Optimal Information Acquisition and Monetary Policy

Thomas E. Cone*

State University of New York, College at Brockport

ABSTRACT

I study optimal monetary policy in an expectational Phillips Curve environment in which private agents optimally choose their amount of information pertinent to predicting policy. ARCH shocks produce interesting information acquisition (IA) dynamics. Under discretion, IA dynamics cause time-varying effectiveness of policy because of the expectational Phillips Curve; policy may be rendered completely ineffective. Greater economic volatility can induce fewer agents to be informed, though only in unstable equilibria. For an agent to become informed increases economic volatility; informed agents therefore impose a negative externality on others. Under commitment policy's effectiveness is again time-varying, but policy is never completely ineffective.

JEL classification: E52; E58; D83; D84

Keywords: information; expectations; monetary policy

*Assistant Professor of Economics, Department of Business Administration and Economics, 119 Hartwell Hall, SUNY Brockport, Brockport, NY, 14420. Phone: (585) 395-5524. Email: tcone@brockport.edu.

1 Introduction

This paper presents optimal monetary policy for a model in which private agents can choose whether to acquire information that helps them predict policy actions. Private sector (PS) information acquisition is important in the model because a monetary policy action has real effects only to the extent that it surprises the PS. This means that as the PS's information acquisition changes over time, the effectiveness of monetary policy also changes, a phenomenon that would appear to an econometrician as an output-inflation tradeoff with a time-varying slope. This topic is important for monetary policy in empirical economies; to the extent that policy neutralization occurs in such economies when private agents become informed, their information acquisition affects the results produced by monetary policy.[1]

In the model, interesting dynamics of information acquisition are produced by output shocks that have autoregressive variance—autoregressive conditional heteroskedasticity, or ARCH. The monetary authority's offsetting of ARCH shocks causes inflation to also exhibit ARCH; furthermore, inflation is the variable the PS might become informed about, but its ARCH behavior implies that becoming informed is worth the cost only in some periods.

This paper advances the literature on this topic in several ways. One way is that PS agents are allowed to have different costs of acquiring information, which seems to be unique in the "information" strand of the monetary policy literature. Another way is the incorporation of ARCH shocks. The third way is the inclusion of both private agents who optimally choose

[1] Since Lucas (1972, 1973) formalized the view that only surprise changes in the money supply have real effects, some work has disputed that view. Woodford (2002) rescues this "surprise" strand of the transmission literature from the principal objections which have been raised against it. Those objections were (1) the theoretical objection that monetary aggregates were almost immediately and freely available, so that private agents could not plausibly be assumed to be surprised, and (2) the empirically-based objection that the time profile of output's response to a monetary shock was much more spread out than the very short lag with which monetary data were published. Woodford models agents as not observing monetary data in real time with perfect precision because of finite bandwidth in the information-theoretic sense. For reasonable parameterizations his model implies a time profile of output's response to a monetary shock that agrees well with the data.

their information and a monetary authority (MA) that optimally conducts monetary policy. The only other paper with both optimal information acquisition (IA) and an optimizing MA is that of Branch *et al* (2005) (discussed below); the rest of the literature has only one of these features. There is of course a vast literature on optimal monetary policy under exogenous PS information sets. Regarding optimal IA, Evans and Ramey (1992) study a model in which private agents endogenously (optimally in some cases) acquire information under *ad hoc* monetary policy. Hahm (1987) studies the PS's optimal IA but does not describe optimal monetary policy response to output shocks. Neither does Hahm consider private agents with differential costs of becoming informed, ARCH shocks, or the differences between discretionary and committed policy, which I do below. Branch *et al* (2005) study optimal monetary policy in a model with endogenous IA but assume all PS agents have the same cost of IA. They also assume the shocks to which the MA responds are homoskedastic, so that the IA problem is not intrinsically time-varying.

I study a representative agent PS and a many-agent PS. Results for a representative agent are as follows: Under discretionary policy, time-varying PS information acquisition causes periods of effective policy to alternate with periods of completely ineffective policy. Under policy with commitment the PS never acquires information in equilibrium; the MA commits to a policy that induces the PS to remain ignorant, in order to preserve policy's ability to affect the real economy.

