MARKET DISTORTIONS AND GOVERNMENT TRANSPARENCY

Facundo Albornoz
Birmingham University

Joan Esteban
Instituto de Análisis Económico (CSIC) and Barcelona Graduate School of Economics

Paolo Vanin
Università di Bologna

Abstract
In this paper, we investigate 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 cannot constitute an equilibrium when 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: D82, E61)

1. Introduction

Many governments are better informed than the private sector about future realizations of macroeconomic variables (see Romer and Romer 2000; Kohn and Sack 2004; Athey, Atkeson, and Kehoe 2005; Kurz 2005). 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, which confirms that individuals find them informative (see Oh and Waldman 1990; Rodríguez Mora and Schulstald 2007). In contrast, the

The editor in charge of this paper was George-Marios Angeletos.

Acknowledgments: For useful comments, we thank Gaetano Antinolfi, Francesca Barigozzi, Jordi Caballé, Antonio Cabrales, Hector Calvo Pardo, Giacomo Calzolari, Matteo Cervellati, Russell Cooper, Nicola Pavoni, Debraj Ray, and Joel Sobel, as well as seminar and conference participants in Bologna, Barcelona (EEA-ESEM), Buenos Aires (LACEA-LAMES), Istanbul (ASSET), and Lucca. 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 under the title "Government Information Transparency".

E-mail: f.albornoz@bham.ac.uk (Albornoz); joan.esteban@iae.csic.es (Esteban); paolo.vanin@unibo.it (Vanin)

Journal of the European Economic Association February 2014 12(1):200–222

DOI: 10.1111/jeea.12052

widespread skepticism attending contemporary Argentine or Greek official statistics exemplifies nontransparent and noncredible government announcements. There are likely to be various opportunistic reasons for governments not to be transparent, but here 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 induce 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 and first-best economy, a benevolent government would always reveal its private information. Yet, in a second-best world with unavoidable distortions, a benevolent government might hope to increase social welfare by appropriately distorting its communication of information. For instance, suppose that monopoly power or income taxes render labor supply suboptimal. If the government knows the economy is heading for a recession but does not reveal this information, it may hope that the increase in labor supply as a result of ignorance will compensate for the undersupply of labor caused by market distortions.

If individuals mechanically believe the government's announcements—that is, if they are credulous—then the government may even be able to achieve the first-best outcome through an appropriately overoptimistic communication strategy. However, if individuals are rational, misleading information about a recession will reduce the government's credibility. Individuals will then discount government's announcements that the economy is approaching a boom phase, even when truthful. This discounting will, in turn, further worsen the undersupply of labor in booms and thus reduce social welfare in good times.

In recessions, by concealing information the government is able to increase the labor supply—relative to its level under perfect information—and thereby to compensate (at least partially) for the welfare loss caused by market distortions. In doing so, however, the government may raise labor supply so much that it causes oversupply of labor (relative to the first-best case), the welfare costs of which may even be higher than those resulting from market distortions under perfect information. The greater the market distortions, the less likely it is that this outcome occurs. Roughly, then, high levels of distortions should induce the government to conceal 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 noninformative 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 in which there is a continuum of heterogeneous receivers, each with a continuum of actions.[1]

1. Farrell and Gibbons (1989) have extended Crawford and Sobel (1982) to two audiences. They have shown that, in what they call a coherent game, the sender prefers separating to pooling ex post and therefore

Transparency allows individuals to react to different information in different states of the world. Opaqueness, in contrast, makes individuals' actions more stable because 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, then our results imply that, even with a benevolent social planner, opaqueness may be an equilibrium despite transparency being more desirable ex ante.

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

In the monopoly power model, we again show that a noninformative equilibrium always exists, whereas an informative equilibrium exists if and only if the monopolistic distortion is sufficiently small (in which case, transparency is the most natural prediction). The model shows that an increase in average productivity hinders the emergence of a transparent equilibrium, 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; however, there is no presumption that economic development per se brings about transparency.

In the taxation model, results concerning existence, efficiency, and equilibrium selection parallel those of the monopoly power model. Again, an increase in shock magnitude favors transparency when shocks are small, and so, all else equal, countries with lower taxation and higher shocks are more likely to have a truthful government. The main novelty regarding income taxation is that it is a natural environment in which to study the effects of labor income inequality on transparency. We show that such effects depend on the elasticity of labor supply. If labor supply is rigid, then an increase in labor income inequality favors transparency; if labor supply is elastic, such an increase favors opaqueness. Only with linear labor supply does inequality have no effect on whether or not transparency emerges in equilibrium.

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

Our results are related first of all to the body of literature on the welfare effects of information. In an influential paper, Morris and Shin (2002) have shown that noisy public information, if used to coordinate actions, may lead individuals to disregard alternative valuable private information; hence, public information that is

also ex ante. This ex-post preference does not hold in our model. Thus, our analysis yields some insights of technical interest to game theorists by showing that the result of Farrell and Gibbons hinges critically on its two-action assumption and does not generalize.

more precise may actually reduce welfare.[2] Angeletos and Pavan (2007) have clarified that this result—and, more generally, the welfare effects of public information in an abstract class of linear-quadratic games with heterogeneous information—depends on whether the equilibrium degree of coordination is inefficiently high or low. We abstract from heterogeneous information and also from strategic complementarity or substitutability (we simply assume that the equilibrium degree of coordination is zero). Our motive for adopting this approach is to focus on a question heretofore unexplored in this body of literature: Does a benevolent social planner have ex-post incentives to conceal information even when revealing it is ex ante efficient? Thus, the novel aspect of our contribution is that, in economies where 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 to be ex post suboptimal.

The importance of transparency for economic policy, and in particular for monetary policy, has long been recognized and is a recurrent theme in theoretical, empirical, and policy debates. Not only the virtues but also the limits of central bank transparency have received substantial attention. Two prominent examples are Cukierman and Meltzer (1986) and Stein (1989), who have considered a central bank's private information on its own policy preferences and have highlighted the advantages of ambiguity and imprecise communication (the latter within a cheap-talk game).[3] 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 or not equilibrium communication is transparent.

