Market Distortions and Government Transparency*

Facundo Albornoz[†] Joan Esteban[‡] Paolo Vanin[§]

May 31, 2012

Abstract

This paper investigates how government transparency depends on economic distortions. We first consider an abstract class of economies, in which a benevolent policy maker is privately informed about the exogenous state of the economy and contemplates whether to release this information. Our key result is that distortions limit communication: even if transparency is ex-ante Pareto superior to opaqueness, it is not an equilibrium whenever distortions are sufficiently high. We next confirm this broad insight in two applied contexts, in which monopoly power and income taxes are the specific sources of distortions.

JEL-Classification: D82, E61

Key-words: Government announcements, Transparency, Economic distortions, Cheap talk, Inequality

---

*We thank Gaetano Antinolfi, Francesca Barigozzi, Jordi Caballé, Antonio Cabrales, Hector Calvo Pardo, Giacomo Calzolari, Matteo Cervellati, Russell Cooper, Nicola Pavoni, Debraj Ray, Joel Sobel and seminar and conference participants in Bologna, Barcelona (EEA-ESEM), Buenos Aires (LACEA-LAMES), Istanbul (ASSET) and Lucca for useful comments. We are also grateful to Geroge-Marios Angeletos and two anonymous referees. Albornoz is grateful for support from the ESRC (RES-062-23-1360), Esteban for support from the Generalitat de Catalunya and the CICYT (SEJ2006-00369) and from the Instituto de Estudios Fiscales, and Vanin for support from the University of Padua (CPDA071899). A previous version circulated with the title "Government Information Transparency".

[†]Department of Economics, Birmingham University; email: f.albornoz@bham.ac.uk

[‡]Instituto de Análisis Económico (CSIC), Barcelona, and Barcelona Graduate School of Economics; email: joan.esteban@iae.csic.es

[§]Corresponding author: Department of Economics, Università di Bologna; email: paolo.vanin@unibo.it; tel.: +39-0512098120; fax: +39-0512098040

1 Introduction

Many governments are better informed than the private sector about future realizations of macroeconomic variables.[1] Often they transparently convey this information to the public, but at other times they do not. For instance, the US government's announcements on current or future activity have a positive real effect on the economy, confirming the fact that individuals find them informative.[2] In contrast, the widespread skepticism on contemporary Argentine or Greek official statistics provides an example of non transparent and non credible government announcements. While there may be different opportunistic reasons for governments not to be transparent, in the present paper we investigate whether a benevolent government would always reveal its private information on real macroeconomic variables. For the sake of concreteness, we focus on the case in which the government has prior information on exogenous aggregate productivity shocks that produce uniform positive (in booms) or negative (in recessions) shifts in productivity. Would a benevolent government always fully reveal this information? Is it efficient to do so? What are the determinants of government information transparency?

In an otherwise perfectly competitive, first best economy a benevolent government would always reveal its private information. But, in a second best world with unavoidable distortions, a benevolent government might hope to increase social welfare by appropriately distorting information communication. For instance, suppose that monopoly power or income taxes make labor supply sub-optimal. If the government knows that the economy is hitting a recession and does not reveal such information, it may hope that the increase in labor supply caused by ignorance compensates the under-supply of labor provoked by the existence of distortions.

If individuals mechanically believe its announcements (if they are credulous), the government may even be able to restore the first best outcome through an appropriately over-optimistic communication strategy. However, if individuals are rational, this misleading information about a recession will make the government lose credibility. In particular, when the economy is hitting a boom and the government announces it, individuals will discount such announcement. This, in turn, will further worsen the under-supply of labor in booms, and thus reduce social welfare in good times.

[1]See Romer and Romer (2000); Kurz (2005); Kohn and Sack (2004); Athey et al. (2005).
[2]See Oh and Waldman (1990); Rodríguez Mora and Schulstald (2007).

In recessions, by hiding information, the government raises labor supply, relative to what it would be under perfect information, so that it may (at least partially) compensate for the welfare loss caused by the existing distortion. Yet, it may also raise labor supply so much, that it indeed causes an over-supply of labor (relative to the first best), whose welfare costs are higher than those due to the distortion under perfect information. The higher the distortion, the less likely it is that this happens. Thus, roughly, high levels of distortion would induce the government to hide negative information.

We start by making our essential point in an abstract model with distortions, in which a benevolent social planner has private information on productivity shocks and sends payoff-irrelevant messages to uninformed individuals. We characterize the equilibria of this cheap talk game. We find that non informative equilibria always exist, whereas an informative (separating) equilibrium exists if and only if distortions are sufficiently small. This analysis extends the standard cheap-talk game (Crawford and Sobel, 1982) to a context where there is a continuum of heterogeneous receivers, each with a continuum of actions.[3]

Transparency allows individuals to react to different information in different states of the world. Opaqueness in turn makes their actions more stable, since different states of the world belong to the same information set. We show that transparency is Pareto optimal, from an ex ante point of view, if and only if full-information social welfare satisfies an appropriate convexity condition in productivity shocks. If this condition is satisfied, our results imply that, even with a benevolent social planner, opaqueness may be an equilibrium although transparency is ex ante desirable.

We next extend the abstract model and focus on two specific sources of distortions: monopoly power and income taxation. This extension shows the direct relevance of the mechanism we analyze in standard macroeconomic and public economics models.

In the monopoly power model, we again show that a non informative equilibrium always exists, whereas existence of an informative equilibrium requires that the monopolistic distortion is sufficiently small (in which case transparency is the most natural prediction). The model shows that an

[3]Farrell and Gibbons (1989) extend Crawford and Sobel (1982) to two audiences. They show that in what they call a 'coherent game', the sender prefers separating to pooling, ex post and therefore also ex ante. This is not true in our model. In this sense, our analysis also yields some insights of technical interest for game theorists, since it shows that this result critically hinges on the two-action assumption and doe not generalize.

increase in average productivity harms transparency, whereas an increase in shock magnitude favors transparency, at least when shocks are small. Thus, ceteris paribus, countries with more competitive product markets and larger shocks are more likely to have a truthful government, whereas there is no presumption that economic development per se brings about transparency.

In the taxation model, results on existence, efficiency and equilibrium selection parallel those of the monopoly power model. Again, an increase in shock magnitude favors transparency, at least when shocks are small. Hence, ceteris paribus, countries with lower taxation and higher shocks are more likely to have a truthful government. The main novelty in the context of income taxation is that it is a natural environment to study the effects of labor income inequality on transparency. We show that such effects depend on labor supply elasticity. If labor supply is rigid, an increase in inequality favors transparency and the opposite if it is elastic (only with linear labor supply inequality does not affect transparency).

Both in the monopoly power and in the taxation model, transparency is ex ante desirable (since the convexity condition of social welfare in shocks is satisfied), but may not be feasible in equilibrium, because even a benevolent government may want to hide negative information.

Our results are first of all related to the recent literature on the welfare effects of information. In an influential paper, Morris and Shin (2002) show that noisy public information, if used to coordinate actions, may lead individuals to disregard alternative valuable private information, so that more precise public information may reduce welfare.[4] Angeletos and Pavan (2007) clarify that this result, and more generally the welfare effects of public information in an abstract class of linear-quadratic games with heterogeneous information, depend on whether the equilibrium degree of coordination is inefficiently high or low. While we abstract from both heterogeneous information and from strategic complementarity or substitutability (we simply assume that the equilibrium degree of coordination is zero), this is done in order to focus on a different question, which has remained unexplored in this

[4]This reasoning has been used to warn against Central Bank transparency (Amato et al., 2002), a conclusion that has been disputed in a lively recent debate, developed both in the abstract context of 'beauty contest' models (Svensson, 2006; Morris et al., 2006; James and Lawler, 2011) and in the applied context of New Keynesian models (Woodford, 2003; Hellwig, 2005; Roca, 2010). Amador and Weill (2010) emphasize that public information may slow social learning, as it jeopardizes the price system's ability to aggregate and transmit private information, and that it may thus result in welfare losses.

literature, namely whether a benevolent social planner has ex post incentives to reveal its information, even when it would be ex ante efficient to do so. The main novelty of our contribution is therefore to show that, in economies in which transparency is ex ante optimal, it may not arise in equilibrium because, when distortions are large, even a benevolent social planner would find revealing bad news ex post sub-optimal.