With many PS agents, results are as follows: First, under discretion, changing IA causes policy's effectiveness to change over time; unlike the case of the representative agent, policy's effectiveness is more general than either completely effective or completely ineffective, because the fraction of the PS that is informed can be between zero and one. This provides a possible explanation for time-varying output-inflation tradeoffs, e.g., the weakening of the tradeoff in the US in the 1970s, when the economy became more volatile. In empirical economies there is a similar relationship across countries; price level shocks have less effect on output in countries with higher variances of the price level; see e.g., Lucas's classic (1973) paper, Apergis and Miller (2004). Second, an increased variance of the shock can induce fewer private agents to become informed and can raise welfare, but such counterintuitive results occur only in equilibria which are implausible on stability grounds. Third, the higher is the fraction of agents who are informed, the higher is inflation's equilibrium volatility; informed agents therefore impose a negative externality on uninformed ones.[2] Fourth, under commitment, in contrast to the representative agent case, some private agents may acquire information in equilibrium, though again the MA never chooses policy that would induce all private agents to be informed.

A result that emerges for both the representative agent PS and the many-agent PS is that policy tends to be rendered ineffective precisely when it would be most useful: The PS becomes informed when output volatility, which the MA would like to offset, is high. With discretion this IA behavior tends to neutralize policy in high-variance periods, and can completely neutralize policy, even with the many-agent PS. With commitment policy retains some effectiveness in high-variance periods, but this is because the MA is forced to be less responsive to shocks to limit the PS's IA. For this reason policy is of limited help in high-volatility periods, not because the policy action has a weak effect on output, but because the policy action itself is limited.

2 Benchmark results

2.1 Uninformed private sector

Before studying endogenous IA I review some benchmark results from the extant literature with exogenous information and iid shocks. The basic model is similar to ones used in much of the time inconsistency literature; it is based on the model in Blanchard and Fischer (1992,

[2] This result is fairly novel. It is more usual for IA to have positive external effects on the uninformed, because of the information-revealing behavior of the informed, cf. the free rider problem in financial markets. The only other paper I know of in which a negative externality from information processing can arise is Evans and Ramey (1992). In that paper, though, the reason is different; it is not caused by the MA's optimal response to the fraction of informed agents, which is what drives the effect here.

chapter 11), which is based on Barro and Gordon (1983) and Canzoneri (1985). Time is discrete and infinite. In period t the MA chooses inflation to minimize the loss function

$$l_t = a\pi_t^2 + (y_t - k\overline{y})^2 \tag{1}$$

where $a > 0$, $k \geq 1$, π_t is the rate of inflation at time t, y_t is output at time t, and $\overline{y}$ is the natural level of output. If $k > 1$, given the aggregate supply curve below, the standard positive inflation bias arises. The MA minimizes this loss subject to the constraint imposed by the aggregate supply curve

$$y_t = \overline{y} + b(\pi_t - E[\pi_t|I_t]) + u_t \tag{2}$$

where the output shock u_t is an iid random variable with mean zero, variance σ^2, and pdf $f(u_t)$, which the MA observes contemporaneously and before setting inflation. $E[\pi_t|I_t]$ is the rational expectation of date t inflation conditional on the PS's start-of-date-t information set I_t, the set of all variables dated $t-1$ or earlier. One rationale for the MA's superior information is that it has first access to economic data. Another is that the MA generates a private forecast of output that is superior to the PS's forecast, as in Canzoneri (1985), Walsh (1995), and Romer and Romer (2000), although in this interpretation the MA would base policy not on u_t but on its forecast of u_t. Aggregate supply curves of the above form can be derived from microfoundations in many models, e.g., nominal wage contracts models as in Fischer (1977) and separated markets models as in Lucas (1973).

Under discretionary policy with a myopic MA[3] the game between the PS and the MA has a unique rational expectations equilibrium (REE) level of inflation,

$$\pi_t = \frac{b}{a}(k-1)\overline{y} - \frac{b}{a+b^2}u_t \tag{3}$$

and REE output is

$$y_t = \overline{y} + \frac{a}{a+b^2}u_t. \tag{4}$$

[3] In an interemporal context the MA will not necessarily play the one-period discretionary outcome unless it is myopic, because of the possibility of reputation-building. I do not study reputation-building in this paper.