Our two extensions (monopoly power and taxation) of the abstract model are also related to the specific body of literature on information in models of the business cycle and of public and political economy. We share with Angeletos and La'O (2009) a focus on information about real rather than monetary shocks in a business cycle context, but their analysis concerns the role of information dispersion. More closely related to our work is that of Angeletos, Iovino, and La'O (2011), who have shown that—in the context of a dynamic stochastic general equilibrium model with dispersed information—more precise public information increases welfare if the business cycle is driven by technology shocks (as in our paper) or preference shocks, but not necessarily if it is driven by shocks to monopoly markups or to labor market distortions. Gavazza and Lizzeri (2009, 2011) have shown that transparency may generate

2. This reasoning has been used to warn against central bank transparency (Amato, Morris, and Shin 2002). This topic is disputed in a lively debate that has developed both in the abstract context of "beauty contest" models (Morris, Shin, and Tong 2006; Svensson 2006; James and Lawler 2011) and in the applied context of New Keynesian models (Woodford 2003; Hellwig 2005; Roca 2010). Amador and Weill (2010) have argued that public information may slow social learning by jeopardizing the price system's ability to aggregate and transmit private information, which may result in welfare losses.

3. See Blinder et al. (2008) and Cukierman (2009) for a review of the literature on central bank communication and for an assessment of the limits to central bank transparency.

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 rest of this paper is organized as follows. In Sections 2, 3, and 4 respectively, we describe the abstract model and the two extensions to monopoly power and to income taxation. In Section 5, we present a concluding discussion. The Online Appendix contains technical results for the two applied models.

2. Transparency and Distortions

In this section, we abstract from the origin of the distortion in the allocation of resources. Distortion is here considered to be exogenous and to affect individual payoffs via individual choices. In later sections, we extend the analysis to specific types of distortion, enabling us to examine the link between market distortion's source and the government's information policy.

2.1. The Economy

Consider an economy with a mass 1 of identical individuals, who make simultaneous choices. Individual $i \in [0, 1]$ chooses an action $\ell_i \geq 0$ and obtains payoff $u(\ell_i, L, \lambda, \theta)$; here, $L = \int_0^1 \ell_i \, \mathrm{d}i$ 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) that 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—unless stated otherwise—apply throughout the entire domain.

1. The individual problem has an interior maximum:
   $u_{11} < 0$ and, for all L, λ, θ, there exists $\ell > 0$ such that $u_1(\ell, L, \lambda, \theta) = 0$.
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, for all λ, θ, there exists $\ell > 0$ such that $u_1(\ell, \ell, \lambda, \theta) = 0$.
5. The distortion-capturing parameter λ strengthens externalities and reduces actions, but it 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 $\ell > 0$ or if both $L > 0$ and $\lambda > 0$.
7. The random variable 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 renders the (Benthamite) social optimum independent of λ. This property makes sense under Assumption 3—for instance, in models where externalities arise from distribution rather than from production.[4] In such cases, stronger externalities amplify the underprovision 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 on 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 is intuitive because it means 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 featuring additive productivity shocks (as in our subsequent models) and simplifies the analysis.

Under perfect information on θ, each individual i would choose $\ell^*(\lambda,\theta) = \text{argmax}_\ell u(\ell, L, \lambda, \theta)$, which is determined by the first-order condition $u_1(\ell, L, \lambda, \theta) = 0$ and satisfies $\ell^*_\lambda < 0$, $\ell^*_\theta > 0$, and $\ell^*_L = 0$.[5] 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 first-order conditions $u_1(\ell_i, L, \lambda, \theta) + \int_0^1 u_2(\ell_j, L, \lambda, \theta)\,\mathrm{d}j = 0$ for all $i \in [0,1]$. Given $u_{12} = 0$, such conditions 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 that satisfies $u_1(\ell, \ell, \lambda, \theta) + u_2(\ell, \ell, \lambda, \theta) = 0$ as well as $\hat{\ell}_\theta > 0$ and $\hat{\ell}_\lambda = 0$. Observe that, for any $\lambda > 0$, we have $\ell^*(\lambda,\theta) < \ell^*(0,\theta) = \hat{\ell}(\theta)$, and thus a positive λ distorts individual actions (effort) downward relative to the social optimum.

2.2. *The Announcements Game*

We investigate what happens when only the social planner knows the realization of θ, whereas individuals are not perfectly informed and must decide on the basis of beliefs, which in turn may be influenced by the planner's announcements. Specifically, we assume that information is disseminated as follows. First, θ is drawn by nature from

4. Two examples are the standard models considered in Sections 3 and 4, where externalities arise from the distribution of either profits or tax revenues. In models with strategic complementarity, however, Assumption 5 would not hold in general. We avoid future confusion by anticipating that in the model of Section 3, Assumption 3 holds with respect to labor supply choices (which are the only decisions made under uncertainty). In that model, no relevant role is played by strategic complementarities in price setting.

5. We use subscript letters (instead of numbers) to denote partial derivatives whenever doing so facilitates reading and does not create confusion.

the following distribution, which is common knowledge:

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

Here, $\vartheta > 0$ and $p \in (0,1)$.[6] Then, the planner observes this realization of θ and chooses a (payoff-irrelevant) message m from a set $\{M, N\}$ of feasible messages. Next, individuals observe m, but not θ, and then simultaneously choose actions that maximize their expected payoffs.

We consider a signaling equilibrium of this cheap-talk game that incorporates the additional but natural requirement that out-of-equilibrium beliefs be the same for everyone. Thus, a pure strategy equilibrium consists of: (a) a message function $m(\theta)$ mapping realizations of the random shock into messages, such that the planner's objective (i.e., ex-post social welfare) is maximized, given individual strategies; (b) posterior beliefs $\Pr(\vartheta|m)$, which map each message into a subjective probability about the realization of the random variable and are derived from messaging strategies through Bayes' rule along the equilibrium path of play (and are the same for everyone following out-of-equilibrium messages); and (c) individual strategies $s(m)$ mapping messages into actions that maximize each individual's expected payoff, given posterior beliefs (and other individuals' strategies).[7]

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 $m = M, N$; we let $\bar{x}(\theta) = px(\vartheta) + (1-p)x(-\vartheta)$ denote that function's ex-ante expected value when expectations are based on prior beliefs.[8] Finally, let $\lambda^*(\mu, \nu)$ be the solution by λ of $u(\ell^*(M), \ell^*(M), \lambda, -\vartheta) = u(\ell^*(N), \ell^*(N), \lambda, -\vartheta)$, which is defined for $\mu \neq \nu$ and where $\ell^*(m)$ is each individual's best response to message $m = M, N$.

The following proposition characterizes pure strategy equilibria (mixed strategy equilibria are characterized in footnote 11 and are not especially illuminating).

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 equilibria.*

6. The parameter ϑ can be interpreted as the amplitude of the cycle, which is assumed (for analytical simplicity) to be symmetric.