The importance of transparency for economic policy, and in particular for monetary policy, has long been recognized and has been a recurrent theme of theoretical, empirical and policy debates. Besides the virtues of central bank's transparency, its limits have also received substantial attention. Two prominent examples are Cukierman and Meltzer (1986) and Stein (1989), who consider central bank's private information on its own policy preferences and highlight the advantages of ambiguity and imprecise communication (the latter within a cheap talk game).[5] We differ from this line of research in that we emphasize real rather than monetary channels and assume private information about the macroeconomic outlook rather than about policy goals. Perhaps more importantly, we abstract from information precision to focus on the role of economic distortions in determining whether equilibrium communication is transparent or not.

The two extensions of the abstract model are also related to the specific literature on information in models of business cycle and of public and political economy. We share the focus on information about real rather than monetary shocks in a business cycle context with Angeletos and La'O (2009), but their analysis concerns the role of information dispersion. More closely related to our work is Angeletos et al. (2011), who show that, in the context of a DSGE model with dispersed information, more precise public information raises welfare if the business cycle is driven by technology or preference shocks (as in the present paper), but not necessarily if it is driven by shocks to monopoly markups or to labor wedges. Gavazza and Lizzeri (2009, 2011) show that transparency may generate economic distortions (wasteful spending of inefficient public debt) if voters are misinformed about government spending and revenues. We tackle the complementary question and show how transparency is endogenously determined by pre-existing distortions.

The remainder of this paper is organized as follows. Sections 2, 3 and 4

[5]See Cukierman (2009) and Blinder et al. (2008) for a recent assessment of the limits to central bank's transparency and for a review of the literature on central bank's communication.

display the abstract model and the two extensions to monopoly power and to income taxation, respectively. Section 5 presents a concluding discussion. The Online Appendix contains technical results for the two applied models.

2 Transparency and distortions

In this section we shall abstract from the origin of the distortion in the allocation of resources. The distortion will be considered exogenous and affecting individual payoffs via individual choices. In the following sections we shall extend the analysis of two specific cases of distortion. This will permit to examine the link between the origin of the distortion and the government's information policy.

2.1 The economy

Consider an economy with a mass one of identical individuals, who make simultaneous choices. Individual $i \in [0,1]$ chooses an action $\ell_i \geq 0$ and obtains payoff $u(\ell_i, L, \lambda, \theta)$, where $L = \int_0^1 \ell_i \, \mathrm{d}i$ is the average (or aggregate) action in the population, $\lambda \geq 0$ is a parameter capturing distortions and θ is a random variable (the state of the world), which affects the productivity of individual actions. Using subscript numbers to denote a function's partial derivatives, we make the following assumptions on $u(\ell_i, L, \lambda, \theta)$, which are assumed to hold on the entire domain, unless otherwise specified.

1. The individual problem has an interior maximum:
   $u_{11} < 0$ and $\forall L, \lambda, \theta,\ \exists \ell > 0:\ u_1(\ell, L, \lambda, \theta) = 0$.

2. There are positive externalities: $u_2 \geq 0$, with equality only for $\lambda = 0$.

3. There is strategic independence: $u_{12} = 0$.

4. The social planner's problem has an interior maximum:
   $u_{11} + u_{22} < 0$ and $\forall \lambda, \theta,\ \exists \ell > 0:\ u_1(\ell, \ell, \lambda, \theta) = 0$.

5. λ strengthens externalities and reduces actions, but does not affect the social optimum: $u_{23} = -u_{13} > 0$.

6. The exogenous state θ is beneficial through own actions and externalities: $u_4(\ell, L, \lambda, \theta) > 0$ if either $\ell > 0$ or $L > 0$ and $\lambda > 0$.

7. θ boosts individual actions and makes externalities stronger: $u_{14} > 0$ and $u_{24} \geq 0$, with equality only for $\lambda = 0$.

8. Given actions (and λ), θ has a linear effect on the payoff: $u_{44} = 0$.

Assumptions 1 and 4 are straightforward. Assumption 2 introduces distortions in the form of positive externalities. Assumption 3 rules out strategic complementarity, which is the mechanism emphasized by Morris and Shin (2002) and others. Assumption 5 makes λ a measure of the strength of externalities, but it makes the (Benthamite) social optimum independent of λ. This property makes sense under Assumption 3, for instance in models in which externalities arise from distribution rather than from production.[6] In such cases, stronger externalities raise the under-provision of effort (relative to the social optimum) at a decentralized solution, but do not affect the socially efficient allocation of effort. Assumption 6 conditions the beneficial effect of the exogenous state θ to the presence of a positive level of either own actions or externalities from other people's actions. A reasonable interpretation, which is consistent with all the models developed in this paper, is that θ represents an aggregate productivity shock. Under this interpretation, Assumption 7 becomes natural, as it states that a higher productivity stimulates individual activity and also reinforces the externalities from other people's activity. Finally, Assumption 8 makes the direct effect of θ (given actions and λ) linear. This is consistent with a linear technology with additive productivity shocks, as in the models in the next sections, and it simplifies the analysis.

Under perfect information on θ, each individual i would choose $\ell^*(\lambda, \theta) = \text{argmax}_\ell\, u(\ell, L, \lambda, \theta)$, determined by the first order condition $u_1(\ell, L, \lambda, \theta) = 0$ and satisfying $\ell^*_\lambda < 0$ and $\ell^*_\theta > 0$ and $\ell^*_L = 0$.[7] A benevolent social planner with a utilitarian welfare function would choose for each individual the same action $\hat{\ell}(\theta)$ by solving $\max_{\{\ell_i : i \in [0,1]\}} \int_0^1 u(\ell_i, L, \lambda, \theta)\,\mathrm{d}i$, which yields the FOC system $u_1(\ell_i, L, \lambda, \theta) + \int_0^1 u_2(\ell_j, L, \lambda, \theta)\,\mathrm{d}j = 0$ for all $i \in [0, 1]$. Given $u_{12} = 0$,

[6] Two examples are the standard models considered in the next sections, in which externalities arise from the distribution of either profits or tax revenues. In models with strategic complementarity, by contrast, Assumption 5 would not generally hold. To avoid future confusion, we anticipate that in the model of Section 3 Assumption 3 holds as referred to labor supply choices, which are the only ones made under uncertainty. The existence of strategic complementarity in price setting in that model does not play a relevant role.

[7] We use subscript letters instead of numbers to denote partial derivatives whenever this facilitates reading and does not create confusion.

this can be written as $u_1(\ell_i, L, \lambda, \theta) + u_2(\ell_i, L, \lambda, \theta) = 0$ for all $i \in [0,1]$, which has a symmetric solution solving $u_1(\ell, \ell, \lambda, \theta) + u_2(\ell, \ell, \lambda, \theta) = 0$, and satisfying $\hat{\ell}_\theta > 0$ and $\hat{\ell}_\lambda = 0$. Notice that for any $\lambda > 0$, $\ell^*(\lambda, \theta) < \ell^*(0, \theta) = \hat{\ell}(\theta)$, so that λ introduces a downward distortion in actions relative to the social optimum.

2.2 The Announcements Game

We investigate what happens when the social planner knows the realization of θ, whereas individuals are not perfectly informed about it, and have to decide on the basis of beliefs, which in turn may be influenced by the planner's announcements. Specifically, we assume that information is as follows. First Nature draws θ from the following distribution, which is common knowledge:

$$\theta = \begin{cases} \vartheta & , \text{ with probability } p \\ -\vartheta & , \text{ with probability } (1-p) \end{cases}, \tag{1}$$

with $\vartheta > 0$ and $p \in (0,1)$.[8] The planner observes the realization of θ and then chooses a (payoff irrelevant) message m from a set of feasible messages $\{M, N\}$. Individuals observe m, but not θ, and then simultaneously choose their actions to maximise expected payoff.

We consider a signaling equilibrium of this cheap talk game, with the additional but natural requirement that out of equilibrium beliefs are the same for everybody. Thus, a pure strategy *equilibrium* consists of: (i) a message function $m(\theta)$ mapping realizations of the random shock into messages, such that the planner's objective (ex-post social welfare) is maximized, given individual strategies; (ii) posterior beliefs $\Pr(\vartheta|m)$, which map each message into a subjective probability about the realization of the random variable, and which are derived from messaging strategies through Bayes' rule along the equilibrium path of play (and are the same for everybody following out-of-equilibrium messages); and (iii) individual strategies $s(m)$ mapping messages into actions, which maximize individual expected payoff, given posterior beliefs (and other individuals' strategies).[9]

[8] In this and in the following sections, the parameter ϑ can be interpreted as the amplitude of the cycle, which is assumed symmetric for analytical simplicity.

[9] The mixed strategy extension is immediate.