The MA partially offsets shocks' effect on output; monetary policy cannot be used to systematically raise output. Inflation's variance is $[b/(a+b^2)]^2\sigma^2$; this is also the PS's expected squared inflation forecast error.

With commitment the MA commits to a policy $\pi(u_t)$ to minimize the expected loss in a given period (since there are no connections between periods the one-period optimum is also the dynamic optimum) subject to the constraint that $E[\pi_t|I_t]$ is the rational expectation of $\pi(u_t)$. The optimal committed policy is

$$\pi(u_t) = -\frac{b}{a+b^2}u_t \tag{5}$$

and output is

$$y_t = \overline{y} + \frac{a}{a+b^2}u_t, \tag{6}$$

the same as output in the discretionary case. Under commitment the MA responds to shocks just as strongly as under discretion but average inflation is zero instead of positive.

2.2 Informed private sector

In this section I derive results that are useful with endogenous IA. For the PS to acquire information means it observes u_t contemporaneously and, since it knows the MA's preferences, it can then predict period t inflation. Under discretion REE inflation, which is also the PS's inflation expectation, is

$$\pi_t = \frac{b}{a}(k-1)\overline{y} - \frac{b}{a}u_t, \tag{7}$$

and output is $y_t = \overline{y} + u_t$. REE inflation is more responsive to shocks when the PS is contemporaneously informed about the shock. Since actual inflation equals expected inflation monetary policy has no effect on output and so the output shock is felt fully on output. In the canonical time inconsistency model inflation is positive on average without raising average output; here inflation also varies without stabilizing output. The optimal committed monetary policy with an always-informed PS is simply $\pi_t = 0$ and output is again $y_t = \overline{y} + u_t$. If the PS is informed monetary policy cannot affect output, so the best policy is simply to commit to zero inflation.

3 Information acquisition with ARCH shocks

3.1 Discretionary policy

This section introduces autoregressive variance of the output shock and endogenous PS choice of whether to observe the shock contemporaneously. Some recent empirical studies finding autoregressivity in output shocks' variances are Altissimo and Violante (2001), Ho and Tsui (2002), and Sims and Zha (2006). ARCH output shocks generate nontrivial IA dynamics because they cause the MA's optimal inflation to also be ARCH, so the PS's optimal choice of whether to become informed varies over time.

Let $z \geq 0$ denote the time, effort and resource cost necessary to become informed, i.e., to observe u_t or to infer it from other data. Let $i_t \in \{0, 1\}$ denote the PS's date t IA choice, $i_t = 0$ meaning the PS is uninformed and $i_t = 1$ meaning the PS is informed. The PS's date t loss is

$$l_t^{PS} = i_t z + (E[\pi_t|\Omega_t] - \pi_t)^2 . \tag{8}$$

The second term is the squared inflation forecast error, a standard form for such losses; see, e.g., Evans and Ramey (1992), Branch *et al* (2005). Ω_t is the endogenous date-t information set; it is I_t if the PS remains ignorant and $I_t \cup u_t$ if the PS becomes informed. In a given period one of the terms in (8) is zero; if the PS remains ignorant the first term is zero and if it engages in IA its forecast error is zero. The PS minimizes the expectation of (8) conditional on its default information set I_t. This entails comparing the loss if it were to acquire information, z, with its expected loss if it were to remain ignorant, $E[\pi_t^2|I_t] - \{E[\pi_t|I_t]\}^2$, the mean-centered variance of π_t conditional on I_t. Note this variance changes over time due to the ARCH behavior of the output shock.