7. The mixed strategy extension is immediate.

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

- *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 proceeds in three steps. First, given posterior beliefs, we determine individual best responses $\ell^*(m)$ to the social planner's messages. We then determine the planner's best response to individual strategies $m^*(\theta)$ in each state of the world. Finally, we impose that posterior beliefs are obtained through Bayes' rule along the equilibrium path of play.

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

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

In the good state, $m^*(\vartheta) = N$. This is because $\ell^*(M) < \ell^*(N) \leq \hat{\ell}(\vartheta)$, where the last inequality is strict for $\lambda > 0$, and because $u_1(\ell, L, \lambda, \vartheta) > 0$ 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 observe that this is a continuous and inverted–U-shaped function of ℓ and that its maximum point, $\hat{\ell}(-\vartheta)$, is independent of λ. The planner compares $W(\ell^*(M), \lambda, -\vartheta)$ with $W(\ell^*(N), \lambda, -\vartheta)$. Because $\ell^*(m)$ is strictly increasing in the posterior belief $\Pr(\vartheta|m)$, it follows 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 $\Pr(\vartheta|m) = 1$, with $m = M, N$. Moreover, because $\ell^*(m)$ is strictly decreasing in λ, the same is true for $\ell^0(\lambda)$ and $\ell^1(\lambda)$. For $\lambda = 0$, we have $\ell^0(0) = \hat{\ell}(-\vartheta)$ and so, for all μ and ν such that $0 \leq \mu < \nu \leq 1$, $W(\ell^*(M), 0, -\vartheta) > W(\ell^*(N), 0, -\vartheta)$. For λ large enough, eventually $\ell^1(\lambda) \leq \hat{\ell}(-\vartheta)$. (This follows because $u_1(\ell^1(\lambda), \ell^1(\lambda), \lambda, -\vartheta) + u_2(\ell^1(\lambda), \ell^1(\lambda), \lambda, -\vartheta) > 0$, because both terms are strictly positive.) Then, for all μ and ν such that $0 \leq \mu < \nu \leq 1$, $W(\ell^*(M), \lambda, -\vartheta) < W(\ell^*(N), \lambda, -\vartheta)$.

The result can now be proved by observing that, given any μ and ν with $0 \leq \mu < \nu \leq 1$, we have $\ell^*(M) < \ell^*(N)$ for any finite λ 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. Hence, there exists a unique $\lambda^*(\mu, \nu)$ such that, for $\lambda = \lambda^*(\mu, \nu)$, we have $\ell^*(M) < \hat{\ell}(-\vartheta) < \ell^*(N)$ and $W(\ell^*(M), \lambda, -\vartheta) = W(\ell^*(N), \lambda, -\vartheta)$. Observe that $\lambda^*(\mu, \nu) > 0$. Moreover, $m^*(-\vartheta) = M$ if $\lambda < \lambda^*(\mu, \nu)$, $m^*(-\vartheta) = N$ if $\lambda > \lambda^*(\mu, \nu)$, and the planner is indifferent between M and N if $\lambda = \lambda^*(\mu, \nu)$.

Step 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 so must base their decisions on prior beliefs. Bayes' rule implies $\nu = p$; therefore, by Step 2, the planner does not deviate in the good state if and only if $\mu \leq p$. The pooling equilibrium constitutes babbling if $\mu = \nu$ (nonbabbling if $\mu < \nu$).[9] When $\mu < p$, the planner does not deviate in the bad state if and only if $\lambda \geq \lambda^*(\mu, p)$. At the other extreme, the planner never deviates in the bad state when $\mu = p$. Hence, 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, after which Bayes' rule implies that $\mu = 0$ and $\nu = 1$. Given this, the planner never deviates in the good state, and does not deviate in the bad state if and only if $\lambda \leq \lambda^*(0, 1)$.[10] □

Proposition 1 shows that there always exists an equilibrium in which the planner is noninformative. Moreover, a transparent equilibrium in which the planner's private information is revealed also exists if and only if distortions are sufficiently small. If distortions are large—and thus individual actions (effort) are seriously distorted downward with respect to the social optimum—then a benevolent planner has a strong incentive to hide bad news, which reduces the likelihood of transparency in equilibrium.

2.4. Efficiency

Let us now compare the pooling and the separating equilibria from an ex-ante point of view (i.e., 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 payoffs) 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 the 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

9. A babbling equilibrium is an equilibrium in which (a) individual strategies disregard the planner's announcement and (b) the planner's signaling strategy disregards the realization of the shock. The only difference between babbling and nonbabbling pooling equilibria involves out-of-equilibrium beliefs.

10. There are only two possible types of mixed strategy equilibria: (a) 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 (b) semiseparating 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 $\nu = p/(p + (1 - p)\rho)$. Mixed strategy babbling equilibria always exist but a semiseparating equilibrium exists if and only if $\lambda = \lambda^*(0, p/(p + (1 - p)\rho))$. To see this, note that in the good state the social planner is willing to mix if and only if $\mu = \nu$, which is compatible with only babbling equilibria. Hence, we can say, without loss of generality, that nonbabbling equilibria in mixed strategies imply $\mu < \nu$ and $m(\vartheta) = N$; in other words, they may be only semiseparating. For $\rho \in \{0, 1\}$, we have the two pure strategy equilibria considered previously. For $\rho \in (0, 1)$, Bayes' rule implies that $\mu = 0$ and $\nu = p/(p + (1 - p)\rho)$. Given these posterior beliefs, the planner does not deviate in the good state and does not deviate in the bad state if and only if $\lambda = \lambda^*(0, p/(p + (1 - p)\rho))$.

chosen under perfect information about the state of the world (as would be obtained at a separating equilibrium), which is the solution by ℓ of $u_1(\ell, L, \lambda, \theta) = 0$. 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 $u_{144} = 0$), which implies that $\ell^P = \ell^S(\bar{\theta})$, because ℓ^P is the solution by ℓ of $pu_1(\ell, L, \lambda, \vartheta) + (1-p)u_1(\ell, L, \lambda, -\vartheta) = u_1(\ell, L, \lambda, \bar{\theta}) = 0$, where $\bar{\theta} = p\vartheta + (1-p)(-\vartheta)$. Linearity also implies that $\bar{u}^P = u^S(\bar{\theta})$. To see this, observe that $u^S(\bar{\theta}) = u(\ell^S(\bar{\theta}), \ell^S(\bar{\theta}), \lambda, \bar{\theta}) = pu(\ell^P, \ell^P, \lambda, \vartheta) + (1-p)u(\ell^P, \ell^P, \lambda, -\vartheta)$ and consider the function $W(\theta) = u(\ell^S(\theta), \ell^S(\theta), \lambda, \theta)$. Then, $\bar{u}^S = pW(\vartheta) + (1-p)W(-\vartheta)$ and $\bar{u}^P = W(\bar{\theta})$, so that $\bar{u}^S > \bar{u}^P$ if and only if $W''(\theta) > 0$. □

Proposition 2 implies that a benevolent planner sees transparency as being preferable ex ante 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 (owing to the convexity condition). By the same argument, concavity would make opaqueness ex ante preferable.

