

Optimal Transparency and Policy Intervention with Heterogeneous and ‘Sticky’ Information

Jonathan G. James

Swansea University

and

Phillip Lawler*

Swansea University

June 2013

Abstract:

The paper examines the optimal degree of monetary policy transparency in the context of a model of price-setting firms. The latter comprise two distinct groups: first, those whose pricing decisions are based on current information, which takes the form of observations of both a public signal, whose source is the central bank, and a private (i.e. firm-specific) signal; second, those which do not update their information before setting prices. The non-adjustment of prices by firms that fail to revise their information provides a source of price dispersion additional to that arising from heterogeneity of the information possessed by firms which do update. The objective of monetary policy is to maximize expected welfare. The central bank is assumed to adjust its instrument according to a pre-determined rule which defines the setting of the instrument as a linear function of the central bank's current information, consisting of both the actual realization of its private signal and the publicly announced value of the latter. In addition to choosing its instrument setting, the central bank must also decide on how precisely to reveal its own information to the private sector. The focus of the paper is then on the issue of the optimal degree of transparency.

Keywords: strategic complementarity; public disclosure; policy intervention

JEL Classification: C72; D62; D82; E58.

* Corresponding Author. Department of Economics, Swansea University, Singleton Park, Swansea SA2 8PP, United Kingdom.

Telephone: +44-1792-295-168. Fax: +44-1792-295-872. E-mail: p.lawler@swansea.ac.uk

1. Introduction

A major issue in monetary policy design is that of whether central banks should publicly reveal the information on which their policy decisions are based. As widely documented (see for example, Dincer and Eichengreen, 2007; Blinder et al. 2008), modern central bank practice reflects, to a large degree, a belief that transparency is conducive to the attainment of the aims of policy. Nonetheless, theoretical support for the alleged efficacy of central bank disclosure is far from categorical.

A particular qualification to the case for transparency is identified in the literature concerned with the consequences of imperfect heterogeneous information regarding economic fundamentals in economies in which strategic complementarities are present. This literature, initiated by Morris and Shin's (2002) important contribution, has shown that in such an environment individual incentives to coordinate actions with those of other agents may either undervalue or overstate the social benefits from coordination. The associated departure from efficiency which potentially arises in the context of dispersed information then gives rise to the possibility that improvements in private sector information will be detrimental to welfare. In such an instance, any public communication from the central bank regarding its assessment of the prospective economic state can then be counter-productive.

Within the related literature, two critical factors have been identified as determining the welfare effects of central bank transparency. First, whether or not there is a divergence between what Angeletos and Pavan (2007a) define as 'the equilibrium degree of coordination' and 'the socially optimal degree of coordination'. Such a divergence underpins the potential, referred to above, for inefficiency in private sector responses to available information. The second consideration is then whether the central bank uses its private information, that might otherwise be made public, to undertake active monetary policy intervention. Optimal stabilization policy is able to ensure that the dispersed information of the private sector and the central bank's own information are together exploited efficiently: however, its ability to accomplish this depends crucially on the latter not being publicly disclosed.¹ These principles

¹ We note that central bank disclosure might be detrimental even in the absence of stabilization policy if the equilibrium degree of coordination exceeds the socially optimal degree of coordination: Morris and Shin's (2002) 'beauty contest' model provides an example of such an eventuality, though as argued by Svensson (2006) their 'anti-transparency' conclusion depends crucially on parameter values. On the other hand, in the presence of policy intervention, if such policy is not directed at maximizing a measure of social welfare which is consistent

are established in James and Lawler (2011, 2012a) in the context of abstract representations of an economy, and are reflected in the findings of Baeriswyl and Cornand's (2010) analysis of transparency using a micro-founded macroeconomic model.

The present paper examines the robustness of these principles to the presence of some private sector agents who do not base their economic choices on continuously updated information. Specifically, applying the simple general equilibrium framework due to Woodford (2002, 2003) and Adam (2007), in which firms' pricing decisions are made in an uncertain macroeconomic environment, only a subset of these firms adjust their prices in response to current signals regarding the economic state. The remaining firms are then assumed to set prices at values which maximize the expected value of profits, conditional on information available in the previous period. The approach adopted can be rationalized in terms of either Mankiw and Reis's (2002) concept of 'sticky information' or Carroll's (2003) epidemiological model of expectations, which both assume that (possibly due to information acquisition and processing costs) the updating of information sets is not immediate across all agents.

Firms which update their information in the current period observe a firm-specific private signal of the realized value of an aggregate demand shock. They may also, depending on the transparency regime, observe a public signal communicated by the central bank, and derived from its own private signal, of the shock. For simplicity, the central bank is assumed to be the only source of public information;² the degree of transparency is then represented by the information content of the signal transmitted by the central bank. In addition to potentially being communicated (possibly in a modified form) to the private sector, the central bank's private signal informs its setting of monetary policy.

The presence of firms which do not update their information prior to setting prices gives rise to a source of heterogeneity in information additional to that associated with the idiosyncratic private signals observed by firms which do update. In this respect, the present analysis is connected to that contained in James and Lawler (2012b), which assumes differences in

with the payoff functions of individual agents then it will generally be associated with the possibility that greater transparency is welfare-improving: see, for example, Walsh (2007).

² The presence of further sources of public information would not affect the conclusions drawn from the analysis that follows, providing that the errors in the central bank's signal contained a component that was orthogonal to the errors in any additional public signals.

information quality across two groups of private sector agents in the setting of Morris and Shin's (2002) framework.³ In the context of the current model, heterogeneity of information is reflected in dispersion of prices: such dispersion represents a source of welfare loss additional to that arising from aggregate output volatility. The latter, like price dispersion, is amplified by the presence of firms that do not set prices on the basis of current information.

The study's focus is on the following question: given that, conditional on the central bank's private signal, monetary policy is conducted optimally from the perspective of social welfare, how does the release of some or all of the central banks information to the private sector impact on welfare outcomes? If all firms adjusted prices in response to current signals, social welfare would, in accordance with the analyses of Baeriswyl and Cornand (2010) and James and Lawler (2011, 2012a), be strictly declining in the degree of transparency. However, in the presence of firms whose prices do not respond to current information, we find that there is no definitive answer to the question: indeed, it is possible for zero, full or some intermediate degree of transparency to be optimal. Nonetheless, reference to empirical evidence regarding the values of key parameters suggests that full transparency is likely to be dominated, in a welfare sense, by zero transparency. Whether or not the latter is to be preferred to partial transparency is shown to depend on the proportion of firms whose pricing decisions do not incorporate current information.

The remainder of the paper is structured as follows. Section 2 outlines the model structure. Our main results are identified and discussed in Section 3, which determines the properties of the model's equilibrium and analyses the welfare effects of transparency. In Section 4, the implications of a corrective tax scheme, as originally considered by Angeletos and Pavan (2007b, 2009), are briefly explored. Finally, Section 5 summarizes and presents the conclusions to be drawn from the study.

2. The model

³ In James and Lawler (2012b), the private sector is assumed to comprise two sets of agents, of equal size, distinguished by the precision of the private signals that they observe, though with all agents potentially (that is, depending on the disclosure regime) observing a common public signal.