2.3 Equilibrium with generic distortions

Some notation will be useful throughout the paper. Let $\mu = \Pr(\theta = \vartheta | m = M)$ and $\nu = \Pr(\theta = \vartheta | m = N)$ describe individual posterior beliefs (for which we also use the notation $\Pr(\vartheta | m)$, for $m = M, N$). Let $E(x(\theta)|m) = \Pr(\vartheta|m)x(\vartheta) + [1 - \Pr(\vartheta|m)]x(-\vartheta)$ denote the expected value of a generic function $x(\theta)$, when expectations are based on posterior beliefs after receiving message $m = M, N$. And let $\bar{x}(\theta) = px(\vartheta) + (1-p)x(-\vartheta)$ denote its ex ante expected value, when expectations are based on prior beliefs.[10] Finally, let $\lambda^*(\mu, \nu)$ be the solution by λ of $u(\ell^*(M), \ell^*(M), \lambda, -\vartheta) = u(\ell^*(N), \ell^*(N), \lambda, -\vartheta)$, defined for $\mu \neq \nu$, where $\ell^*(m)$ is each individual's best response to message $m = M, N$.

The next proposition characterizes pure strategy equilibria (mixed strategy equilibria are characterized in footnote 12 and are not very insightful).

Proposition 1 (Equilibrium with generic distortions)
An equilibrium in pure strategies always exists. Given μ and ν, individual strategies are $\ell^*(m) = \operatorname{argmax}_\ell E(u(\ell, L, \lambda, \theta)|m)$, for $m = M, N$. There are two possible types of pure strategy equilibrium.

- *At a pooling equilibrium $m^*(\vartheta) = m^*(-\vartheta) = N$ and $\mu \leq \nu = p$. A pooling equilibrium always exists.*
- *At a separating equilibrium $m^*(-\vartheta) = M$, $m^*(\vartheta) = N$, $\mu = 0$ and $\nu = 1$. A separating equilibrium exists if and only if $\lambda \leq \lambda^*(0, 1)$.*

Proof The proof consists of three steps: (i) given posterior beliefs, we determine individual best responses to the planner's messages, $\ell^*(m)$; (ii) we determine the planner's best response to individual strategies in each state of the world, $m^*(\theta)$; (iii) we impose that posterior beliefs are obtained through Bayes' rule along the equilibrium path of play.

1. Given μ and ν, $\ell^*(M)$ is the solution by ℓ of $\mu u_1(\ell, L, \lambda, \vartheta) + (1 - \mu)u_1(\ell, L, \lambda, -\vartheta) = 0$ and analogously, with ν in place of μ, for $\ell^*(N)$. Notice that L is immaterial to individual choices. Individual strategies $\ell^*(m)$ satisfy $\ell^*_\lambda(m) < 0$, for $m = M, N$, and $\ell^*_\mu(M) > 0$ and $\ell^*_\nu(N) > 0$.

[10] So expected utility after $m = M$ is $E(u(\ell, L, \lambda, \theta)|M) = \mu u(\ell, L, \lambda, \vartheta) + (1 - \mu)u(\ell, L, \lambda, -\vartheta)$ and analogously for $E(u|N)$.

2. $m^*(\theta) = \text{argmax}_{m\in\{M,N\}}\, u(\ell^*(m), \ell^*(m), \lambda, \theta)$. The planner is indifferent (and can thus randomize with any probability) between the two messages if $\mu = \nu$. If $\mu \neq \nu$, consider without loss of generality the case of $\mu < \nu$.

   In the good state, $m^*(\vartheta) = N$. This is due to $\ell^*(M) < \ell^*(N) \leq \hat{\ell}(\vartheta)$, with the last inequality strict for $\lambda > 0$, and to $u_1(\ell, L, \lambda, \vartheta) > 0$ for $\ell < \hat{\ell}(\vartheta)$.

   In the bad state, for any μ and ν such that $0 \leq \mu < \nu \leq 1$, there exists a unique $\lambda^*(\mu, \nu) > 0$, such that $m^*(-\vartheta) = M$ if $\lambda < \lambda^*(\mu, \nu)$, $m^*(-\vartheta) = N$ if $\lambda > \lambda^*(\mu, \nu)$ and the planner is indifferent between M and N if $\lambda = \lambda^*(\mu, \nu)$. To see this, let $W(\ell, \lambda, -\vartheta) = u(\ell, \ell, \lambda, -\vartheta)$ and notice that it is a continuous and inverted-U shaped function of ℓ and that its point of maximum, $\hat{\ell}(-\vartheta)$, is independent of λ. The planner compares $W(\ell^*(M), \lambda, -\vartheta)$ with $W(\ell^*(N), \lambda, -\vartheta)$. Since $\ell^*(m)$ is strictly increasing in the posterior belief $\Pr(\vartheta|m)$, we have that $\ell^*(m) \in [\ell^0(\lambda), \ell^1(\lambda)]$, where $\ell^0(\lambda) = \ell^*(m)$ for $\Pr(\vartheta|m) = 0$ and $\ell^1(\lambda) = \ell^*(m)$ for $\Pr(\vartheta|m) = 1$, for $m = M, N$. Moreover, since $\ell^*(m)$ is strictly decreasing in λ, the same is true for $\ell^0(\lambda)$ and $\ell^1(\lambda)$. For $\lambda = 0$, we have $\ell^0(0) = \hat{\ell}(-\vartheta)$, so that $\forall \mu, \nu :\ 0 \leq \mu < \nu \leq 1,\ W(\ell^*(M), 0, -\vartheta) > W(\ell^*(N), 0, -\vartheta)$. For λ large enough, we have that eventually $\ell^1(\lambda) \leq \hat{\ell}(-\vartheta)$ (to see this, notice that $u_1(\ell^1(\lambda), \ell^1(\lambda), \lambda, -\vartheta) + u_2(\ell^1(\lambda), \ell^1(\lambda), \lambda, -\vartheta) > 0$, because both terms are strictly positive), so that $\forall \mu, \nu :\ 0 \leq \mu < \nu \leq 1,\ W(\ell^*(M), \lambda, -\vartheta) < W(\ell^*(N), \lambda, -\vartheta)$. The result is then proven by observing that, given any $\mu, \nu :\ 0 \leq \mu < \nu \leq 1$, we have that for any finite λ, $\ell^*(M) < \ell^*(N)$, and that both $\ell^*(M)$ and $\ell^*(N)$ decrease continuously in λ, passing from being both above $\hat{\ell}(-\vartheta)$ (at least weakly for $\ell^*(M)$) when $\lambda = 0$ to being both below it (at least weakly for $\ell^*(N)$) when λ is sufficiently large. Therefore there exists a unique $\lambda^*(\mu, \nu)$, such that for $\lambda = \lambda^*(\mu, \nu)$, $\ell^*(M) < \hat{\ell}(-\vartheta) < \ell^*(N)$ and $W(\ell^*(M), \lambda, -\vartheta) = W(\ell^*(N), \lambda, -\vartheta)$. And we have that $\lambda^*(\mu, \nu) > 0$; that $m^*(-\vartheta) = M$ if $\lambda < \lambda^*(\mu, \nu)$; that $m^*(-\vartheta) = N$ if $\lambda > \lambda^*(\mu, \nu)$; and that the planner is indifferent between M and N if $\lambda = \lambda^*(\mu, \nu)$.

3. Consider a candidate pooling equilibrium. The planner sends the same message N in both states of the world. Along the equilibrium path of play, i.e., upon receiving N, individuals do not learn anything and have to base decisions on their prior beliefs: Bayes' rule implies $\nu = p$. Then

by the previous point, the planner does not deviate in the good state if and only if $\mu \leq p$. The pooling equilibrium is babbling if $\mu = \nu$ and non babbling if $\mu < \nu$.[11] If $\mu < p$, the planner does not deviate in the bad state if and only if $\lambda \geq \lambda^*(\mu, p)$. In turn, if $\mu = \nu$, the planner never deviates in the bad state. So a babbling equilibrium always exists.

Now consider a candidate separating equilibrium. The planner announces N in the good state and M in the bad state. Bayes' rule then implies $\mu = 0$ and $\nu = 1$. Given this, the planner never deviates in the good state. It does not deviate in the bad state either, if and only if $\lambda \leq \lambda^*(0, 1)$.[12]

■

Proposition 1 shows that there always exists an equilibrium in which the planner is non informative. Furthermore, a transparent equilibrium in which the planner reveals its private information also exists if and only if distortions are sufficiently small. If distortions are large, and thus individual actions are seriously downward distorted with respect to the social optimum, a benevolent planner has a strong incentive to hide bad news, and this disrupts the possibility that in equilibrium it is transparent.

