E. S. (1983). The trouble with Rational Expectations and the problem of inflation stabilization. In: Frydman, R., Phelps, E. S. (eds.), *Individual Forecasting and Aggregate Outcomes: 'Rational Expectations' Examined*, Cambridge: Cambridge University Press.

Romer, C.D., Romer, D.H. (2000). Federal reserve information and the behavior of interest rates. *American Economic Review*, 90 (3), 429-457.

Sims, C. A. (1998). Stickiness. *Carnegie-Rochester Conference Series on Public Policy*, 49, 317-356.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50, 665-690.

Svensson, L.E.O. (2006). Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con. *American Economic Review*, 96, 448-452.

Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. *Journal of Political Economy*, 88(1), 1-23.

Townsend, R. M. (1983). Forecasting the forecasts of others. *Journal of Political Economy*, 91(4), 546-588.

Woodford, M. (2003a). Imperfect common knowledge and the effects of monetary policy. In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (eds.), *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps.* Princeton University Press, Princeton.

Woodford, M. (2003b). *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton University Press, Princeton.

Woodford, M. (2005). Central Bank Communication and Policy Effectiveness. Working Paper No. 11898, *NBER Working Paper Series.*