The model that provides the vehicle for the analysis of this paper derives from the micro-founded general equilibrium framework developed in the studies of Woodford (2002, 2003) and Adam (2007), and applied in the context of the transparency issue by Baeriswyl and Cornand (2010), Hahn (2010) and Roca (2010). There is a continuum of monopolistically competitive firms, uniformly distributed over the unit interval with each employing a common production technology. The representative household's utility is defined over consumption, described by a Dixit-Stiglitz aggregator over the varieties of differentiated goods, and leisure. The framework allows household welfare to be represented, as a second-order approximation, by the following:⁴

$$W = -\left[\int_0^1 (p_i - p)^2 di + \lambda y^2\right] \quad (1)$$

where p_i is the price set by firm i , $p(\equiv \int_0^1 p_i di)$ identifies the average price across all firms, i.e. the price level, and y represents the output gap, i.e. the deviation of actual output from its 'natural' or full information level.⁵ Realized welfare is therefore determined both by the degree of price dispersion across firms and by the magnitude of the output gap. Within the model, heterogeneity of product prices is relevant to welfare because of its consequences for the dispersion of output levels across firms. In this context, changes in the price level impact on welfare only insofar as they affect the extent of price (and therefore output) dispersion. The relative importance of the two components of the right-hand-side of (1), is determined by the structural parameter λ , whose value depends on the degree of risk aversion of the household as well as the elasticity of substitution, $\theta(> 1)$, between different varieties of good.

Firm i 's optimal price, \hat{p}_i , is found from consideration of its profit-maximization problem to be a linear function of its expectations of the price level and the output gap:

⁴ For a derivation of this approximation see, for example, Woodford (2002) and Adam (2007).

⁵ Because of the type of shock considered in this paper, i.e. to aggregate demand, the full information output level is constant. The latter's value is, for convenience, normalized at zero and so the terms 'output gap' and 'output' can be used interchangeably. Aggregate demand shocks provide an example of a broader class of shocks, including those to technology and to household preferences, that do not affect the relationship between full information equilibria and the corresponding socially optimal outcomes. Shocks to mark-ups in the goods or labour markets do impact on this relationship and introduce additional considerations which tend to reinforce the broad tenor of the results identified in this paper.

$$\hat{p}_i = E_i(p + \beta y) \quad (2)$$

where the parameter β , which determines the responsiveness of \hat{p}_i to the expected output gap, is related to λ by $\beta = \lambda\theta$, and thus $\lambda < \beta$.

Aggregate nominal demand, n , within the economy is determined by the setting of the central bank's policy instrument, g , and the realization of an aggregate demand shock, ϕ , i.e. $n (\equiv p + y) = g + \phi$. In conducting the analysis, we assume realizations of the shock to be drawn, as in Morris and Shin (2002), from a uniform prior over the real line. Although this appears a somewhat special case and, as will be discussed, has certain implications for the optimal conduct of monetary policy, all of the principal results derived in what follows extend to the case of a normally distributed disturbance.

The above expression for nominal income can be rearranged as:

$$y = g - p + \phi \quad (3)$$

allowing (2) to be expressed as:

$$\hat{p}_i = E_i[(1 - \beta)p + \beta(g + \phi)] \quad (2')$$

Equation (2') makes clear the dependence of firm i 's optimal price on the prices set by all other firms. We consider the case of $\beta \in (0,1)$, in which case prices are strategic complements, i.e. there is a benefit (declining in β) to the individual firm of aligning its own price with those of other firms.

The significance of a strategic complementarity in price setting derives from the presence in the model, discussed below, of dispersed private sector information. These aspects of the framework connect the current study to the literature initiated by Morris and Shin's (2002) identification of the potential dangers of public disclosure of information by policymakers. Important in this context is Angeletos and Pavan's (2007a) distinction between the equilibrium degree of coordination and the socially optimal degree of coordination. The former identifies the equilibrium response of agents' actions to the expected aggregate value of such actions, while the latter indicates the socially efficient response. Should these coordination concepts diverge in value then, in the presence of information that is

heterogeneous across agents, the private incentives to align actions do not appropriately reflect the social benefits of such alignment.

Morris and Shin's (2002) important contribution employs a model in which the equilibrium degree of coordination exceeds the socially optimal degree of coordination. The key implication of this is that public information is accorded an inefficiently high weight relative to private information and provides the source of the conclusion that greater precision of public information might be damaging to welfare. The present framework, in contrast, is characterized by an equilibrium degree of coordination which lies below the socially optimal degree of coordination.⁶ Reflecting this, the individual firm does not fully recognize the social benefit of aligning its own price with other firms' prices, with the consequence that price setting entails an excessive weight being placed on idiosyncratic private information.

If each firm's individual product price was continuously adjusted to its (ex ante) optimal value, \hat{p}_i , in response to current information, price (and output) dispersion would purely reflect the idiosyncratic information errors associated with heterogeneous firm-specific signals. However, we assume that a proportion $\mu \in (0,1)$ of firms do not update their information each period: consequently, given the assumed stochastic properties of the shock, such firms leave prices unchanged at their previous period's value.⁷ This assumption, also present in Baeriswyl and Cornand (2010) and Hahn (2010), can be rationalized in terms of Mankiw and Reis's (2002) concept of 'sticky information', associated with the presumed costs of information acquisition and processing or Carroll's (2003) epidemiological model. The non-adjustment of prices by firms that fail to revise their information provides a source of price dispersion additional to that arising from heterogeneity of the information possessed by firms which do update, while also impacting on volatility of the output gap.

⁶ The equilibrium degree of coordination is $1-\beta$, while its socially optimal counterpart is $1-\lambda$, with the relationship between β and λ implying $1-\beta < 1-\lambda$. Note that for $\beta = 1$ individual firms do not perceive any strategic complementarity, despite the social benefit of price alignment.

⁷ It is straightforward to extend the analysis of the paper to consider autocorrelated disturbances, now determined according to $\phi_t = \gamma\phi_{t-1} + \zeta_t$, where $0 < \gamma < 1$ and $\zeta_t \sim U(-\infty, \infty)$. Assuming all firms learn of the actual value of ϕ at the end of the period in which it occurs and incorporate this information into the subsequent period's prices, while a fraction $1-\mu$ also use current information regarding the innovation ζ , leaves the analysis which follows essentially unchanged.

Firms are ordered such that those who fail to update their information each period are distributed along the interval $[0, \mu)$, while those who employ current information to determine their optimal price lie in the interval $[\mu, 1]$. The common price of the goods produced by the former group is normalized for convenience, but without any loss of generality⁸, at zero. It follows that the price level is described by:

$$p = \int_{\mu}^1 p_i di = \int_{\mu}^1 \hat{p}_i di \quad (4)$$

Using \bar{p} to denote the average price of firms which adjust prices in response to current information, i.e. $\bar{p} = \frac{1}{1-\mu} \int_{\mu}^1 p_i di = \frac{1}{1-\mu} p$, the measure of price dispersion relevant for welfare, as identified in (1), can be written as:

$$\int_0^1 (p_i - p)^2 di = \mu(1-\mu)\bar{p}^2 + \int_{\mu}^1 (p_i - \bar{p})^2 di \quad (5)^9$$

The first term on the right hand side of (5) captures the cross-firm price dispersion resulting from the presence both of firms whose own prices do not incorporate current information and of firms who use their observations of contemporaneous signals in setting prices. The second term then reflects the price dispersion within the latter group as a result of the heterogeneity of those signals.

⁸ The normalization is completely inconsequential, but simplifies notation. Allowing for differences in the prices of non-updating firms would add a further source of price dispersion, distinct from those arising with a common price, but one that gives rise to a deadweight loss which cannot be influenced by policy.