2.5. Implications of the Main Results

We have assumed an economy that is subject to a distortion and experiences random shocks about which a benevolent social planner has private information. This planner chooses an information policy that aims to maximize social welfare. Note that the most preferred policy before knowing the realization of the random shock may be different from the one preferred after knowing it. We assume that the government cannot credibly commit and therefore chooses the policy that maximizes ex-post social welfare. Our results show that if the distortion is large enough, then the only equilibrium policy will involve opaqueness because the planner always finds it preferable ex post to conceal negative shocks. As a consequence, individuals in equilibrium fully distrust the planner's announcements. If the distortion is small, then truth telling, too, can be an equilibrium.[11] However, transparency will be seen as preferable to opaqueness ex ante if and only if, under full information, the individual payoff is convex in the random variable. If this condition is satisfied, then the planner would ex ante prefer to be transparent; however, if distortions are substantial, then the equilibrium will be one of opaqueness. It would then be beneficial to delegate information policy to a separate agency that is committed to transparency, because such commitment would preclude ex-post decisions that are ex ante suboptimal.

In what follows, we examine two possible sources of distortions and apply our general results to obtain the corresponding equilibrium information policy. This approach allows us to demonstrate not only how the basic mechanism works in standard economic contexts but also the new insights it provides.

11. This result has a clear second-best flavor. After all, concealing information is, in itself, a distortion relative to full information. Yet, given the other distortions already in the economy, it may well be that this additional distortion ends up increasing aggregate welfare.

3. Transparency and Monopoly Power

In our first extension of the abstract model, we consider an economy with a monopolistic distortion. The model is a simplified version of the canonical real business cycle (RBC) model with no capital and a continuum of differentiated goods. To focus on the information analysis, we abstract from dynamic considerations (see, e.g., Angeletos and La'O 2009).[12]

3.1. The Economy and the Announcements Game

There is a mass 1 of identical individuals who work, consume, and own shares of a mass 1 of firms, each of which produces a different variety of a consumption good. Utility depends on consumption and labor:

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

Here,

$$c = \left(\int_0^1 c_i^{(\sigma-1)/\sigma}\, di \right)^{\sigma/(\sigma-1)}$$

is the Dixit–Stiglitz aggregator, where 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: using ℓ_i units of labor, firm i produces $y_i = A\ell_i$ units of its variety of good. Labor productivity is $A = \tilde{A} + \theta$ and so is the same for every firm, but it does depend on two factors: the observable component $\tilde{A} > 0$ and the ex-ante unobservable component θ (i.e., being in a boom or in recession), which is distributed according to distribution (1). We assume $\vartheta \in (0, \tilde{A})$ to ensure that productivity is always positive.[13]

Under perfect information on θ, the equilibrium of this economy is simple. Individuals choose $\{c_i : i \in [0, 1]\}$ and ℓ to maximize $u(c, \ell)$ under the budget constraint $\int_0^1 p_i c_i\, di = w\ell + \pi$, while taking as given the wage rate w, prices p_i, and distributed profits π.[14] Labor supply in this economy is $\ell = (w/P)^{1/(\delta-1)}$, where

$$P = \left(\int_0^1 p_i^{1-\sigma}\, di \right)^{1/(1-\sigma)}$$

is the aggregate price index; demand for good i is given by $c_i = (p_i/P)^{-\sigma} c$.

12. Relative to their analysis, we also abstract from dispersed information (as we have done in our abstract model).

13. The model can be immediately extended to the case of heterogeneous firm productivity. We present the identical-firms version for expositional simplicity and because it is sufficient to convey the main insights.

14. Here, we assume that individuals are identical in both productivity and shareholding. We discuss the role of heterogeneous productivity in Section 4. Heterogeneity in shareholding would not introduce any substantive changes in the present model, because it would affect the distribution of income but it would have no effect on individual behavior (because individuals take distributed profits as given).

Each firm i sets its product's price p_i so as to maximize profits $\pi_i = p_i y_i - w\ell_i$, while taking technology, demand, other firms' prices, and the wage rate as given. Thus, each firm prices according to the markup rule $p_i = (\sigma/(\sigma - 1))(w/A)$ and demands labor $\ell_i = c_i/A$. Because prices are the same for every firm, quantities are also the same: for all i, $p_i = P$ and $c_i = c$. The real wage $w/P = ((\sigma - 1)/\sigma)A$ is below labor productivity. Labor supply is then

$$\ell = \left(\frac{\sigma - 1}{\sigma} A\right)^{1/(\delta-1)},$$

consumption is

$$c = \left(\frac{\sigma - 1}{\sigma}\right)^{1/(\delta-1)} A^{\delta/(\delta-1)},$$

real profits are $\pi/P = c/\sigma$, and equilibrium utility is

$$u(c, \ell) = \left(\frac{\sigma - 1}{\sigma}\right)^{1/(\delta-1)} A^{\delta/(\delta-1)} \left(1 - \frac{\sigma - 1}{\delta\sigma}\right).$$

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 and thereby induces a suboptimal downward distortion in individual labor supply—relative to the (first-best) social optimum, which would require $\ell = A^{1/(\delta-1)}$. Second, this distortion is decreasing in the elasticity of substitution σ. Indeed, profit distribution creates a positive externality from labor supply, for which individual choices fail to account, and (real) profits are decreasing in σ. Third, equilibrium utility is convex in A and therefore in θ. This last observation makes the planner's transparency Pareto superior to opaqueness from the ex-ante point of view because, as is easy to check, this model adds economic structure to the specification of actions and their relation to utility while continuing to satisfy all of the previous model's assumptions.