[11] A babbling equilibrium is an equilibrium in which individual strategies disregard the planner's announcement, and the planner's signaling strategy disregards the realization of the shock. The only difference between babbling and non babbling pooling equilibria is in terms of out of equilibrium beliefs.

[12] There are only two possible types of mixed strategy equilibria: (i) babbling equilibria, in which the planner randomizes with any $\Pr(m(-\vartheta) = N) = \Pr(m(\vartheta) = N) = \rho \in (0, 1)$ and $\mu = \nu = p$; and (ii) semi-separating equilibria, in which $m(\vartheta) = N$ and the planner randomizes in bad times with some $\Pr(m(-\vartheta) = N) = \rho \in (0, 1)$, with posterior beliefs $\mu = 0$ and $\nu = \frac{p}{p+(1-p)\rho}$. Mixed strategies babbling equilibria always exist. A semi-separating equilibrium exists if and only if $\lambda = \lambda^*\left(0, \frac{p}{p+(1-p)\rho}\right)$. To see this, notice that in the good state the planner is willing to mix if and only if $\mu = \nu$, which is only compatible with babbling equilibria. So, without loss of generality, non babbling equilibria in mixed strategies imply $\mu < \nu$ and $m(\vartheta) = N$, i.e., they may only be semi-separating. For $\rho \in \{0, 1\}$, we have the two pure strategy equilibria considered above. For $\rho \in (0, 1)$, Bayes' rule implies $\mu = 0$ and $\nu = \frac{p}{p+(1-p)\rho}$. Given this, the planner does not deviate in the good state. It does not deviate in the bad state either, if and only if $\lambda = \lambda^*\left(0, \frac{p}{p+(1-p)\rho}\right)$.

2.4 Efficiency

Let us now compare the pooling and the separating equilibria from an ex ante point of view, that is, when averages (or expected values) are based on the prior distribution of shocks. Let $\bar{u}^S$ and $\bar{u}^P$ denote the ex ante expected levels of social welfare (equivalently, of individual payoff), at a separating and at a pooling equilibrium, respectively.

Proposition 2 (Ex ante Pareto dominance)
The separating equilibrium ex ante Pareto dominates the pooling equilibrium ($\bar{u}^S > \bar{u}^P$), if and only if equilibrium payoff under perfect information is a convex function of the random variable θ.

Proof Let $u^S(\theta)$ and $u^P(\theta)$ be social welfare at a separating and at a pooling equilibrium, respectively, when the state of the world is θ. Let $\ell^S(\theta)$ be the action chosen under perfect information about the state of the world (as it happens at a separating equilibrium), which solves by ℓ $u_1(\ell, L, \lambda, \theta) = 0$. Let ℓ^P be the action chosen at a pooling equilibrium. Linearity of the payoff in θ ($u_4 > 0$ and $u_{44} = 0$) implies linearity of the marginal payoff in θ ($u_{14} > 0$ and $u_{144} = 0$) and so it implies that $\ell^P = \ell^S(\bar{\theta})$, since it solves $pu_1(\ell, L, \lambda, \vartheta) + (1-p)u_1(\ell, L, \lambda, -\vartheta) = u_1(\ell, L, \lambda, \bar{\theta}) = 0$, where $\bar{\theta} = p\vartheta + (1-p)(-\vartheta)$. It also implies that $\bar{u}^P = u^S(\bar{\theta})$. To see this, observe that $u^S(\bar{\theta}) = u(\ell^S(\bar{\theta}), \ell^S(\bar{\theta}), \lambda, \bar{\theta}) = pu(\ell^P, \ell^P, \lambda, \vartheta) + (1-p)u(\ell^P, \ell^P, \lambda, -\vartheta)$. Consider now the function $W(\theta) = u(\ell^S(\theta), \ell^S(\theta), \lambda, \theta)$. We have that $\bar{u}^S = pW(\vartheta) + (1-p)W(-\vartheta)$ and $\bar{u}^P = W(\bar{\theta})$, so that $\bar{u}^S > \bar{u}^P \iff W''(\theta) > 0$. ■

Proposition 2 implies that a benevolent planner ex-ante sees transparency as preferable to opaqueness if and only if the shocks, if publicly observed, would have a convex effect on the equilibrium payoff. In this case, opaqueness would induce actions responding to the expected values of the random variable yielding a lower payoff because of the convexity condition. Because of the same argument, concavity would make opaqueness ex ante preferable.

2.5 Implications of the main results

We have assumed an economy subject to a distortion and that experiences random shocks, over which the planner has private information. The planner decides on its information policy in view to maximize social welfare. Notice that the most preferred policy before knowing the realization of the random shock might be different from the one preferred after knowing it. We assume that the government cannot credibly commit and hence chooses the policy that maximizes ex-post social welfare. Our results say that if the distortion is large enough, the only equilibrium policy entails opaqueness because the planner always finds it ex-post preferable to hide negative shocks. The consequence of this is that individuals in equilibrium fully distrust the planner's announcements. If the distortion is small, truth-telling can be an equilibrium, together with opaqueness.[13] However, transparency will be ex ante seen as preferable to opaqueness if and only if, under full information, the individual payoff is convex in the random variable. If this condition is satisfied, the planner would ex ante prefer to be transparent, but, if distortions are substantial in magnitude, will be opaque in equilibrium. In this case, delegating information policy to a separate agency, committed to transparency, would be beneficial, because such commitment would prevent from making ex post a decision that ex ante is sub-optimal.

In the next sections we examine two different potential sources of distortions and apply our general results to obtain the corresponding equilibrium information policy. This allows to show both how the basic mechanism works in standard economic contexts and what new insights it provides.

3 Transparency and monopoly power

In the first extension of the abstract model we consider an economy with a monopolistic distortion. The model is a simplified version of the canonical RBC model with no capital and a continuum of differentiated goods. To focus on the information analysis, we abstract from dynamic considerations.[14]

3.1 The economy and the announcements game

There is a mass one of identical individuals, who work, consume and own shares of a mass one of firms, each producing a different variety of a consumption good. Utility depends on consumption and labor:

[13] Note that this result has a clear "second best" flavor. Hiding information is in itself a distortion, relative to full information. However, given the presence of other distortions in the economy, it may well be that this additional distortion ends up increasing aggregate welfare.

[14] See, e.g., Angeletos and La'O (2009). Relative to their analysis, we also abstract from dispersed information (as we did in our abstract model).

$$u(c,\ell) = c - \frac{\ell^\delta}{\delta}, \tag{2}$$

where c is the Dixit-Stiglitz aggregator $c = \left(\int_0^1 c_i^{\frac{\sigma-1}{\sigma}}\,\mathrm{d}i\right)^{\frac{\sigma}{\sigma-1}}$, c_i represents consumption of variety i, the parameter $\sigma > 1$ is the elasticity of substitution between any two varieties and the parameter $\delta > 1$ captures the degree of convexity of labor supply, which is linear in the wage for $\delta = 2$, strictly convex for $\delta \in (1,2)$ and strictly concave for $\delta > 2$. Firms produce with an identical linear technology: by using ℓ_i units of labor, firm i produces $y_i = A\ell_i$ units of its variety of good. Labor productivity is $A = \tilde{A} + \theta$, so it is the same for every firm, but it depends on two factors: the observable component $\tilde{A} > 0$ and the ex-ante unobservable component θ (say, being in a boom or in recession), distributed according to (1). We assume $\vartheta \in (0, \tilde{A})$ to assure that productivity is always positive.[15]

Under perfect information on θ, the equilibrium of this economy is very simple. Individuals choose $\{c_i : \; i \in [0,1]\}$ and ℓ to maximise $u(c,\ell)$ under the budget constraint $\int_0^1 p_i c_i \,\mathrm{d}i = w\ell + \pi$, taking the wage rate w, prices p_i's and distributed profits π as given.[16] Their labor supply is $\ell = \left(\frac{w}{P}\right)^{\frac{1}{\delta-1}}$, where $P = \left(\int_0^1 p_i^{1-\sigma} di\right)^{\frac{1}{1-\sigma}}$ is the aggregate price index. Their demand of good i is $c_i = \left(\frac{p_i}{P}\right)^{-\sigma} c$.