⁹ If the initial average price of updating firms is identical to that of the non-updating firms, then this expression can be written, as in Hahn (2010), as $\int_0^1 (p_i - p)^2 di = \frac{\mu}{1-\mu} \pi^2 + \int_{\mu}^1 (p_i - \bar{p})^2 di$, where $\pi (\equiv p - p_{-1})$ represents the inflation rate. Thus, price inflexibility (whatever its source) on the part of a sub-set of firms introduces a route through which inflation impacts on social welfare in contexts where complete price flexibility would imply inflation *per se* was welfare-neutral (see Woodford, 2003). Note that substitution of the above expression into (1) gives rise to a welfare function which, apart from the term associated with information heterogeneity across the updating firms, is identical to that employed in numerous studies where the social loss function is not derived explicitly from an underlying micro-founded model.

Individual prices set by updating firms are based on imperfect information regarding the realization of the aggregate demand shock. Specifically, firm i ($i \in [\mu, 1]$) observes an idiosyncratic signal, η_i (unobserved by any other firm), of ϕ , where $\eta_i = \phi + \xi_i$, with $\xi_i \sim N(0, \sigma_\xi^2)$ and $E(\xi_i \xi_j) = 0$ for $j \neq i$, while $\int_{\mu}^1 \xi_i di = 0$. In addition to this private signal, firm i also potentially has access to a public signal, whose source is the central bank. This public signal is then combined with the firm's private signal to form optimal estimates (conditional on the observed signals) of ϕ , p and g to determine \hat{p}_i as defined by (2').

Prior to both making any announcement and implementing monetary policy, the central bank observes its own noisy signal, δ , of the aggregate demand shock, with $\delta = \phi + \varepsilon$, where $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ and $E(\varepsilon \xi_i) = 0, \forall i$. This signal is private, in the sense that it cannot be observed directly by any firm. However, depending on the transparency regime, the central bank may choose to disclose its value, possibly in a modified form, to the private sector. Transparency is modeled in terms of how precisely the central bank reveals its own information to the private sector. The public signal, ν , potentially introduces additional noise to the central bank's private signal before being communicated to the public. In particular, the central bank is assumed to commit to a disclosure rule of the form: $\nu = \delta + \omega$ where $\omega \sim N(0, \sigma_\omega^2)$, $E(\omega \xi_i) = 0, \forall i$ and $E(\omega \varepsilon) = 0$. This characterization of disclosure policy is that followed in James and Lawler (2010) and used extensively in the monetary policy transparency literature,¹⁰ allowing the extreme cases of zero ($\sigma_\omega^2 \rightarrow \infty$) and full ($\sigma_\omega^2 \rightarrow 0$) disclosure to be captured, as well as intermediate degrees of central bank transparency.

Firm i 's ($i \in [\mu, 1]$) observation of the public signal is used together with its private signal¹¹ to form an expectation of ϕ , $E_i(\phi) = [(\sigma_\varepsilon^2 + \sigma_\omega^2)\eta_i + \sigma_\xi^2 \nu] / (\sigma_\varepsilon^2 + \sigma_\omega^2 + \sigma_\xi^2)$. Similarly, combining

¹⁰ Notable examples include Cukierman and Meltzer (1986) and Faust and Svensson (2001, 2002). An alternative representation of transparency is in terms of Cornand and Heinemann's (2008) concept of the 'degree of publicity' which identifies the extent of transparency with the proportion of private sector agents to whom the central bank's signal is disclosed. Cornand and Heinemann's approach is applied in Walsh (2007) and in James and Lawler (2012a, 2012b).

¹¹ It is assumed that no firm observes the value of the central bank's policy instrument before setting its price. Thus we abstract from the potential signaling role, emphasized in Baeriswyl and Cornand's (2010) study, that

η_i and ν allows an updating firm to improve its estimate of the central bank's private signal, δ , compared to that arrived at using the public signal alone, with the optimal estimate of δ described by $E_i(\delta) = [\sigma_\omega^2 \eta_i + (\sigma_\varepsilon^2 + \sigma_\xi^2) \nu] / (\sigma_\varepsilon^2 + \sigma_\omega^2 + \sigma_\xi^2)$. The precision of ν as a signal of δ (as measured by σ_ω^{-2}) plays a crucial role in determining the strength of private sector pricing responses to central bank announcements and thereby in influencing the impact of monetary policy.

The objective of monetary policy is to maximize expected welfare. The central bank is assumed to adjust its instrument according to a pre-determined rule; in the context of the present model, the optimal rule defines the setting of the instrument as a linear function of the central bank's current information, which consists of both the actual realization of its private signal, δ , and the publicly announced value of the latter, ν .¹² Hence, the setting of the monetary policy instrument is determined according to:

$$g = \rho_1 \delta + \rho_2 \nu \quad (6)$$

The values of the rule parameters, i.e. ρ_1 and ρ_2 are chosen to maximize expected welfare, with these values assumed to be public knowledge.

3. Equilibrium and the Welfare Effects of Transparency

3.1 Equilibrium and optimal policy

We begin by identifying the equilibrium pricing decisions of those firms which set prices using current information, taking as given both the values of the rule parameters and the

policy might play. The analysis of James and Lawler (2011) suggests that allowing for such a signaling role does not change the conclusions regarding optimal transparency compared to when such a role is absent.

¹² The central bank is assumed not to observe any p_i ($i \in [\mu, 1]$) prior to choosing its setting of policy; while it has knowledge of the common price of non-updating firms' output our normalization means that its value does not appear explicitly in the policy rule as described by (6). We further note that if all firms used current information in setting prices, a policy response to ν would be irrelevant to welfare outcomes and the optimal value of ρ_2 would consequently be indeterminate.

‘degree of transparency’, as captured by σ_ω^2 . The model structure implies that the optimal price of an updating firm can be expressed as a linear function of the two signals that it observes, i.e.:

$$p_i = \kappa_1 \eta_i + \kappa_2 \nu \quad (7)$$

From the property $\int_0^1 \xi_i di = 0$, it follows that the average price set by updating firms, \bar{p} , is given by $\bar{p} = \kappa_1 \phi + \kappa_2 \nu$, implying $p = (1 - \mu)(\kappa_1 \phi + \kappa_2 \nu)$. Using the latter, together with the policy rule (6), to substitute for p and g respectively in (2'), firm i 's ($i \in [\mu, 1]$) expectations of ϕ and δ , as previously identified, can then be used in conjunction with (2') to express the firm's optimal price as a function of η_i and ν . With $p_i = \hat{p}_i$, equating coefficients on the two signals between this expression and equation (7) allows the equilibrium values of the κ_i to be determined:

$$\kappa_1 = \frac{\beta[\sigma_\varepsilon^2 + (1 + \rho_1)\sigma_\omega^2]}{[\sigma_\xi^2 + \Delta(\sigma_\varepsilon^2 + \sigma_\omega^2)]} \quad \kappa_2 = \frac{\beta}{\Delta} \left[\frac{\sigma_\xi^2 + (\sigma_\xi^2 + \Delta\sigma_\varepsilon^2)\rho_1}{\sigma_\xi^2 + \Delta(\sigma_\varepsilon^2 + \sigma_\omega^2)} + \rho_2 \right]$$

where $\Delta \equiv [\beta + \mu(1 - \beta)]$.