Now we consider imperfect information, under 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; however, employment decisions are based on expectations, because, at this stage, both parties know m but neither knows θ. Fourth, once employment is determined, the economic fundamentals are publicly revealed, production is realized, and commodity prices adjust so as to clear the commodity markets (thus prices p_i, consumption choices c_i, the wage rate w, and profits π are all determined under full information). Ex-post social welfare is $W(\theta, m) = \int_0^1 u\left(c(\theta, m, \ell^*(m)), \ell^*(m)\right) di$, where $c(\theta, m, \ell^*(m))$ and $\ell^*(m)$ are the equilibrium values of consumption and labor supply, respectively.

We remark that, here, employment is a proxy for all kinds of input and production choices that are made before the perfect realization of aggregate uncertainty. Therefore, employment can also be viewed 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 differs from how employment is modeled in the standard New Keynesian approach, which assumes (for simplicity, if not for realism) that all employment can freely adjust to the true realized state.

3.2. Equilibrium with Monopoly Power

Our next result parallels Proposition 1. Let

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

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

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

PROPOSITION 3 (Equilibrium with monopoly power). *Given μ and ν, individuals' strategies are*

$$\begin{aligned} \ell^*(m) &= E\left[\frac{w}{P}|m\right]^{1/(\delta-1)} = \left(\left(\frac{\sigma-1}{\sigma}\right)\left(\tilde{A} + E[\theta|m]\right)\right)^{1/(\delta-1)}, \\ c_i(\theta, m, \ell) &= \left(\frac{p_i(\theta, m, \ell)}{P(\theta, m, \ell)}\right)^{-\sigma} c(\theta, m, \ell). \end{aligned}$$

Firms' strategies are

$$p_i(\theta, m, \ell) = \left(\frac{\sigma}{\sigma - 1}\right)\frac{w(\theta, m, \ell)}{\tilde{A} + \theta}.$$

There are two possible types of pure strategy equilibria.[15]

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

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

15. The structure of mixed strategy equilibria is analogous to that of the abstract model and is not reiterated here.

Proof. See the Online Appendix. □

3.3. *Efficiency and Equilibrium Selection*

We 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, respectively, at a separating (S) and a pooling (P) equilibrium. Then, irrespective of whether a separating equilibrium exists, the following statement holds.

PROPOSITION 4 (Ex-ante Pareto dominance). *For any constellation of parameters, we have (a) $\bar{\ell}^S < \bar{\ell}^P$ if and only if $\delta > 2$, (b) $\bar{y}^S > \bar{y}^P$, and (c) $\bar{u}^S > \bar{u}^P$.*

Proof. See the Online Appendix. □

Proposition 4 establishes that, for any degree of monopoly power, the transparent and credible revelation of information is ex-ante Pareto superior to concealing information. This is not surprising in light of Proposition 2, given that (as already mentioned) the full-information equilibrium utility is convex in θ.[16] That being said, Proposition 3 indicates that high monopoly power may prevent the transparent outcome from materializing in equilibrium.[17] The intuition is simple. Transparency allows individuals to work more when they are more productive and to work less when they are less productive. These effects unequivocally raise 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 (because workers dislike fluctuations in labor effort), this latter effect is always more than compensated by the higher expected level of consumption—hence the positive effect on welfare.[18]

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

Ex-ante Pareto dominance selects the separating equilibrium whenever it exists. Yet this is not always a good selection criterion in the present context because, whenever

16. This superiority is also in line with Angeletos, Iovino, and La'O (2011), since technology shocks drive the business cycle. Although we assume that the source of distortions (the elasticity of substitution here, and the tax rate in Section 4) is acyclical, the wedge between first-best and full-information labor supply is procyclical in both models. Angeletos and Pavan (2007) have shown that, in economies that are inefficient even under complete information, if this wedge covaries positively with full-information equilibrium strategies, then welfare is increasing (ex ante) in the accuracy of public information. The results reported here are coherent with those in Angeletos and Pavan.

17. For any parameter constellation, the ex-ante expected levels of individual labor supply, production, and indirect utility (whose relationships are identified in Proposition 4) are well defined, regardless of whether or not a separating equilibrium exists.

18. Transparency increases expected leisure time if the elasticity of labor supply is $\gamma \equiv 1/(\delta - 1) < 1$ (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 labor supply increases in booms. In contrast, if the elasticity of labor supply is $\gamma > 1$ (i.e., for $\delta < 2$), then labor supply is a convex function of expected wages. In that case, transparency increases expected labor supply relative to opaqueness.

$\sigma \in [\sigma^*(0, 1), \sigma^*(0, p)]$, the planner's preferences with respect to different equilibria are reversed in different states of the world.[19] Hence, it is worthwhile to examine different equilibrium refinements.

In cheap-talk games, the standard refinements based on Kohlberg and Mertens (1986), which restrict out-of-equilibrium beliefs, have little power because mixed strategy babbling equilibria always survive them. We consider a more recent refinement—no incentive to separate (NITS), which was introduced by Chen, Kartik, and Sobel (2008) for the express purpose of selecting equilibria in cheap talk—as well as a stronger refinement, neologism proof (NP) equilibrium proposed by Farrell (1993). In the Online Appendix, we define these concepts within the context of our model and show that, whenever it exists, the separating equilibrium satisfies both NITS and NP, and is the only NP equilibrium.

Overall, these results suggest that, for $\sigma \geq \sigma^*(0, 1)$, the separating equilibrium is the game's most natural prediction. Therefore, we conduct the following discussion while 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 reduces transparency.[20] To answer this question in the simplest way, it is convenient to focus on the case of linear labor supply ($\delta = 2$). In that case, expression (3) simplifies to $\sigma^*(0, 1) = \tilde{A}/\vartheta$. If we measure the relative amplitude of the business cycle by the ratio $\vartheta/\tilde{A}$ of shock size to structural productivity, then it is clear that economies subject to more pronounced business cycles have a larger support for transparency. This result may be surprising, but from the theoretical point of view it follows directly from how the support for transparency is determined—namely, by the welfare comparison that a credible planner makes when deciding whether or not to reveal information about a recession. When $\sigma = \sigma^*(0, 1)$, a credible planner is indifferent between the two alternatives. This is because the monopolistic distortion (the underwork resulting from monopoly power under truthful revelation) is equivalent, in welfare terms, to the information distortion (the overwork that follows when the planner lies). An increase in relative shock magnitude raises the information distortion relative to the monopolistic distortion, so it likewise increases the relative cost of lying and expands the support for transparency. The full generalization of this result to arbitrary

19. In booms, the social planner would prefer a separating equilibrium in which private information is revealed, thus boosting labor supply and welfare. In recessions, the planner would prefer a pooling equilibrium in which information is not revealed, so that labor supply and welfare are higher than under perfect information. The proof of this claim follows immediately from the proof of Lemma 2 in the Online Appendix; that is, $\sigma^*(0, p) > \sigma^*(0, 1)$ follows from Lemma 1 (also proved in the Online Appendix).