Each firm i sets price p_i to maximise profits $\pi_i = p_i y_i - w\ell_i$, taking technology, demand, other firms' prices and the wage rate as given. It thus prices according to the mark-up rule $p_i = \frac{\sigma}{\sigma-1}\frac{w}{A}$ and demands labor $\ell_i = \frac{c_i}{A}$. Since prices are the same for every firm, the same holds for quantities: $\forall i,\; p_i = P$ and $c_i = c$. The real wage $\frac{w}{P} = \frac{\sigma-1}{\sigma}A$ is below labor productivity. Labor supply is then $\ell = \left(\frac{\sigma-1}{\sigma}A\right)^{\frac{1}{\delta-1}}$, consumption is $c = \left(\frac{\sigma-1}{\sigma}\right)^{\frac{1}{\delta-1}} A^{\frac{\delta}{\delta-1}}$, real profits are $\frac{\pi}{P} = \frac{c}{\sigma}$ and equilibrium utility is $u(c,\ell) = \left(\frac{\sigma-1}{\sigma}\right)^{\frac{1}{\delta-1}} A^{\frac{\delta}{\delta-1}} \left(1 - \frac{\sigma-1}{\delta\sigma}\right)$.

[15] It is immediate to extend the model to the case in which firm productivity is heterogeneous. We present the identical firms version for expositional simplicity, as it is sufficient to convey the main insights.

[16] Here we assume that individuals are identical both in productivity and in shareholding. We discuss the role of heterogeneous productivity in the next section. Heterogeneity in shareholding would make no relevant changes in the present model, since it would affect the distribution of income but not individual behavior, as individuals take distributed profits as given.

Taking c as the numeraire, so that $P = 1$, the nominal part of the economy is also easily determined.

Three observations are in order. First, monopoly power drives a wedge between real wage and productivity, imposing a suboptimal downward distortion in individual labor supply, relative to the (first best) social optimum, which would require $\ell = A^{\frac{1}{\delta-1}}$. Second, this distortion is higher, the lower the elasticity of substitution σ. Indeed, profit distribution creates a positive externality from labor supply, which is not taken into account by individual choices, and (real) profits are decreasing in σ. Third, equilibrium utility is convex in A and therefore in θ. This last observation makes planner's transparency Pareto-superior to opaqueness from an ex-ante point of view, since, as it is easy to check, this model adds economic structure to the specification of actions and their relation to utility, but it satisfies all the assumptions of the previous one.

Consider now imperfect information, with the following structure. First, Nature draws θ from distribution (1), which is common knowledge. Second, the planner observes the realization of θ and then chooses a (payoff irrelevant) message $m \in \{M, N\}$ to maximize (ex post) social welfare. Third, the labor market opens, firms demand labor, workers supply labor, and the (expected) wage adjusts to clear the market. Firms and workers contract on a state-contingent real wage, but employment decisions are based on expectations, since at this stage they know m but not yet θ. Fourth, once employment is sunk, the fundamental is publicly revealed, production is realized, and commodity prices adjust so as to clear the commodity markets (thus, prices p_i's, consumption choices c_i's, the wage rate w, and profits π are all determined under full information). Ex post social welfare is $W(\theta, m) = \int_0^1 u(c(\theta, m, \ell^*(m)), \ell^*(m))\, \mathrm{d}i$, where $c(\theta, m, \ell^*(m))$ and $\ell^*(m)$ are the equilibrium values of consumption and labor supply.

Notice that "employment" is here a proxy for all kinds of input and production choices that are made before the perfect realization of aggregate uncertainty. One can thus think of this also as a proxy for investment. The key point is that in this formulation real economic decisions are made on the basis of incomplete information about the state of the economy. This is different from the congenital way employment is modeled in the standard New-Keynesian model, where it is assumed, for simplicity but not for realism, that all employment is free to adjust to the true realized state.

3.2 Equilibrium with monopoly power

The following proposition parallels Proposition 1. Let

$$\begin{aligned} x_\mu &= E(A|M)^{\frac{1}{\delta-1}} = [\tilde{A} + (2\mu-1)\vartheta]^{\frac{1}{\delta-1}} \\ x_\nu &= E(A|N)^{\frac{1}{\delta-1}} = [\tilde{A} + (2\nu-1)\vartheta]^{\frac{1}{\delta-1}} \end{aligned}$$

and, for $\mu \neq \nu$, let

$$\sigma^*(\mu,\nu) = \frac{x_\nu^\delta - x_\mu^\delta}{x_\nu^\delta - x_\mu^\delta - \delta(x_\nu - x_\mu)(\tilde{A} - \vartheta)}. \qquad (3)$$

Proposition 3 (Equilibrium with monopoly power)
Given μ and ν, individuals' strategies are
$\ell^*(m) = E\left(\frac{w}{P}|m\right)^{\frac{1}{\delta-1}} = \left\{\left(\frac{\sigma-1}{\sigma}\right)\left[\tilde{A} + E(\theta|m)\right]\right\}^{\frac{1}{\delta-1}}$ *and*
$c_i(\theta,m,\ell) = \left[\frac{p_i(\theta,m,\ell)}{P(\theta,m,\ell)}\right]^{-\sigma} c(\theta,m,\ell)$.
Firms' strategies are $p_i(\theta,m,\ell) = \frac{\sigma}{\sigma-1}\frac{w(\theta,m,\ell)}{\tilde{A}+\theta}$.
There are two possible types of pure strategy equilibrium.[17]

- *At a pooling equilibrium* $m(\vartheta) = m(-\vartheta) = N$ *and* $\mu \leq \nu = p$. *A pooling equilibrium always exists.*

- *At a separating equilibrium* $m(-\vartheta) = M,\ m(\vartheta) = N,\ \mu = 0$ *and* $\nu = 1$. *A separating equilibrium exists if and only if* $\sigma \geq \sigma^*(0,1)$.

Proof See the Online Appendix. ■

3.3 Efficiency and equilibrium selection

Let us now compare the different equilibria from an ex ante point of view. Let $\bar{\ell}^S$, $\bar{y}^S$, $\bar{u}^S$, and $\bar{\ell}^P$, $\bar{y}^P$, $\bar{u}^P$, denote the ex ante expected levels of labor supply, production and indirect utility, at a separating and at a pooling equilibrium, respectively. Independently of equilibrium existence, the following holds.

[17] The structure of mixed strategy equilibria is analogous to that of the abstract model and is not reported.

Proposition 4 (Ex ante Pareto dominance)
For any parameter constellation, the following holds: (i) $\bar{\ell}^S < \bar{\ell}^P \iff \delta > 2$; *(ii)* $\bar{y}^S > \bar{y}^P$; *(iii)* $\bar{u}^S > \bar{u}^P$.

Proof See Online Appendix. ■

Proposition 4 establishes that, for any degree of monopoly power, transparent and credible revelation of information is ex ante Pareto superior to information hiding. This is not surprising in light of Proposition 2, since, as noted above, full-information equilibrium utility is convex in θ.[18] Yet, as we know from Proposition 3, high monopoly power may prevent the transparent outcome from materializing in equilibrium.[19] The intuition is very simple. Transparency allows individuals to work more when they are more productive and less when they are less productive. This unequivocally raises the ex ante level of production (and welfare), relative to no information disclosure. Although it also raises the ex ante level of disutility from labor (since workers dislike fluctuations in labor effort), this latter effect is always more than compensated by the higher expected level of consumption (hence the effect on welfare).[20]

As in the abstract model, when distortions are not too large, i.e., here, for $\sigma \geq \sigma^*(0,1)$, both a separating and a pooling equilibrium exist. It is then natural to ask which equilibrium is more plausible in this case.

[18] This is also in line with Angeletos et al. (2011), since technology shocks are the source of the business cycle. Notice that, although we assume that the source of distortions (the elasticity of substitution in the present model as well as the tax rate in the next section) is a-cyclical, this is not true for the wedge between first-best and full-information labor supply, which is pro-cyclical in both models. Angeletos and Pavan (2007) show that, in economies that are inefficient even under complete information, if this wedge positively co-varies with full-information equilibrium strategies, then more precise public information is (ex ante) welfare increasing. Our results are coherent with theirs.

[19] Notice that, for any parameter constellation, ex ante expected levels of individual labor supply, production and indirect utility, whose relationships are identified in Proposition 4, are well defined, independently of whether a separating equilibrium exists.

[20] Transparency raises expected leisure time if the elasticity of labor supply is $\gamma \equiv \frac{1}{\delta - 1} < 1$ (i.e., for $\delta > 2$). In this case, labor supply is a concave function of expected wages. This implies that, relative to the case of no information, labor supply reductions in recessions are more pronounced than increases in booms. If the elasticity of labor supply is $\gamma > 1$ (i.e., for $\delta < 2$), by contrast, labor supply is a convex function of expected wages. In this case, transparency raises expected labor supply, relative to information hiding.