It follows that each updating firm sets its individual product price according to:

$$p_i = \beta \left\{ \frac{[\sigma_\varepsilon^2 + (1 + \rho_1)\sigma_\omega^2]}{[\sigma_\xi^2 + \Delta(\sigma_\varepsilon^2 + \sigma_\omega^2)]} \eta_i + \frac{1}{\Delta} \left[\frac{\sigma_\xi^2 + (\sigma_\xi^2 + \Delta\sigma_\varepsilon^2)\rho_1}{\sigma_\xi^2 + \Delta(\sigma_\varepsilon^2 + \sigma_\omega^2)} + \rho_2 \right] \nu \right\} \quad (8)$$

The average price across such firms is therefore:

$$\bar{p} = \beta \left\{ \frac{[\sigma_\varepsilon^2 + (1 + \rho_1)\sigma_\omega^2]}{[\sigma_\xi^2 + \Delta(\sigma_\varepsilon^2 + \sigma_\omega^2)]} \phi + \frac{1}{\Delta} \left[\frac{\sigma_\xi^2 + (\sigma_\xi^2 + \Delta\sigma_\varepsilon^2)\rho_1}{\sigma_\xi^2 + \Delta(\sigma_\varepsilon^2 + \sigma_\omega^2)} + \rho_2 \right] \nu \right\} \quad (9)$$

with the economy-wide price level following directly from (9):

$$p = (1 - \mu)\beta \left\{ \frac{[\sigma_\varepsilon^2 + (1 + \rho_1)\sigma_\omega^2]}{[\sigma_\xi^2 + \Delta(\sigma_\varepsilon^2 + \sigma_\omega^2)]} \phi + \frac{1}{\Delta} \left[\frac{\sigma_\xi^2 + (\sigma_\xi^2 + \Delta\sigma_\varepsilon^2)\rho_1}{\sigma_\xi^2 + \Delta(\sigma_\varepsilon^2 + \sigma_\omega^2)} + \rho_2 \right] \nu \right\} \quad (10)$$

Equations (8), (9) and (10) can now be used, in combination with equations (3) and (6), to substitute for the price dispersion and output gap terms in (1). To identify the optimal values of the two policy rule parameters, the unconditional expectation of the resulting welfare expression is taken, from which it is immediately evident that expected welfare is bounded below only if ρ_1 and ρ_2 are constrained to sum to -1 . This feature of the framework follows from the specified distribution for ϕ , in combination with the assumed presence of a subset of firms which do not respond to current information regarding the realization of the aggregate demand shock. In this context, only if the central bank fully offsets its own expectation of the shock (which requires $\rho_1 + \rho_2 = -1$) is the variance of the output of non-updating firms bounded.

Imposing the constraint in the expression for expected welfare allows the optimal values of the two policy parameters to be found from the corresponding first order conditions. Denoting these values by ρ_1^* and ρ_2^* , we find:

$$\rho_1^* = -\frac{[\sigma_\xi^2 + \Psi(\sigma_\varepsilon^2 + \sigma_\omega^2)][\sigma_\xi^2 + \mu(\sigma_\varepsilon^2 + \sigma_\omega^2)]}{\sigma_\xi^4 + (\Delta^2\sigma_\varepsilon^2 + \mu\Psi\sigma_\omega^2)(\sigma_\varepsilon^2 + \sigma_\omega^2) + [2\Delta\sigma_\varepsilon^2 + (\mu + \Psi)\sigma_\omega^2]\sigma_\xi^2} \quad (11a)$$

$$\rho_2^* = \frac{(1-\mu)\beta\{(\theta-2)\sigma_\xi^2 + [\mu(\theta-1) - \Delta](\sigma_\varepsilon^2 + \sigma_\omega^2)\}\sigma_\xi^2}{\sigma_\xi^4 + (\Delta^2\sigma_\varepsilon^2 + \mu\Psi\sigma_\omega^2)(\sigma_\varepsilon^2 + \sigma_\omega^2) + [2\Delta\sigma_\varepsilon^2 + (\mu + \Psi)\sigma_\omega^2]\sigma_\xi^2} \quad (11b)$$

where $\Psi \equiv \beta\theta + \mu(1 - \beta\theta)$.

The identified values of ρ_1^* and ρ_2^* allow us to determine equilibrium expected welfare with policy set optimally:

$$E(W | \phi) \Big|_{\rho_1=\rho_1^*, \rho_2=\rho_2^*} = -\left(\frac{\beta\sigma_\varepsilon^2}{\theta}\right) \frac{[\sigma_\xi^2 + \Psi(\sigma_\varepsilon^2 + \sigma_\omega^2)][\sigma_\xi^2 + \mu(\sigma_\varepsilon^2 + \sigma_\omega^2)]}{\sigma_\xi^4 + (\Delta^2\sigma_\varepsilon^2 + \mu\Psi\sigma_\omega^2)(\sigma_\varepsilon^2 + \sigma_\omega^2) + [2\Delta\sigma_\varepsilon^2 + (\mu + \Psi)\sigma_\omega^2]\sigma_\xi^2} \quad (12)$$

The above expression provides the basis for the analysis of how central bank transparency, as measured by the quality of the public signal provided by the central bank, i.e. σ_ω^2 , impacts on welfare. We begin by comparing the welfare implications of the limiting cases of zero ($\sigma_\omega^2 \rightarrow \infty$) and full ($\sigma_\omega^2 \rightarrow 0$) disclosure. Besides being prominent cases in the literature in

their own right, these extreme instances of transparency also provide the appropriate benchmarks when identifying the globally optimal degree of transparency.

3.2 Limiting cases of transparency and expected welfare

3.2.1 The welfare implications of zero transparency

Under a regime of zero transparency, the public signal conveys no useful information regarding the central bank's observation of δ : reflecting this, optimal policy is described in terms of a monetary policy response solely to the central bank's private signal, with $\rho_1^* = -1$ and $\rho_2^* = 0$. To determine the value of expected welfare under zero disclosure, which we denote by $E(W_{ZD}^*)$, we find the limit of (12) as $\sigma_\omega^2 \rightarrow \infty$:

$$E(W_{ZD}^*) = -\frac{\beta\sigma_\varepsilon^2}{\theta} \quad (\equiv -\lambda\sigma_\varepsilon^2) \quad (13)$$

In interpreting this expression, we first note from (8) that (with $\rho_1 = -1$, $\rho_2 = 0$) as $\sigma_\omega^2 \rightarrow \infty$

$p_i = 0$, i.e. updating firms do not adjust their prices in light of the observed values of their

private signals. This is because, in the absence of any informative public announcement, firm i 's best estimate of any other updating firm's price is its own price, p_i , implying from (2')

and the relationship between p and \bar{p} that $p_i = \frac{\beta E_i}{i \in [\mu, 1]} (\phi + g) / \Delta$. Moreover, with the

setting of monetary policy directed at fully neutralizing the central bank's own expectation of ϕ , $E_i (\phi + g) = 0$.¹³ Consequently, the prices of updating firms are, under zero

transparency, unresponsive to observations of private signals and therefore remain perfectly aligned, both relative to each other and with respect to non-updating firms' prices. Thus, there is no price dispersion and the loss in welfare identified by (13) arises purely from aggregate output fluctuations associated with central bank expectational errors.

3.2.2 Expected welfare with full transparency

¹³ This follows from the property $E_i(\delta | \eta_i) = E_i(\phi | \eta_i) = \eta_i$.