20. We refer to the interval $[\sigma^*(0, 1), \infty)$ as 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 (and is, indeed, the most natural prediction of the game) for $\sigma \geq \sigma^*(0, 1)$, that is, when the distortion is small.

parameter values is complex; even so, 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.[21] Whereas 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 of interest that our model features a fixed elasticity of labor supply with respect to expected wages: $\gamma = 1/(\delta - 1)$. However, although actual wages do not depend on whether or not equilibrium information is transparent, expected wages (and hence the actual labor supply) do depend on the information regime. In particular, they fluctuate over the business cycle under transparency but not under opaqueness. Over the business cycle, then, the elasticity of labor supply to actual wages is zero for an economy with large monopolistic distortions and thus characterized by opaque information. In contrast, this elasticity is positive and equal to γ for an economy with the same value of δ but with less monopoly power and thus with transparent information. The model therefore predicts that, ceteris paribus, income fluctuations over the business cycle will be more pronounced in economies with more competitive product markets—and thus with transparent rather than opaque information. The importance of information and expectations is an aspect that tends to be ignored in the empirical debate over estimating labor supply elasticity, which in our view deserves more attention.

4. Transparency, Taxation, and Inequality

As a second extension of the abstract model, we now consider an economy in which distortions arise from taxation rather than from monopoly power. Income taxes push net wages below individual productivity and thus render labor effort inefficiently low, because individuals fail to internalize the externality emerging from the redistribution of tax revenues—just 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 1 of individuals who have identical preferences over a homogeneous consumption good c and over labor effort ℓ, as 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 the individual supply of efficiency units of labor. Individual productivity depends on two factors: an idiosyncratic observable component (ability or human capital), which is denoted by β and is distributed according to the cumulative distribution function F with support on the nonempty interval $[b, B) \subset \mathbb{R}_+$; and the aggregate, ex ante

21. To prove this result, write $\sigma^*(0, 1) = 1/[1 - q(0, 1)]$ as in Lemma 1 (see the Online Appendix) and observe that $\partial\sigma^*(0, 1)/\partial\vartheta$ has the same sign as $\partial q(0, 1)/\partial\vartheta$. By applying l'Hôpital's rule, it is easy to obtain $\lim_{\vartheta\to 0}(\partial q(0, 1)/\partial\vartheta) = -1/\tilde{A} < 0$, so that $(\partial\sigma^*(0, 1)/\partial\vartheta)\Big|_{\vartheta=0} < 0$.

unobservable, and random component θ (being in boom or in recession), which is distributed according to distribution (1) and with $\vartheta \in (0, b)$ to ensure that individual productivity is always positive.

Thus, individual labor income 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 the tax revenues $T = \int_b^B t y_\beta \, dF(\beta)$ are equally redistributed, so individual consumption is equal to $c_\beta = (1 - t)y_\beta + T$. Because the population is continuous, each individual takes T as given.[22]

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

$$W = \int_b^B u_\beta \, dF(\beta) = \frac{\delta - 1 + t}{\delta}(1 - t)^{1/(\delta-1)} \int_b^B (\beta + \theta)^{\delta/(\delta-1)} \, dF(\beta),$$

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 distribution (1). Both the c.d.f. 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. The realization of θ is observed ex post by all individuals, who are paid accordingly. The aim of the benevolent social planner is to maximize social welfare $W = \int_b^B u_\beta \, dF(\beta)$. Here, u_β denotes the utility of an individual with ability β; this utility depends on t, on θ, on individual labor effort ℓ_β, and on the labor effort chosen by the entire population (because T depends on that effort). The equilibrium concept and our notation for beliefs and expectations remain as before.

4.2. Equilibrium and Efficiency

The main results on equilibrium and efficiency parallel those obtained previously, 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

$$x_\mu = (\beta + E[\theta|M])^{1/(\delta-1)} = (\beta + (2\mu - 1)\vartheta)^{1/(\delta-1)},$$
$$x_\nu = (\beta + E[\theta|N])^{1/(\delta-1)} = (\beta + (2\nu - 1)\vartheta)^{1/(\delta-1)};$$

22. The tax collection per capita (T) depends on the realization of θ. Therefore, individuals will entertain conjectures about the value of T. As we shall see, our assumption on individual preferences makes these conjectures immaterial because they have no effect on labor supply.

for $\mu \neq \nu$,

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

PROPOSITION 5 (Equilibrium with taxation and inequality). *Given μ and ν, equilibrium labor supply strategies are described by*

$$\ell_\beta^*(m) = \{(1-t)(\beta + E[\theta|m])\}^{1/(\delta-1)}.$$

There are two possible types of pure strategy equilibria.

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

(2) 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 the Online Appendix. □

Let us now compare the different equilibria from an ex-ante point of view. For an individual with ability β, let $\bar{\ell}_\beta^S$, $\bar{y}_\beta^S$, $\bar{u}_\beta^S$ and $\bar{\ell}_\beta^P$, $\bar{y}_\beta^P$, $\bar{u}_\beta^P$ denote the ex ante expected levels of labor supply, production, and indirect utility at a separating (S) and at a pooling (P) equilibrium, respectively. Then, whether or not a separating equilibrium exists, the following proposition holds.

PROPOSITION 6 (Ex-ante Pareto dominance). *For any parameter constellation, any ability distribution, and any level β of individual ability, we have (a) $\bar{\ell}_\beta^S < \bar{\ell}_\beta^P$, if and only if $\delta > 2$, (b) $\bar{y}_\beta^S > \bar{y}_\beta^P$, and (c) $\bar{u}_\beta^S > \bar{u}_\beta^P$.*

Proof. See the Online Appendix. □

The results for equilibrium selection are also very similar to those obtained in the monopoly power model, and these are discussed in the Online Appendix. The principal insight is again that, whenever it exists (i.e., for $t \leq t^*(0, 1)$), the separating equilibrium is the game's most natural prediction. In light of this finding, for the discussion to follow we assume that the economy coordinates on the separating equilibrium when there is one.