Ex ante Pareto dominance selects the separating equilibrium whenever it exists, but it is not (always) a good selection criterion in the present context, because, whenever $\sigma \in [\sigma^*(0,1), \sigma^*(0,p)]$, planner's preferences over equilibria are reversed in different states of the world.[21] It is therefore worthwhile to look at different equilibrium refinements.

In cheap talk games, standard refinements based on Kohlberg and Mertens (1986), which restrict off-the-equilibrium-path beliefs, have little power, because mixed strategy babbling equilibria always survive them. We consider a recent refinement, explicitly introduced by Chen et al. (2008) to select equilibria in cheap-talk games, called No Incentive to Separate (NITS); and a stronger refinement, the Neologism Proof (NP) equilibrium, proposed by Farrell (1993). In the Online Appendix we define such concepts in the context of our model and show that, whenever it exists, the separating equilibrium satisfies both NITS and NP. Moreover, whenever existing, it is the only NP equilibrium.

The implication is that, for $\sigma \geq \sigma^*(0,1)$, the separating equilibrium appears the most natural prediction of the game. In light of this, we conduct the following discussion assuming that the economy coordinates on the transparent equilibrium whenever it exists.

3.4 Transparency and the business cycle

One natural question is whether the amplitude of the business cycle favors or harms transparency.[22] To answer this question in the simplest way, it is convenient to focus on the case of linear labor supply ($\delta = 2$). In this case, expression (3) simplifies to $\sigma^*(0,1) = \frac{\tilde{A}}{\vartheta}$. If we measure the relative amplitude of the business cycle by the ratio of shock size to structural productivity, $\frac{\vartheta}{\tilde{A}}$, it is clear that economies subject to more pronounced business cycles have a larger support for transparency. While this may be surprising, from the

[21] In booms the planner would prefer to be in a separating equilibrium, in which it reveals its private information, thus boosting labor supply and welfare; in recessions it would prefer to be in a pooling equilibrium, in which information is not revealed, so that labor supply and welfare are higher than with perfect information. The proof of this claim immediately follows from the proof of Lemma 2 in the Online Appendix, whereas the fact that $\sigma^*(0,p) > \sigma^*(0,1)$ follows from Lemma 1 in the Online Appendix.

[22] We refer to the interval $[\sigma^*(0,1), \infty)$ as to the support of transparency and say that parameter changes favor (reduce) transparency if they decrease (increase) $\sigma^*(0,1)$. Recall from Proposition 3 that information transparency is an equilibrium policy (indeed, the most natural prediction of the game) for $\sigma \geq \sigma^*(0,1)$ —i.e., when the distortion is "small".

theoretical point of view it is a direct consequence of the way in which the support for transparency is determined, namely by the welfare comparison that a credible planner makes between revealing or not information on a recession. When $\sigma = \sigma^*(0,1)$, a credible planner is indifferent between the two alternatives. This is because the monopolistic distortion (the underwork caused by monopoly power under truthful revelation) is equivalent (in welfare terms) to the information distortion (the overwork arising if the planner lies). An increase in relative shock magnitude raises the information distortion relative to the monopolistic distortion, and therefore raises the relative cost of lying and expands the support for transparency. The full generalization of this result to arbitrary parameter values proves complex, but for small shocks it is easy to show that, for any parameter constellation (not only for $\delta = 2$), an increase in shock magnitude favors transparency.[23] While shock magnitude matters for transparency, shock frequency does not, because the threshold $\sigma^*(0,1)$ is determined conditionally on being in recession.

It is also interesting to notice that our model features a fixed elasticity of labor supply to expected wages, equal to $\gamma = \frac{1}{\delta-1}$. Yet, while actual wages do not depend on whether equilibrium information is transparent or not, expected wages, and thus actual labor supply, do depend on the information regime. In particular, they fluctuate over the business cycle under transparency but not under opaqueness. As a result, the elasticity of labor supply to actual wages over the business cycle is zero for an economy with large monopolistic distortions and therefore opaque information, whereas it is positive and equal to γ for an economy with the same value of δ, but with lower monopoly power and therefore transparent information. As a consequence, all else equal, the model predicts that income fluctuations over the business cycle will also be more pronounced in economies with more competitive product markets and therefore transparent rather than opaque information. The importance of information and expectations is an aspect that tends to be overlooked in the empirical debate on the estimates of labor supply elasticity, which in our view deserves more attention.

[23] To prove this result, write $\sigma^*(0,1) = \frac{1}{1-q(0,1)}$ as in Lemma 1 in the Online Appendix and observe that $\frac{\partial \sigma^*(0,1)}{\partial \vartheta}$ is equal in sign to $\frac{\partial q(0,1)}{\partial \vartheta}$. By applying L'Hôpital's rule it is easy to obtain that $\lim_{\vartheta \to 0} \frac{\partial q(0,1)}{\partial \vartheta} = -\frac{1}{\bar{A}} < 0$, so that $\left.\frac{\partial \sigma^*(0,1)}{\partial \vartheta}\right|_{\vartheta=0} < 0$.

4 Transparency, taxation and inequality

As a second extension of the abstract model, consider now an economy in which distortions arise from taxation rather than from monopoly power. Income taxes push net wages below individual productivity and thus make labor effort inefficiently low, as individuals do not internalize the externality emerging from the redistribution of tax revenues, exactly as they did not internalize the externality arising from the distribution of firms' profits in the previous model.

4.1 The economy and the announcements game

There is a mass one of individuals, who have identical preferences over a homogeneous consumption good c and over labor effort ℓ, described by (2), but who differ in productivity. Individuals earn competitive wages and produce with a linear technology, so that labor income y (equivalently, production, taken as numeraire) is simply equal to individual supply of efficiency units of labor. Individual productivity depends on two factors: an idiosyncratic observable component (ability or human capital), denoted β and distributed according to the cumulative distribution function F, with support on the non empty interval $[b, B) \subset \mathbb{R}_+$; and the aggregate, ex-ante unobservable, random component θ (being in boom or in recession), distributed according to (1), with $\vartheta \in (0, b)$ to assure that individual productivity is always positive.

Individual labor income therefore depends on effort, ability and aggregate conditions, $y_\beta = (\beta + \theta)\ell_\beta$. Labor income is taxed at a constant marginal rate $t \in (0, 1)$ and tax revenues $T = \int_b^B t y_\beta dF(\beta)$ are equally redistributed, so that individual consumption is equal to $c_\beta = (1 - t)y_\beta + T$. Since the population is continuous, each individual takes T as given.[24]

From our assumption on preferences it is immediate to obtain that, if individuals could observe the realization of θ before choosing their effort level, they would choose $\ell_\beta = [(1-t)(\beta+\theta)]^{\frac{1}{\delta-1}}$ and produce $y_\beta = (1 - t)^{\frac{1}{\delta-1}}(\beta + \theta)^{\frac{\delta}{\delta-1}}$. Taxes impose a downward distortion in individual effort supply, relative to the social optimum, which would require $\ell_\beta = (\beta + \theta)^{\frac{1}{\delta-1}}$.

[24] The tax collection per capita T will depend on the realization of θ. Therefore, individuals will entertain conjectures about their value. As we shall see, because of our assumption on individual preferences these conjectures are immaterial because they have no effect on labor supply.

Equilibrium social welfare under perfect information is $W = \int_b^B u_\beta dF(\beta) = \left(\frac{\delta-1+t}{\delta}\right)(1-t)^{\frac{1}{\delta-1}} \int_b^B (\beta+\theta)^{\frac{\delta}{\delta-1}} dF(\beta)$, which is convex in θ. Therefore, we can expect transparency to be Pareto-superior to opaqueness from an ex-ante point of view.

Consider now imperfect information. First Nature draws θ from (1). Both F and the distribution of θ are common knowledge. The planner observes the realization of θ and then chooses a (payoff irrelevant) message $m \in \{M, N\}$. Individuals observe m, but not θ, and then simultaneously choose their labor effort to maximize utility. Ex-post the realization of θ is observed by all individuals, who are paid accordingly. The aim of the planner is to maximise social welfare $W = \int_b^B u_\beta dF(\beta)$, where u_β denotes the utility of an individual with ability β and depends on t, on θ, on individual labor effort ℓ_β, and on the labor effort chosen by the entire population (since T depends on it). The equilibrium concept and the notation on beliefs and expectations are as above.