In this case, the central bank fully discloses the observed value of its private signal to the public, i.e. $\nu \equiv \delta$, though by assumption only a proportion $1-\mu$ of firms use the announcement to update the information set on which their pricing decisions are based. Denoting expected welfare in the case of full disclosure, and with policy set optimally,¹⁴ by $E(W_{FD}^*)$, we find, taking the limit of (12) as $\sigma_\omega^2 \rightarrow 0$:

$$E(W_{FD}^*) = -\left(\frac{\beta\sigma_\varepsilon^2}{\theta}\right) \frac{(\sigma_\xi^2 + \mu\sigma_\varepsilon^2)(\sigma_\xi^2 + \Psi\sigma_\varepsilon^2)}{(\sigma_\xi^2 + \Delta\sigma_\varepsilon^2)^2} \quad (14)$$

As for the case of zero transparency, central bank expectational errors, which lead to imperfect stabilization of the impact of the aggregate demand shock, are the proximate source of the welfare loss represented by (13). However, in the present case, additional forces come into play.

By combining the information content of their private signals with that conveyed by the public signal, updating firms are able to form a superior estimate of ϕ compared to that based on the public signal alone. Hence, in this case, each firm adjusts its price in response to its own expectation of the central bank's policy error. This exploitation of the information present in private signals has a beneficial welfare consequence, since it leads to a reduction in aggregate output volatility relative to the zero disclosure case.¹⁵ At the same time, however, price dispersion is introduced, both within the set of updating firms and between this group and the set of non-updating firms: of course, such dispersion is detrimental to welfare.

3.2.3 Comparing welfare outcomes in the extreme cases of transparency

¹⁴ Note that, although taking the limit of (11a) and (11b) as $\sigma_\omega^2 \rightarrow 0$ identifies unique values of the two policy parameters, with ν identical to δ the individual values of ρ_1 and ρ_2 are immaterial, so long as they are constrained to sum to unity. In this instance, as for the case of zero transparency (though for a different reason), optimal policy can be represented in terms of a single response coefficient relating to the adjustment of the instrument g to the central bank's private signal.

¹⁵ The variance of aggregate output under full disclosure is given by $E(y_{FD}^2) = (\sigma_\xi^2 + \mu\sigma_\varepsilon^2)^2 \sigma_\varepsilon^2 / (\sigma_\xi^2 + \Delta\sigma_\varepsilon^2)^2$. It is straightforward to show that this is smaller in value than the corresponding variance under zero disclosure, σ_ε^2 .

To evaluate the relative welfare performances of the limiting cases of disclosure, we subtract equation (14) from (13). Following some straightforward, though tedious, algebraic manipulation, we arrive at:

$$E(W_{ZD}^*) - E(W_{FD}^*) = \frac{(1-\mu)\beta^3\sigma_\varepsilon^6(\mu + \sigma_\xi^2\sigma_\varepsilon^{-2})}{\theta(\sigma_\xi^2 + \Delta\sigma_\varepsilon^2)^2} \left[\frac{(\theta-2)}{\beta} - \frac{(1-\mu)}{(\mu + \sigma_\xi^2\sigma_\varepsilon^{-2})} \right] \quad (15)$$

This expression cannot be signed on *a priori* grounds and, thus, which of the two extreme instances of transparency delivers a superior welfare outcome is, in general, indeterminate. However, noting that $(1-\mu)/(\mu + \sigma_\xi^2\sigma_\varepsilon^{-2})$ is strictly decreasing in μ for $\mu \in [0,1]$, it follows that a sufficient condition for zero disclosure to outperform full disclosure is $\sigma_\varepsilon^2 / \sigma_\xi^2 < (\theta-2) / \beta$. We now turn to consider whether this inequality is likely to be satisfied in practice.

The left hand side of the inequality is simply the precision of the private signals observed by updating firms relative to the precision of the central bank's signal. Some insight into the range of values within which this ratio is likely to lie is provided by Romer and Romer's (2002) empirical study of the comparative accuracy of private-sector and Federal Reserve forecasts of U.S. inflation. They find that the latter are more accurate than the former, rationalizing this conclusion in terms of the considerable resources devoted by the Federal Reserve to forecasting. This argument, which seems likely to have validity beyond the U.S., is drawn on by Svensson (2006) in his critique of the interpretation of Morris and Shin's (2002) findings as representing an 'anti-transparency' result. Indeed, he identifies a situation in which the precisions of private sector and central bank signals are equal as representing "a rather conservative benchmark case".¹⁶ In developing our arguments, we adopt this reference point and take the maximum value of the relative precision of the two signals to be unity.

Applying this benchmark to the sufficient condition identified above for zero disclosure to be associated with welfare outcomes superior to those arising under full disclosure, the inequality becomes $\beta < (\theta-2)$. As previously noted, the case in which goods are strategic complements corresponds to $\beta \in (0,1)$: with the upper limit of this range imposed, the sufficient condition for $E(W_{ZD}^*) > E(W_{FD}^*)$ becomes simply $3 < \theta$. Consideration of whether

¹⁶ Svensson (2006), p.449.

the elasticity of substitution between varieties of goods satisfies this inequality requires reference to empirical evidence. Ball and Romer (1990), in their study of real rigidities and monetary non-neutrality, refer to such evidence in applying a mark-up value of 0.15 to their model, a figure which is broadly consistent with the values reported in Oliveira Martins, Scarpetta and Pilat (1996), and the recent study by De Loecker and Warzynski (2012).¹⁷ Such a value implies an elasticity of substitution of 7.7, a magnitude which is comfortably in excess of that needed to ensure that zero disclosure is preferable to full disclosure. Indeed, observed mark-ups exceeding the implausibly large value of 0.5 would be necessary to imply a value of θ less than 3. Thus, empirical evidence appears to be consistent with the notion that zero transparency gives rise to superior welfare outcomes compared to those resulting under full transparency.

The conclusion that for plausible values of key parameters zero transparency is associated with higher welfare than full transparency might appear, at first sight, surprising. This is particularly so since, as previously discussed, within a regime of zero transparency the information content of each updating firm's private signal is 'lost', in the sense that there is no price response to it and, consequently, attained welfare is identical to that which would arise if no firm updated its information set in the current period. The explanation for the finding lies in the feature which is at the center of the related literature. That is, when both public and (heterogeneous) private signals are observed in an economy characterized by strategic complementarities, private sector actions based on that information are potentially inefficient from the viewpoint of social welfare. James and Lawler (2010, 2011) using abstract models which incorporate strategic complementarities and heterogeneous information show that, in the presence of stabilization policy, *any* (otherwise private) information released by the policymaker is detrimental to welfare. Although this result is derived in the context of models which assume all private sector agents to base their decisions on current information, it clearly has relevance (albeit in a modified form) to the current setting.