4.3. Transparency and Inequality

The main new insight provided by this model concerns the role of inequality and its effects on transparency.[23] Indeed, this is a natural environment for asking whether

23. The effects of the business cycle amplitude are similar to those discussed for the previous model and are analyzed in the Online Appendix. The taxation model also confirms that, all else equal, both output

inequality favors or reduces transparency.[24] In the following proposition, we use Lorenz dominance (second-order stochastic dominance) as a criterion for establishing whether one given distribution has more inequality than another. Let $\gamma \equiv 1/(\delta - 1)$ 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, for a shift from skill distribution F to a more unequal distribution G that is Lorenz-dominated by F, the following statements hold.*

- *If $\gamma = 1$, then the increase in inequality has no effect on transparency.*
- *If $\gamma < 1$, then the increase in inequality favors transparency.*
- *If $\gamma > 1$, then for*

$$\hat{\gamma} = \frac{2}{1 - t^*(0, 1)} > 2,$$

  we find that $\gamma \in (1, \hat{\gamma}]$ is a sufficient condition for the increase in inequality to reduce transparency.

Proof. See the Online Appendix. □

First, we recall that most of the literature on information transparency assumes $\gamma = 1$ (i.e., a linear labor supply) and thus assumes away the effects of inequality. However, the emerging general picture is that inequality matters for transparency—and in a way that depends on the shape of the labor supply curve. In particular, if labor supply is rigid ($\gamma < 1$), as most micro estimates suggest, then inequality favors transparency. Yet, if labor supply is elastic, as many macro models assume, then inequality reduces transparency.[25]

To grasp the intuition behind this result, observe that $t^*(0, 1)$ depends on the welfare comparison between (credibly) revealing and not revealing information—conditional on being in a recession.[26] Relative to concealing information, transparency in recessions increases the leisure and decreases the consumption of each individual. It is therefore useful to disentangle the effects of inequality on $t^*(0, 1)$ into those working

and hours worked fluctuate more when the government is transparent. This result is consistent with the evidence provided by Demertzis and Hughes-Hallett (2007) and with the negative relationship between taxation and output volatility found by Debrun, Pisani-Ferry, and Sapir (2008). Yet, because the government tends to be transparent when aggregate shocks are relatively large, the ceteris paribus condition should not be neglected.

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

25. For evidence on labor supply elasticity, see, for example, Fiorito and Zanella (2012).

26. The relative social welfare gain to transparency in recessions depends on the tax rate: it is positive (negative) for low (high) tax distortions. The formal details on such a comparison are provided in footnote 7 of the Online Appendix.

through the consumption differential and those working through the leisure differential between transparency and opaqueness.

Because of the complementarity between skills and effort, an increase in skill inequality raises mean income, and hence mean consumption, independently of labor supply elasticity.[27] Higher consumption drives the welfare advantage of opaqueness over transparency, and so inequality, by raising aggregate consumption, favors opaqueness. However, utility depends not only on consumption but also on leisure, 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 ($\gamma < 1$). So, with a rigid labor supply, an increase in skill inequality also increases aggregate leisure time. In this case, higher leisure drives the welfare advantage of transparency over opaqueness, and so inequality, by raising aggregate leisure, favors transparency. Thus, if the labor supply is inelastic, then the two welfare effects of inequality (through consumption and through leisure) work in opposite directions: one favors opaqueness and the other favors transparency. The overall effect depends on which force dominates. We have shown that, with a 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; hence, aggregate consumption increases and aggregate leisure decreases. Both of these effects work in the same direction, which means that skill inequality favors opaqueness.

To the extent that inequality generates higher tax rates and thus greater distortions, it may also have an indirect effect on transparency. To explore this mechanism, in Albornoz, Esteban, and Vanin (2009), we have provided a political economy extension of the analysis where taxes are chosen by majority voting (along the lines of Meltzer and Richard 1981).[28] The impact of inequality then results from the combination of two effects. On the one hand, greater 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, greater inequality changes the government's valuation of truth telling (in recessions) in a way that depends on the elasticity of labor supply, as discussed in Proposition 7. If labor supply is elastic, then an increase in inequality amplifies the distortion created by taxation and thus raises the incentive to conceal bad news. In this case, the two channels move in the same direction and inequality reduces transparency. However, if labor supply is inelastic, then an increase in inequality decreases the magnitude of the tax distortion and thus raises the incentive to reveal bad news (transparency). Because the two channels work in opposite directions, the net effect cannot be established in general.[29]

27. Although it might be surprising, the fact that skill inequality is welfare increasing is a direct consequence of pairing the skill–effort complementarity with a Benthamite social welfare function.

28. We examine 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. Under the model's assumptions, a benevolent government that is completely free to set the tax rate would set it at zero.

29. Focusing on the case of a unit elasticity of labor supply, we show the following. (a) If the size of the shock is larger than a threshold level, then the equilibrium is informative for all levels of inequality. (b)

5. Concluding Discussion

In this paper, we have investigated how government transparency depends on economic distortions, which 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 effect transparent communication. 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.[30] Although this correlation may reflect the influence of political institutions on both these variables, our theory suggests that causality may also run from distortions directly to the level of transparency. An open avenue for future research is to use available data to test this hypothesis.

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 under other assumptions for future investigation.[31] Within our framework, it is worth noting that precisely when the government "lies" (in the sense that, in recessions, it sends the same message that it sends in booms), individuals prefer ex post that it did so. It is therefore unproblematic that the government's private information is ex post verifiable. Note also that our restriction to two elements each for both the state and message space has the effect of polarizing equilibria as either full revelation or no revelation. An extension to the continuous case would generate equilibria with partial revelation and would thus allow us to study the degree of information precision, but it would not affect the main intuition and the main results.[32]

Perhaps the most interesting extensions of our framework concern the various possible forms of interaction between economic distortions and transparency. First, we have assumed that the business cycle is driven by a productivity shock about which the

However, below this threshold, inequality reduces transparency. Specifically, the equilibrium is informative for low inequality and uninformative for high inequality, yet both an informative and an uninformative equilibrium exist for intermediate levels.

30. This holds, for example, when fiscal transparency is measured by the Open Budget Index (developed by the International Budget Partnership) and distortions by either the Time to Start Business or the Ease of Doing Business (which capture barriers to entry and obstacles to business activity; both are provided by the World Bank). The correlation also holds when distortions are captured by several measures of taxes, including the top marginal tax rate on labor income, taxes on goods and taxes on international trade (all available from the OECD World Tax Indicators).

31. For instance, an incumbent government might prefer being overoptimistic in order to influence individuals' beliefs on its ability. Although this would provide an extra incentive to conceal bad news (beyond the motives emphasized in this paper), we doubt that it would alter our main results.