For the shock's autoregressive variance I assume u_t has pdf $f_t(u_t) \in \{f_{low}(\cdot), f_{high}(\cdot)\}$ where $f_t(\cdot)$ has mean zero and variance $\sigma_t^2 \in \{\sigma_{low}^2, \sigma_{high}^2\}$, σ_{low}^2 being the variance of the pdf $f_{low}(\cdot)$ and σ_{high}^2 $(> \sigma_{low}^2)$ being the variance of the pdf $f_{high}(\cdot)$. Because the variance is the object of interest for the PS, I express the following assumptions in terms of the variance: For some constant $x > 0$, if $u_{t-1}^2 \in [0, x]$ then $\sigma_t^2 = \sigma_{low}^2$ and if $u_{t-1}^2 > x$ then $\sigma_t^2 = \sigma_{high}^2$. I assume $f_{low}(\cdot)$ and $f_{high}(\cdot)$ are such that $0 < \Pr[u_t^2 > x|\sigma_t^2 = \sigma_{low}^2] < \Pr[u_t^2 > x|\sigma_t^2 = \sigma_{high}^2]$ and $0 < \Pr[u_t^2 \leq x|\sigma_t^2 = \sigma_{high}^2] < \Pr[u_t^2 \leq x|\sigma_t^2 = \sigma_{low}^2]$. The variance is a two-state Markov process since—given $f_{low}(\cdot)$ and $f_{high}(\cdot)$—σ_t^2 implies the probability of $u_t^2 \leq x$, and so implies the probabilities of $\sigma_{t+1}^2 = \sigma_{low}^2$ and $\sigma_{t+1}^2 = \sigma_{high}^2$. This specification is quite general. It is not necessary to assume, for example, that $f_{low}(\cdot)$ and $f_{high}(\cdot)$ are continuous or symmetric around zero or that they come from the same family of functions parameterized by σ_t^2.

Having observed u_{t-1} at the start of period t, the PS knows σ_t^2 and it has the opportunity to pay the cost z and observe u_t. I assume that if the cost of observing the shock equals the expected benefit the PS remains ignorant. Then the MA, which always observes u_t contemporaneously, sets inflation and then output is produced. With an uninformed PS inflation's variance conditional on the start-of-date information set is $[b/(a+b^2)]^2\sigma_t^2$. Let var_π^{low} denote that variance for $\sigma_t^2 = \sigma_{low}^2$ and var_π^{high} denote that variance for $\sigma_t^2 = \sigma_{high}^2$. Under discretion the MA's optimal policy is (3) when the PS is uninformed and (7) when it's informed, so the PS's expected loss is $[b/(a+b^2)]^2\sigma_t^2$ if it were not to observe the current shock and z if it were to observe it. Therefore we have the following proposition.

Proposition 1 *If $z < \mathrm{var}_\pi^{low}$ the PS becomes informed, and so monetary policy fails to affect output, every period. If $z \geq \mathrm{var}_\pi^{high}$ the PS never becomes informed, and so monetary policy can always affect output. If $\mathrm{var}_\pi^{low} \leq z < \mathrm{var}_\pi^{high}$ then the PS is ignorant and monetary policy is effective in a positive fraction of periods, and the PS is informed and monetary policy is ineffective in a positive fraction of periods.*

High-volatility periods, in which stabilization policy would be most useful, are precisely those in which it is rendered ineffective. In data sets of such an economy one observed relationship would be that, on average, inflation has a greater effect on output in periods when inflation's variance is low, and is less able to affect output in periods when inflation's variance is high. An empirical episode that comes to mind is the 1970s in the US, when inflation's variance rose and it lost its ability to affect real variables in desired ways.

The PS's IA choice does not only respond to inflation's variance; it also affects it. Given an informed PS, inflation's variance is $[b/a]^2\sigma_t^2$, which is larger than the variance given that the PS is uninformed.

3.2 Welfare

For low values of the shock's variance, the PS's expected loss is increasing in the variance. For values of the variance that induce IA, variances greater than $z[(a+b^2)/b]^2$, the PS's loss is simply z. The higher is the shock's variance, the higher is the MA's expected loss. The direct effect of a higher shock variance of course is to destabilize the economy, but there also may be the indirect effect on the PS's information acquisition. If the PS acquires information, the MA loses any ability to stabilize output, and inflation's equilibrium variance rises. Higher IA cost z is worse for the PS but better for the MA.