To identify this inefficiency, we consider the pricing decisions of updating firms under full transparency. Each such firm's price response to the signals that it observes reflects its estimate of the central bank's expectational error, ε , in respect of the aggregate demand

¹⁷ Oliveira Martins et al's (1996) paper covers 14 OECD countries over the period 1970-92, while De Loecker and Warzynski's (2012) recent study focuses on Slovenian data relating to the years 1994-2000.

shock. However, this response is muted somewhat by the coordination motive associated with the strategic complementarity which characterizes the model. Any price adjustment by updating firm i based on the signals that it observes will inevitably give rise to a deviation from the common price of non-updating firms' output. Furthermore, with firm i 's estimate of ε derived, in part, from its own private signal, the idiosyncratic element inherent in this signal represents a source of departure from the average price set by other updating firms. In the presence of a coordination motive, these factors lead to a reduced sensitivity of price adjustments to firms' estimates of the central bank's expectational errors regarding ϕ . Firm i 's price under full disclosure of its signal by the central bank is given by:

$$p_i = \frac{\beta \sigma_\varepsilon^2 (\xi_i - \varepsilon)}{\sigma_\varepsilon^2 + \Delta \sigma_\varepsilon^2} \quad (16)$$

which, for $\beta \in (0,1)$, is smaller in absolute magnitude than $E_i(\varepsilon)$.¹⁸

Notwithstanding the moderating effect which the coordination motive has on price adjustment, the incentives facing the individual firm to align its product price with those of other firms understate the true social benefit of coordination. Consequently, individual prices remain too responsive to the information content of observed signals. The precise meaning of 'too responsive' in this context can be clarified by reference to the pricing decisions which would maximize social welfare, i.e. the set of prices that would be chosen by a benevolent central planner whose objective was to maximize the expected value of (1), taking as given the existing dispersion of information. Denoting the socially optimal value of p_i by \tilde{p}_i ,

determined according to $\tilde{p}_i = \tilde{\kappa}_1 \eta_i + \tilde{\kappa}_2 \nu$, we identify the values of $\tilde{\kappa}_1$ and $\tilde{\kappa}_2$ which, with full disclosure and policy set optimally, maximize $E(W)$. This exercise yields the following expression for the socially optimal price of each updating firm:

$$\tilde{p}_i = \frac{\lambda \sigma_\varepsilon^2 (\xi_i - \varepsilon)}{\sigma_\varepsilon^2 + \Delta' \sigma_\varepsilon^2} \quad (17)$$

¹⁸ It is evident from equation (2') that for $\beta = 1$, i.e. if there are no perceived strategic complementarities at the individual firm level, the price set by any updating firm is simply its expectation of the central bank's error in estimating ϕ , with $E_i(\varepsilon) = \sigma_\varepsilon^2 (\xi_i - \varepsilon) / (\sigma_\varepsilon^2 + \sigma_\varepsilon^2)$. Of course, setting $\beta = 1$ in (16) yields this value for updating firm i 's price.

where $\Delta' \equiv [\lambda + \mu(1 - \lambda)]$

With $\lambda < \beta$ it is evident, as previously indicated, that with full transparency the equilibrium prices of updating firms are overly-responsive, relative to the social-efficiency benchmark, to estimates of the central bank's expectational error. Such overreaction underlies the finding that, in the presence of optimal policy intervention, zero transparency is, for plausible parameter values, superior to full transparency. Indeed, if it were possible to induce updating firms to adjust their prices in a socially optimal manner, full disclosure would lead to unambiguously better welfare outcomes than zero disclosure. Denoting expected welfare when the central bank fully reveals its estimate of ϕ , while policy is adjusted optimally and updating firms set prices in accordance with (17) by $E(\tilde{W}_{FD}^*)$, we find:

$$E(\tilde{W}_{FD}^*) = -\frac{\lambda(\sigma_\xi^2 + \mu\sigma_\varepsilon^2)\sigma_\varepsilon^2}{\sigma_\xi^2 + \Delta'\sigma_\varepsilon^2} \quad (18)$$

Comparison of (18) with (13) directly reveals $E(\tilde{W}_{FD}^*) > E(W_{ZD}^*)$. If the central bank's public signal and updating firms' private signals are together exploited efficiently in setting prices, then welfare outcomes under full transparency will inevitably be superior to those associated with zero transparency, since in the latter case the information content of private signals is left unutilized.

3.3 The optimal degree of transparency

Although, as discussed, plausible values for key parameters of the model imply that, with policy conducted optimally, zero transparency outperforms full transparency, this leaves open the question of whether some partial degree of transparency (i.e. with the public signal disclosed by the central bank providing imperfect information on its own private signal's value) might dominate both extreme cases of central bank disclosure. This possibility can be considered by examining the properties of the relationship between expected welfare and the quality of the public signal announced by the central bank. As might be expected from comparison of the two extreme cases of transparency, the nature of this relationship is dependent on parameter values. However, for $\sigma_\varepsilon^2 \leq \sigma_\xi^2$ and $2 < \theta$, the following two key properties of the relationship can be established:

(i) $0 < \partial E(W^*) / \partial \sigma_\omega^2 |_{\sigma_\omega^2=0}$: i.e. beginning from a position of full transparency, a ‘small’ decline in the quality of the public signal is welfare improving.

(ii) Defining $\hat{\mu} \equiv \beta / (\beta + \theta - 2)$, then for any μ such that $\hat{\mu} < \mu < 1$, no value of $\sigma_\omega^2 \in \mathbb{R}^+$ satisfies the first order condition $\partial E(W) / \partial \sigma_\omega^2 = 0$, while for any μ such that $0 < \mu < \hat{\mu}$ there exists a unique value of $\sigma_\omega^2 \in \mathbb{R}^+$ that satisfies $\partial E(W) / \partial \sigma_\omega^2 = 0$.

We note that the parameter restrictions jointly sufficient to ensure that these two properties apply are together weaker than those identified as ensuring zero transparency is associated with welfare outcomes superior to those arising under full transparency. Hence the arguments previously presented suggest that it is reasonable to suppose that they will be satisfied in practice.

Turning now to the implications of these properties, the first establishes that full transparency cannot be optimal, since it ensures that there exists a range of strictly positive values of σ_ω^2 for which expected welfare exceeds that attained for $\sigma_\omega^2 = 0$. The second property indicates that, depending on the value of μ relative to the critical value, $\hat{\mu}$, either zero transparency or partial transparency can be optimal. In particular, if the proportion of non-updating firms is sufficiently large, in the sense that it exceeds $\hat{\mu}$, then zero transparency is optimal, while for values of μ below $\hat{\mu}$ partial transparency is superior.¹⁹

To interpret these conclusions, consider an increase in σ_ω^2 beginning from a position of full transparency. A decline in the precision of the public signal will lead to a reduction in the responsiveness of updating firms’ prices to their estimates of the central bank’s expectational error. Although increased aggregate output volatility results, the associated detrimental effect on welfare is initially outweighed by the beneficial consequences of lower price dispersion. If

¹⁹ We note that there is a discontinuity in the relationship between $E(W^*)$ and σ_ω^2 as $\mu \rightarrow 0$. This discontinuity implies $\lim_{\mu \rightarrow 0} E(W_{ZD}^*) \neq E(W_{ZD}^*) |_{\mu=0}$ and, for this reason, the current framework does not encompass the unambiguous (i.e. regardless of parameter values) superiority of zero disclosure when all firms base their pricing decisions on current information. The discontinuity arises from the combination of the assumed distribution of ϕ and the presence of non-updating firms, and is not present if ϕ is taken to be normally distributed.

μ lies in the interval $(\hat{\mu}, 1)$, then the latter effect of reduced public signal quality is dominant for all values of σ_ω^2 and, hence, zero transparency is optimal. However, for $\mu \in (0, \hat{\mu})$, as σ_ω^2 increases from zero, a degree of signal precision will eventually be attained such that the marginal cost of greater output volatility associated with any further deterioration in the quality of the public signal will exceed the marginal benefit of less price dispersion. It follows that, in this case, an intermediate degree of transparency maximizes welfare.