4.2 Equilibrium and efficiency

The main results on equilibrium and efficiency parallel those obtained above, so we present them without discussion. In particular, Proposition 5 parallels Propositions 1 and 3 and Proposition 6 follows Propositions 2 and 4. With a slight abuse of notation, but in the same spirit as in Section 3, let

$$\begin{aligned} x_\mu &= [\beta + E(\theta|M)]^{\frac{1}{\delta-1}} = [\beta + (2\mu-1)\vartheta]^{\frac{1}{\delta-1}}, \\ x_\nu &= [\beta + E(\theta|N)]^{\frac{1}{\delta-1}} = [\beta + (2\nu-1)\vartheta]^{\frac{1}{\delta-1}} \end{aligned}$$

and, for $\mu \neq \nu$,

$$t^*(\mu,\nu) = 1 - \frac{\int_b^B (\beta-\vartheta)(x_\nu - x_\mu)\,\mathrm{d}F(\beta)}{\int_b^B \frac{1}{\delta}(x_\nu^\delta - x_\mu^\delta)\,\mathrm{d}F(\beta)} \tag{4}$$

Proposition 5 (Equilibrium with taxation and inequality) *Given μ and ν, equilibrium labor supply strategies are described by $\ell^*_\beta(m) = \{(1-t)[\beta + E(\theta|m)]\}^{\frac{1}{\delta-1}}$.*

There are two possible types of pure strategy equilibrium.

- *At a pooling equilibrium $m(\vartheta) = m(-\vartheta) = N$ and $\mu \leq \nu = p$. A pooling equilibrium always exists.*

- *At a separating equilibrium $m(-\vartheta) = M$, $m(\vartheta) = N$, $\mu = 0$ and $\nu = 1$. A separating equilibrium exists if and only if $t \leq t^*(0,1)$.*

Proof See Online Appendix. ■

Let us now compare the different equilibria from an ex ante point of view. For an individual with ability β, let $\bar{\ell}^S_\beta$, $\bar{y}^S_\beta$, $\bar{u}^S_\beta$, and $\bar{\ell}^P_\beta$, $\bar{y}^P_\beta$, $\bar{u}^P_\beta$, denote the ex ante expected levels of labor supply, production and indirect utility, at a separating and at a pooling equilibrium, respectively. Independently of equilibrium existence, the following results hold.

Proposition 6 (Ex ante Pareto dominance)
For any parameter constellation, any ability distribution, and any level β of individual ability, the following holds: (i) $\bar{\ell}^S_\beta < \bar{\ell}^P_\beta \iff \delta > 2$; (ii) $\bar{y}^S_\beta > \bar{y}^P_\beta$; (iii) $\bar{u}^S_\beta > \bar{u}^P_\beta$.

Proof See Online Appendix. ■

Results for equilibrium selection are also very close to those obtained in the context of the monopoly power model and are discussed in the Online Appendix. Again, the main insight is that, whenever existing, i.e., for $t \leq t^*(0,1)$, the separating equilibrium appears the most natural prediction of the game. In light of this, for the following discussion we assume that the economy coordinates on the separating equilibrium whenever it exists.

4.3 Transparency and inequality

The main new insight provided by this model concerns the role of inequality and its effects on transparency.[25] Indeed, this is a natural environment to ask

[25] The effects of the amplitude of the business cycle are similar to those discussed for the previous model and are analyzed in the Online Appendix. Moreover, the taxation model also confirms that, all else equal, output and hours worked fluctuate more when the government is transparent. This is consistent with the evidence provided by Demertzis and Hughes-Hallett (2007) and with the recently uncovered negative relationship between taxation and output volatility (Debrun et al., 2008). Yet, since the government tends to be transparent when aggregate shocks are relatively large, the ceteris paribus condition should not be forgotten.

whether inequality favors or harms transparency.[26] In the next proposition we use Lorenz dominance (second order stochastic dominance) as a criterion to establish whether a distribution has more inequality than another one. Let $\gamma \equiv \frac{1}{\delta-1}$ be the elasticity of labor supply.

Proposition 7 (Effects of inequality)
For any parameter constellation and distributional assumption, the effects of skill inequality on transparency depend on labor supply elasticity. In particular, consider a shift from skill distribution F to a more unequal distribution G, dominated by F with respect to second order stochastic dominance.

- *If $\gamma = 1$, such an increase in inequality has no effects on transparency.*
- *If $\gamma < 1$, it favors transparency.*
- *If $\gamma > 1$, letting $\hat{\gamma} = \frac{2}{1-t^*(0,1)} > 2$, we have that $\gamma \in (1, \hat{\gamma}]$ is a sufficient condition for it to reduce transparency.*

Proof See Online Appendix. ■

First notice that most of the literature on information transparency assumes $\gamma = 1$ (which means linear labor supply) and thus assumes away the effects of inequality. Yet, the general picture is that inequality matters for transparency and that the way it does depends on the shape of the labor supply curve. In particular, if labor supply is rigid (i.e., for $\gamma < 1$), as most micro-estimates suggest, inequality favors transparency. Yet, if labor supply is elastic, as many macro models assume, inequality harms transparency.[27]

To grasp the intuition of this result, notice that $t^*(0,1)$ depends on the welfare comparison between (credibly) revealing and not revealing information, conditional on being in a recession.[28] Relative to information hiding, transparency in recessions raises leisure and reduces consumption for each individual. It is therefore useful to disentangle the effects of inequality on

[26] Analogously to what we did in Section 3, we refer to the tax rate interval $[0, t^*(0,1)]$ as to the support of transparency, and say that parameter or distributional changes favor (reduce) transparency if they increase (decrease) $t^*(0,1)$.

[27] For evidence on labor supply elasticity see, e.g., Fiorito and Zanella (2012).

[28] The relative social welfare gain to transparency in recessions depends on the tax rate: it is positive for low tax distortions and negative in the opposite case. The formal details on such comparison are provided in footnote 7 in the Online Appendix.

$t^*(0,1)$ into those working through the consumption differential and those working through the leisure differential between transparency and opaqueness.

Given the complementarity between skills and effort, an increase in skill inequality raises mean income and therefore mean consumption, independently of labor supply elasticity.[29] Since higher consumption is what drives the welfare advantage to opaqueness over transparency, by raising aggregate consumption, inequality favors opaqueness. Yet, utility depends on leisure besides on consumption, and the effects of inequality on leisure are more interesting.

In our model labor supply is concave (in wage or ability) whenever it is rigid ($\gamma < 1$). With rigid labor supply, an increase in skill inequality therefore raises aggregate leisure time. Since higher leisure is what drives the welfare advantage to transparency over opaqueness, by raising aggregate leisure, inequality favors transparency. Thus, under rigid labor supply, the two welfare effects of inequality, through consumption and through leisure, work in opposite directions: one favors opaqueness and the other one transparency. The overall effect depends on which force dominates. What we show is that, with rigid labor supply, the leisure channel dominates and skill inequality indeed favors transparency.

In contrast, with elastic (and therefore convex in ability) labor supply ($\gamma > 1$), an increase in skill inequality raises aggregate labor time and thus, besides increasing aggregate consumption, it reduces aggregate leisure. Both effects now work in the same direction, making skill inequality favor opaqueness.

To the extent that inequality generates higher tax rates and therefore higher distortions, it may also have an indirect effect on transparency. To explore this mechanism, in Albornoz et al. (2009) we provide a political economy extension of the analysis, in which taxes are chosen by majority voting along the lines of Meltzer and Richard (1981).[30] The impact of inequality then results from the combination of two effects. On the one hand, higher

[29]While at first sight surprising, the fact that skill inequality is welfare increasing is a direct consequence of the above mentioned complementarity, paired with a Benthamite social welfare function.

[30]We look at a politico-economic Nash equilibrium, in which the tax rate (chosen by majority voting) and the informational policy (chosen by the government) are mutually consistent. Notice that, given the assumptions of the model, if a benevolent government were completely free to set the tax rate, it would set it equal to zero.

inequality induces the median voter to choose a higher tax rate. This in itself reduces the scope for transparency. On the other hand, for any given tax rate, the way in which higher inequality changes the government's valuation of truth-telling (in recessions) depends on the elasticity of labor supply, as discussed in Proposition 7. If labor supply is elastic, an increase in inequality amplifies the distortion created by taxation and thus raises the incentive to hide bad news. In this case, the two channels move in the same direction and inequality harms transparency. However, if labor supply is inelastic, an increase in inequality reduces the magnitude of the tax distortion and thus raises the incentive to transparently reveal bad news. Now the two channels work in opposite directions and the net effect cannot be established in general.[31]

5 Concluding discussion

This paper investigates how government transparency depends on economic distortions. Distortions drive a wedge between the social optimum and the full-information equilibrium. As a consequence, a benevolent government, with welfare-relevant private information, has an incentive to manipulate communication.