32. This can be seen most clearly in Albornoz, Esteban, and Vanin (2009), where mixed strategies allow for a semiseparating equilibrium.

government has private information. However, we could also assume (as in Angeletos, Iovino, and La'O 2011) that the government's private information instead concerns shocks to monopoly markups or to labor wedges, which might change our results and make opaqueness ex ante desirable. Second, we have assumed that distortions are persistent and that the government cannot easily eliminate them. An investigation of the political economy reasons for this difficulty is a promising research direction. For example, an elected government might be influenced by lobbying activity or by a demand for redistribution. In Albornoz, Esteban, and Vanin (2009), we have explored the latter possibility and have shown that inequality reduces transparency because it generates higher taxes. This finding modifies the results obtained in Section 4, where we have shown 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. Yet, it is certainly worth investigating exactly how the latter interacts with monetary and fiscal policy, especially because the shock's direction and size may well depend on policy actions. Fourth, we have assumed that the government perfectly observes the shock and that individuals have no other source of information. However, a natural extension would be to examine government incentives to put a privately observed noisy signal of the shock in the public domain when individuals also have access to dispersed and noisy information. These lines of research remain open.

References

Albornoz, Facundo, Joan Esteban, and Paolo Vanin (2009). "Government Information Transparency." UFAE and IAE Working Paper No. 774.09, Unitat de Fonaments de l'Anàlisi Econòmica (UAB) and Institut d'Anàlisi Econòmica (CSIC).

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

Amato, Jeffery D., Stephen Morris, and Hyun Song Shin (2002). "Communication and Monetary Policy." *Oxford Review of Economic Policy*, 18, 495–503.

Angeletos, George-Marios and Jennifer La'O (2009). "Noisy Business Cycles." *NBER Macroeconomics Annual*, 24, 318–378.

Angeletos, George-Marios and Alessandro Pavan (2007). "Efficient Use of Information and Social Value of Information." *Econometrica*, 75, 1103–1142.

Angeletos, George-Marios, Luigi Iovino, and Jennifer La'O (2011). "Cycles, Gaps, and the Social Value of Information." NBER Working Paper No. 17229.

Athey, Susan, Andrew Atkeson, and Patrick J. Kehoe (2005). "The Optimal Degree of Discretion in Monetary Policy." *Econometrica*, 73, 1431–1475.

Blinder, Alan S., Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen (2008). "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence." *Journal of Economic Literature*, 46, 910–945.

Chen, Ying, Navin Kartik, and Joel Sobel (2008). "Selecting Cheap-Talk Equilibria." *Econometrica*, 76, 117–136.

Crawford, Vincent P. and Joel Sobel (1982). "Strategic Information Transmission." *Econometrica*, 50, 1431–1451.

Cukierman, Alex (2009). "The Limits of Transparency." *Economic Notes*, 38, 1–37.

Cukierman, Alex and Allan H. Meltzer (1986). "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information." *Econometrica*, 54, 1099–1128.

Debrun, Xavier, Jean Pisani-Ferry, and André Sapir (2008). "Government Size and Output Volatility: Should We Forsake Automatic Stabilization?" IMF Working Paper 08/122.

Demertzis, Maria and Andrew Hughes-Hallett (2007). "Central Bank Transparency in Theory and Practice." *Journal of Macroeconomics*, 29, 760–789.

Farrell, Joseph (1993). "Meaning and Credibility in Cheap-Talk Games." *Games and Economic Behavior*, 5, 514–531.

Farrell, Joseph and Robert Gibbons (1989). "Cheap Talk with Two Audiences." *American Economic Review*, 79(5), 1214–1223.

Fiorito, Riccardo and Giulio Zanella (2012). "The Anatomy of the Aggregate Labor Supply Elasticity." *Review of Economic Dynamics*, 15, 171–187.

Gavazza, Alessandro and Alessandro Lizzeri (2009). "Transparency and Economic Policy." *Review of Economic Studies*, 76, 1023–1048.

Gavazza, Alessandro and Alessandro Lizzeri (2011). "Transparency and Manipulation of Public Accounts." *Journal of Public Economic Theory*, 13, 327–349.

Hellwig, Christian (2005). "Heterogeneous Information and the Benefits of Public Information Disclosures." UCLA Department of Economics Working Paper 283.

James, Jonathan G. and Phillip Lawler (2011). "Optimal Policy Intervention and the Social Value of Public Information." *American Economic Review*, 101(4), 1561–1574.

Kohlberg, Elon and Jean-Francois Mertens (1986). "On the Strategic Stability of Equilibria." *Econometrica*, 54, 1003–1037.

Kohn, Donald L. and Brian Sack (2004). "Central Bank Talk: Does it Matter and Why?" In *Macroeconomics, Monetary Policy and Financial Stability 2003*. Bank of Canada, Ottawa, pp. 175–206.

Kurz, Mordecai (2005). *Diverse Beliefs, Forecast Errors and Central Bank Policy*. Stanford University, CA.

Meltzer, Allan H. and Scott F. Richard (1981). "A Rational Theory of the Size of Government." *Journal of Political Economy*, 89, 914–927.

Morris, Stephen and Hyun Song Shin (2002). "Social Value of Public Information." *American Economic Review*, 92(5), 1521–1534.

Morris, Stephen, Hyun Song Shin, and Hui Tong (2006). "Social Value of Public Information: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con: Reply." *American Economic Review*, 96(1), 453–455.

Oh, Seonghwan and Michael Waldman (1990). "The Macroeconomic Effects of False Announcements." *Quarterly Journal of Economics*, 105, 1017–1034.

Roca, Mauro (2010). "Transparency and Monetary Policy with Imperfect Common Knowledge." IMF Working Paper No. 10/91.

Rodríguez Mora, José Vicente and Paul Schulstald (2007). "The Effect of GNP Announcements on Fluctuations of GNP Growth." *European Economic Review*, 51, 1922–1940.

Romer, Christina and David Romer (2000). "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review*, 90(3), 429–457.

Stein, Jeremy C. (1989). "Cheap Talk and the Fed: A Theory of Imprecise Policy Announcements." *American Economic Review*, 79(1), 32–42.

Svensson, Lars E. O. (2006). "Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con." *American Economic Review*, 96(1), 448–452.

Woodford, Michael (2003). "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, edited by Philippe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford. Princeton University Press, Princeton, NJ, pp. 25–58.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Online Appendix