The potential for partial transparency to be optimal does not arise in the analyses contained in James and Lawler (2011, 2012a), in which all agents choose their actions based on currently available information, both private and public, and where the former, while dispersed, is homogeneous in quality. Rather, in both contributions, zero transparency unambiguously attains the welfare maximum in the presence of optimal policy. However, the possibility that an intermediate degree of transparency, in which the policymaker partially reveals its information, might be desirable is present in James and Lawler (2012b). The key feature of the framework developed therein is that the quality of private sector agents' private information, as represented by the precision of the idiosyncratic signals they observe is heterogeneous. Heterogeneity of quality of information is clearly also present in the current model, though extended to public as well as private information, with the precision of non-updating firms signals being, in effect, zero. Thus, it appears that differences in information precision across agents represent a factor associated with the potential for partial transparency to be optimal.

A further aspect of the implications of heterogeneous information quality observed in James and Lawler (2012b) shared by the present analysis is that optimal stabilization policy, in combination with the optimal transparency regime, is unable to attain the 'first best': i.e. the welfare outcome associated with all agents reacting in a collectively efficient fashion to current information. This contrasts with the findings presented in James and Lawler (2011, 2012a), but reflects the fact that, when information quality is non-homogeneous across agents, the policy rule cannot be tailored to induce all agents to choose their actions in a manner consistent with overall efficiency. This shortcoming of stabilization policy leads to the question of whether an alternative, or additional, form of policy intervention might be capable of improving welfare outcomes and we turn to consider this issue in the next section.

4. An Optimally-Designed Pigouvian Tax

A means of addressing the inefficiencies arising from dispersed information in economies in which strategic complementarities are present is identified by Angeletos and Pavan (2007, 2009). They show that an appropriately-designed tax regime can be used to manipulate the incentives which agents have to respond to information in such a way that, subject to the degree of information dispersion that characterizes the economy, the first best (i.e. socially efficient) outcome is attained. In this respect, such a tax represents an alternative to the direct policy intervention demonstrated to be capable of achieving the social optimum in James and Lawler (2011, 2012a).

The optimal Pigouvian tax, which entails a marginal tax rate that depends on the value of aggregate actions, is derived by Angeletos and Pavan in a setting in which all agents observe, and respond to, current information and where such information is of homogeneous quality across all agents. However, in their analysis of the implications of heterogeneous information quality in the context of the Morris and Shin (2002) framework, James and Lawler (2012b) show that Angeletos and Pavan's conclusions extend to a setting in which the quality of private information potentially differs between agents. As such, the optimal tax regime, in correcting the inefficient use of information, ensures that full transparency maximizes social welfare while obviating the need for stabilization policy.

Here we briefly consider the efficacy of a Pigouvian tax in the current model, in which a proportion of firms do not respond to current information at all. The tax (subsidy) imposed on each updating firm is a linear function of the deviation of its own price from the average price set by other updating firms, with the marginal tax rate assumed to be determined by this average price and the realization of the public signal. Each updating firm's net tax liability,²⁰ t_i is described by $t_i \equiv 2(\tau_{\bar{p}}\bar{p} + \tau_{\delta}\delta)(\bar{p} - p_i)$, implying its optimal price is:

$$p_i = E_i \left[\left(1 - \beta - \frac{\tau_{\bar{p}}}{1 - \mu} \right) p + \beta(g + \phi) - \tau_{\delta}\delta \right] \quad (19)$$

²⁰ Each firm also receives a lump-sum transfer equal to the average tax liability across all firms.

Applying the solution procedure of Section 3, we derive an expression for expected welfare for given values of the policy coefficients, ρ_1 , ρ_2 , $\tau_{\bar{p}}$ and τ_{δ} , then determine the optimal values of these parameters. These are:

$$\rho_1^* = -\left(\frac{\sigma_{\xi}^2 + \mu(\sigma_{\varepsilon}^2 + \sigma_{\omega}^2)}{\sigma_{\xi}^2 + \Delta'\sigma_{\varepsilon}^2 + \mu\sigma_{\omega}^2}\right) \quad (20a)$$

$$\rho_2^* = -\left(\frac{\lambda(1-\mu)\sigma_{\varepsilon}^2}{\sigma_{\xi}^2 + \Delta'\sigma_{\varepsilon}^2 + \mu\sigma_{\omega}^2}\right) \quad (20b)^{21}$$

$$\tau_{\bar{p}}^* = \frac{(\beta-\lambda)}{\lambda} \left(\frac{\sigma_{\xi}^2 + \mu(\sigma_{\varepsilon}^2 + \sigma_{\omega}^2)}{\sigma_{\varepsilon}^2 + \sigma_{\omega}^2}\right) \quad (20c)$$

$$\tau_{\delta}^* = 0 \quad (20d)$$

Substituting these values into the expected welfare expression that they collectively minimize, we find:

$$E(W^*) = -\lambda \left(\frac{\sigma_{\xi}^2 + \mu(\sigma_{\varepsilon}^2 + \sigma_{\omega}^2)}{\sigma_{\xi}^2 + \Delta'\sigma_{\varepsilon}^2 + \mu\sigma_{\omega}^2}\right) \sigma_{\varepsilon}^2 \quad (21)$$

It is straightforward to demonstrate that expected welfare is strictly decreasing in σ_{ω}^2 , and thus the optimal value of σ_{ω}^2 is zero, i.e. implementation of the optimal tax regime ensures that full transparency maximizes welfare. Moreover, by setting $\sigma_{\omega}^2 = 0$ in (21), it is directly evident from comparison with (18) that the optimal Pigouvian tax eliminates the inefficiency in pricing decisions identified in Section 3.2.3 completely.²² This conclusion is consistent with the findings of Angeletos and Pavan (2007b, 2009) and James and Lawler (2012b), which show an appropriately-designed tax regime to be capable of modifying the incentives facing agents in such a way as to induce socially-efficient responses to available information. However, in the current instance, direct monetary policy intervention is also an integral element in attaining a welfare maximum. In the presence of firms which do not base pricing decisions on current information, correcting the inefficiencies which would otherwise characterize the prices set by updating firms is not, in itself, sufficient to attain the social

²¹ As in the absence of Pigouvian tax, the stabilization policy coefficients must sum to -1 .

²² In the presence of the Pigouvian tax (with τ_{δ} set to zero), under full transparency updating firm i 's pricing decision is described by $p_i = \beta\sigma_{\varepsilon}^2(\xi_i - \varepsilon)/[\sigma_{\xi}^2 + (\Delta + \tau_{\bar{p}})\sigma_{\varepsilon}^2]$. The optimal value of $\tau_{\bar{p}}$ ensures that this coincides with the efficient response to the individual firm's estimate of the central bank's expectational error.

optimum. In this context, macroeconomic policy intervention is an essential complement to micro-based Pigouvian taxes.

5. Conclusions

Previous work that has considered the desirability, or otherwise, of central bank transparency has typically examined the issue using models in which all agents are assumed to respond to current information. Motivated by the ideas underlying Mankiw and Reis's (2002) concept of 'sticky information', as well as Carroll's (2003) epidemiological approach to information diffusion, the aim of the present study has been to identify how the presence of some agents who do not continually update the information on which their choices are based influence the welfare consequences of transparency.