If distortions are high, transparency cannot emerge in equilibrium, even when it is ex ante desirable. If distortions are also hard to remove, the policy implication is that the government should find some commitment device to transparency. For instance, announcements over the economic outlook might be delegated to an independent statistical office, committed to transparency.

Our results suggest that, all else equal, we should expect a negative relationship between government transparency and economic distortions. We are not aware of any empirical investigation of the impact of distortions on government transparency. Yet, in the cross-section of countries, there is a strong negative correlation between measures of fiscal transparency and measures of distortions.[32] While this correlation may reflect the joint effect of

[31] Focusing on the case of a unit elasticity of labor supply, we show that, if the size of the shock is larger than a threshold level, the equilibrium is informative for all levels of inequality; but below this threshold, inequality harms transparency. Specifically, the equilibrium is informative for low inequality, uninformative for high inequality, whereas both an informative and an uninformative equilibrium exist for intermediate levels.

[32] For instance, this holds when fiscal transparency is measured by the "Open Budget Index"

political institutions on these variables, our theory suggests that causality may also run from distortions directly to the level of transparency. An open direction for future research is to explore this possibility in the data.

Our theory highlights the limits of equilibrium transparency when the government is benevolent, individuals are rational and no credible commitment is possible. We leave the analysis of transparency outside these assumptions for future investigation.[33] Within our framework, it is worth noting that precisely when the government 'lies' (in the sense that, in recessions, it sends the same message it sends in booms), individuals are ex post happy that it 'lied'. Therefore, the fact that the government's private information is ex post verifiable is not problematic. Moreover, the fact that we restrict to two elements both the state and message space polarizes equilibria on either full revelation or no revelation at all. An extension to the continuum case would generate equilibria with partial revelation and would thus allow to study the degree of information precision, but it would not affect the main intuition and the main results.[34]

Perhaps the most interesting extensions of the present framework concern the various possible forms of interaction between economic distortions and transparency. First, while we assumed that the driver of the business cycle, on which the government has private information, is a productivity shock, one might also imagine that the government's private information concerns shocks to monopoly mark-ups or to labor wedges, as in Angeletos et al. (2011). This might change our results and make opaqueness ex ante desirable. Second, we have assumed that distortions are persistent and the government cannot directly eliminate them. The investigation of the political economy reasons behind this difficulty is a promising research avenue. For instance, an elected government might be influenced by lobbying activity or by a demand for redistribution. In Albornoz et al. (2009) we explore this last

(developed by the International Budget Partnership), and distortions by the "Time to Start Business" or the "Ease of Doing Business" (which capture barriers to entry and obstacles to business activity, both provided by the World Bank), or by several measures of taxes (such as top marginal tax rate on labor income, taxes on goods and taxes on international trade, obtained by the OECD World Tax Indicators).

[33] For instance, an incumbent government might want to be over-optimistic in order to influence individuals' beliefs on its ability, beyond the motives emphasized in this paper. While this would provide an extra incentive to hide bad news, we expect that it would not change our main results.

[34] This can be seen most clearly in Albornoz et al. (2009), where mixed strategies allow for a semi-separating equilibrium.

possibility and show that inequality harms transparency because it generates higher taxes. This modifies the results obtained in section 4 above, where we show how the effects of inequality on transparency depend on labor supply elasticity. Third, we have assumed that the government influences individual choices only through its informational policy, but how the latter interacts with monetary and fiscal policy is certainly worth investigating, since the direction and size of the shock may well depend on policy actions. Fourth, while we have assumed that the government perfectly observes the shock and that individuals have no other source of information, a natural extension would be to look at government incentives to put a privately observed noisy signal of the shock in the public domain, when individuals also have dispersed and noisy information. These lines of research remain open.

References

Albornoz, F., J. Esteban, and P. Vanin (2009). Government information transparency. UFAE and IAE Working Papers 774.09, Unitat de Fonaments de l'Anàlisi Econòmica (UAB) and Institut d'Anàlisi Econòmica (CSIC).

Amador, M. and P.-O. Weill (2010). Learning from prices: Public communication and welfare. *Journal of Political Economy 118*(5), 866 – 907.

Amato, J. D., S. Morris, and H. S. Shin (2002). Communication and monetary policy. *Oxford Review of Economic Policy 18*(4), 495–503.

Angeletos, G.-M., L. Iovino, and J. La'O (2011). Cycles, gaps, and the social value of information. NBER Working Papers 17229.

Angeletos, G.-M. and J. La'O (2009). Noisy business cycles. *NBER Macroeconomics Annual 24*, 318–378.

Angeletos, G.-M. and A. Pavan (2007). Efficient use of information and social value of information. *Econometrica 75*(4), 1103–1142.

Athey, S., A. Atkeson, and P. J. Kehoe (2005). The optimal degree of discretion in monetary policy. *Econometrica 73*(5), 1431–1475.

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. D. Haan, and D.-J. Jansen (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature 46*(4), 910–45.

Chen, Y., N. Kartik, and J. Sobel (2008). Selecting cheap-talk equilibria. *Econometrica 76*(1), 117–136.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica 50*(6), 1431–51.

Cukierman, A. (2009). The limits of transparency. *Economic Notes 38*(1-2), 1–37.

Cukierman, A. and A. H. Meltzer (1986). A theory of ambiguity, credibility, and inflation under discretion and asymmetric information. *Econometrica 54*(5), 1099–1128.

Debrun, X., J. Pisani-Ferry, and A. Sapir (2008). Government size and output volatility: Should we forsake automatic stabilization? *IMF Working Paper* (08/122).

Demertzis, M. and A. Hughes-Hallett (2007). Central bank transparency in theory and practice. *Journal of Macroeconomics 29*(4), 760–789.

Farrell, J. (1993). Meaning and credibility in cheap-talk games. *Games and Economic Behavior 5*(4), 514–531.

Farrell, J. and R. Gibbons (1989). Cheap talk with two audiences. *American Economic Review 79*(5), 1214–1223.

Fiorito, R. and G. Zanella (2012). The anatomy of the aggregate labor supply elasticity. *Review of Economic Dynamics 15*(2), 171 – 187.

Gavazza, A. and A. Lizzeri (2009). Transparency and economic policy. *Review of Economic Studies 76*(3), 1023–1048.

Gavazza, A. and A. Lizzeri (2011). Transparency and manipulation of public accounts. *Journal of Public Economic Theory 13*(3), 327–349.

Hellwig, C. (2005). Heterogeneous information and the benefits of public information disclosures. Mimeo, UCLA Department of Economics.

James, J. G. and P. Lawler (2011). Optimal policy intervention and the social value of public information. *American Economic Review 101*(4), 1561–74.

Kohlberg, E. and J.-F. Mertens (1986). On the strategic stability of equilibria. *Econometrica 54*(5), 1003–37.

Kohn, D. L. and B. Sack (2004). *Central Bank Talk: Does it Matter and Why?*, Volume Macroeconomics, Monetary Policy and Financial Stability. Ottawa: Bank of Canada.

Kurz, M. (2005). Diverse beliefs, forecast errors and central bank policy. Stanford University.

Meltzer, A. H. and S. F. Richard (1981). A rational theory of the size of government. *Journal of Political Economy 89*(5), 914–927.

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

Morris, S., H. S. Shin, and H. Tong (2006). Social value of public information: Morris and shin (2002) is actually pro-transparency, not con: Reply. *American Economic Review 96*(1), 453–455.

Oh, S. and M. Waldman (1990). The macroeconomic effects of false announcements. *Quarterly Journal of Economics 105*(4), 1017–34.

Roca, M. (2010). Transparency and monetary policy with imperfect common knowledge. IMF Working Papers 10/91.

Rodríguez Mora, J. and P. Schulstald (2007). The effect of GNP announcements on fluctuations of GNP growth. *European Economic Review 51*(8), 1922–1940.

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

Stein, J. C. (1989). Cheap talk and the Fed: A theory of imprecise policy announcements. *The American Economic Review 79*(1), 32–42.

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*(1), 448–452.

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*, pp. 25–58. Princeton, NJ: Princeton University Press.