The macroeconomic framework used to analyze this issue has the pricing decisions of individual firms at its center. Within the model, even if all firms set prices on the basis of current information regarding aggregate demand shocks, as embodied in observations of both public and private signals, macroeconomic equilibrium would not correspond to the socially efficient outcome. This feature reflects the interaction between the strategic complementarity which characterizes the framework and heterogeneity of private information. The nature of the strategic complementarity is such that the incentives facing an individual firm do not adequately reflect the social benefits of price alignment. As a consequence, firms place too much weight on their private signals, giving rise to excessive price dispersion. The presence of firms whose prices do not incorporate current information on aggregate demand shock realizations then leads both to increased price dispersion and to greater output volatility. Although monetary policy intervention is able to mitigate the associated welfare losses, stabilization is inevitably imperfect. Thus, the question is raised of whether public disclosure of the central bank's private information can bring the economy closer to the social optimum.

The foregoing analysis has demonstrated that, within the framework, there is no unequivocal answer this question. If all prices responded to current information then, given optimally-designed policy intervention, a regime of zero transparency would invariably maximize welfare, reflecting the principles exemplified in James and Lawler (2011, 2012a). However, the presence of non-updating firms introduces the possibility that either full transparency or some degree of partial transparency might be optimal. In this regard, as noted in Section 3,

the conclusions echo those arrived at in James and Lawler (2012b), which shares the feature of the current model that the quality of information differs across groups of agents.

Nonetheless, for empirically plausible parameter values, full transparency is welfare-dominated by zero transparency. With stabilization policy conducted optimally, a regime of zero transparency eliminates price dispersion completely. Although this comes at the expense of greater aggregate output volatility compared to the case of full transparency, providing reasonable parameter restrictions hold then the welfare costs associated with this feature will be outweighed by the benefits arising from the absence of price dispersion. However, zero transparency implies that the information content of updating firms' private signals is left completely unexploited. Reflecting this, partial disclosure of the central bank's signal, which induces updating firms to utilize their private information in order to improve their estimates of the central bank's forecasting error, can potentially (that is, providing the proportion of updating firms is sufficiently large) improve welfare compared to zero transparency.

The factor ultimately responsible for the likely inferiority of full transparency is the divergence within the model between the socially optimal and the equilibrium degrees of coordination. As discussed, this divergence is reflected in private incentives to align prices that understate the social benefits of such alignment. A Pigouvian tax, as originally considered by Angeletos and Pavan (2007b, 2009), provides a potential means of modifying the incentives facing firms in a manner that promotes an efficient use of information. Under the optimal tax scheme a regime of full transparency can, indeed, ensure that equilibrium outcomes coincide with socially efficient outcomes. However, unlike in Angeletos and Pavan (2007b, 2009) and James and Lawler (2012b), appropriately-formulated stabilization policy is also an essential component in attaining optimal outcomes.

Notwithstanding the transparency implications of Pigouvian taxes, the findings that emerge from our analysis clearly cannot be viewed as providing endorsement of any general case for central bank disclosure. On the contrary, in the absence of corrective taxes, reference to empirical evidence on key model parameters suggests that, despite the theoretical ambiguities present, there should be a presumption against central banks fully revealing their information publicly. This conclusion reflects the principle, identified in James and Lawler (2011, 2012a), that when private sector agents respond to information in a collectively suboptimal manner, policymakers should not disclose their private information but, instead, use it to guide their

conduct of stabilization policy. Although the presence of agents who do not respond at all to current information dilutes this principle to a degree, it nonetheless retains its core validity.

References

Adam, K. (2007). "Optimal monetary policy with imperfect common knowledge." *Journal of Monetary Economics* 54 (2), 276-301.

Angeletos, G.-M., Pavan, A. (2007a). "Efficient Use of Information and Social Value of Information." *Econometrica* 75(4), 1103-42.

- Angeletos, G.-M., Pavan, A. (2007b). "Socially Optimal Coordination: Characterization and Policy Implications." *Journal of the European Economic Association* 5(2-3), 585-593.
- Angeletos, G.-M., Pavan, A. (2009). "Policy with Dispersed Information." *Journal of the European Economic Association* 7(1), 11-60.
- Baeriswyl, R., Cornand, C. (2010). "The Signaling Role of Policy Actions." *Journal of Monetary Economics* 57, 682-695.
- Ball, L., Romer, D. (1990). "Real Rigidities and the Non-Neutrality of Money." *Review of Economic Studies* 57, 183-203.
- Blinder, A.S., Ehrmann, M., Fratzscher, M., De Haan, J., Jansen, D.J. (2008). "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence." *Journal of Economic Literature*, 46, 910-45.
- Carroll, C.D. (2003). "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics*, 118(1), 269-298.
- Cornand, C., Heinemann F. (2008). "Optimal degree of public information dissemination." *Economic Journal* 118, 718-742.
- Cukierman, A., Meltzer, A.H. (1986). "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information." *Econometrica* 54, 1099-1128.
- De Loecker, J., Warzynski, F. (2012). "Markups and Firm-Level Export Status." *American Economic Review* 102(6), 2437-71.
- Dincer, N.N., Eichengreen, B. (2007). "Central Bank Transparency: Where, Why, and with What Effects?" Working Paper No. 13003, National Bureau of Economic Research, Cambridge MA.
- Faust, J., Svensson, L.E.O. (2001). "Transparency and Credibility with Unobservable Goals." *International Economic Review* 42, 369-397.
- Faust, J., Svensson, L.E.O. (2002). "The Equilibrium Degree of Transparency and Control in Monetary Policy." *Journal of Money, Credit and Banking* 34, 520-539.
- Hahn, V. (2010) "Transparency in Monetary Policy, Signaling, and Heterogeneous Information." Working Paper, Center of Economic Research, ETH Zurich.
- James, J.G., Lawler, P. (2011). "Optimal Policy Intervention and the Social Value of Public Information." *American Economic Review*, 101(4), 1561-74.
- James, J.G., Lawler, P. (2012a). "Strategic Complementarity, Stabilization Policy and the Optimal Degree of Publicity." *Journal of Money, Credit and Banking* 44(4), 551-572.
- James, J.G., Lawler, P. (2012b). "Heterogeneous information quality, strategic complementarities and optimal policy design." *Journal of Economic Behavior and Organization* 83, 342-352.

- Mankiw, N.G., Reis, R. (2002). “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve.” *Quarterly Journal of Economics* 117(4), 1295-1328.
- Morris, S., Shin, H.S. (2002). “Social Value of Public Information.” *American Economic Review* 92(5), 1521-1534.
- Oliveira Martins, J., Scarpetta, S., Pilat, D. (1996). “Mark-Up Ratios in Manufacturing Industries: Estimates for 14 OECD Countries.” OECD Economics Department Working Papers, nº 162.
- Roca, M.F. (2010). “Transparency and Monetary Policy with Imperfect Common Knowledge.” IMF Working Paper WP/10/91.
- Romer, C.D., Romer, D.H. (2000). “Federal Reserve Information and the Behavior of Interest Rates.” *American Economic Review* 90, 429-457.
- Svensson, L.E.O. (2006). “Social Value of Public Information: Morris and Shin (2002) is Actually Pro Transparency, not Con.” *American Economic Review* 96, 448-452.
- Walsh, C.E. (2007). “Optimal Economic Transparency.” *International Journal of Central Banking* 3, 5-36.
- Woodford, M. (2002). “Inflation Stabilization and Welfare.” *Contributions to Macroeconomics* 2, 1-53
- Woodford, M. (2003). “Imperfect Common Knowledge and the Effects of Monetary Policy”, in P.Aghion, R.Frydman, J.Stiglitz and M.Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton, NJ: Princeton University Press, 25– 58.