3.3 Committed policy

Suppose the MA can commit to a policy in which each period's inflation is a function of the history of shocks and the PS's IA choice that period. The PS, knowing the committed policy, chooses in each period whether to observe the current shock, then the MA, having observed the shock and the PS's choice, sets inflation according to the committed policy. At date 0, before observing u_0, the MA chooses its committed policy to minimize the intertemporal loss function

$$\sum_{t=0}^{t=\infty} \beta^t E[a\pi_t^2 + (y_t - k\overline{y})^2 | I_0] \tag{9}$$

where $\beta \in (0,1)$ is the MA's discount factor, though as will be seen, the intertemporal element to the MA's problem is trivial. The PS's IA strategy is a function $i_t = i(\sigma_t^2, \pi(\cdot))$ where $\pi(\cdot)$ is the MA's committed policy. The MA's policy is a function $\pi(\sigma_t^2, u_t, i(\cdot))$ that specifies an inflation rate for every combination of σ_t^2, u_t, and private sector IA strategy. I restrict attention to functions $\pi(\sigma_t^2, u_t, i(\cdot))$ that are twice differentiable in u_t.[4]

[4] For both players I restrict attention to strategies that are pure and time-invariant.

Definition 1 *A Nash equilibrium is a pair of strategies* $\pi(\sigma_t^2, u_t, i(\cdot))$ *and* $i(\sigma_t^2, \pi(\cdot))$ *such that* $\pi(\sigma_t^2, u_t, i(\cdot))$ *minimizes (9) if the private sector plays* $i(\sigma_t^2, \pi(\cdot))$ *and* $i(\sigma_t^2, \pi(\cdot))$ *minimizes the expectation of (8) conditional on* I_t *if the MA has committed to* $\pi(\sigma_t^2, u_t, i(\cdot))$.

The PS part of the equilibrium is simple; its optimal strategy is simply to become informed if next period's inflation variance, given the committed policy, is sufficiently high and to remain ignorant if it is sufficiently low. The MA's part is more involved. Before proceeding to a proposition about this game's equilibria it is helpful to have some terminology. Let "fully responsive inflation" refer to the inflation in (5). Let "fully responsive policy" mean a committed policy which has fully responsive inflation in periods in which the PS is uninformed, and zero inflation in periods in which the PS is informed. Let "bounded inflation" denote inflation smaller in absolute value than fully responsive inflation in a particular period, and "bounded policy" denote a committed policy in which inflation's variance conditional on I_t is always bounded to prevent IA, $E[(\pi_t - \widehat{\pi})^2 | I_t] \leq z$ for all t, where $\widehat{\pi}$ is inflation's unconditional mean under a given policy.

In the following proposition I assume $[b/(a+b^2)]^2 \sigma_{low}^2 \leq z < [b/(a+b^2)]^2 \sigma_{high}^2$, so the MA's choice of a committed policy is non-trivial. The other two possibilities, $z < [b/(a+b^2)]^2 \sigma_{low}^2$ and $[b/(a+b^2)]^2 \sigma_{high}^2 \leq z$, are simple and I address them below.

Proposition 2 *There exists exactly one Nash equilibrium. In the equilibrium the private sector never observes the current shock, and if* $\sigma_t^2 = \sigma_{low}^2$ *then* $\pi_t = [-b/(a+b^2)]\, u_t$*, and if* $\sigma_t^2 = \sigma_{high}^2$ *then* $\pi_t = -\sqrt{z/\sigma_{high}^2}\, u_t$.

Proof: Suppose the MA's inflation conditional on PS ignorance is such that the PS's optimal choice at date t is to be informed. Optimal committed policy has zero inflation in periods in which the PS is informed, which entails a loss of $(u_t - (k-1)\overline{y})^2$. So committing to bounded inflation, for periods in which fully responsive inflation would induce IA, is optimal if the expected loss with bounded inflation is less than $E[(u_t - (k-1)\overline{y})^2]$, where the expectation is conditional on $\sigma_t^2 = \sigma_{high}^2$. The optimal bounded policy (its optimality is proven in the Appendix) is: (1) If the PS is informed this period, $\pi_t = 0$. (2) If the PS is uninformed this period: (a) If $\sigma_t^2 = \sigma_{low}^2$ then $\pi_t = [-b/(a+b^2)]\, u_t$, so inflation's variance is $[b/(a+b^2)]^2\sigma_{low}^2 \leq z$. (b) If $\sigma_t^2 = \sigma_{high}^2$ then $\pi_t = -\sqrt{z/\sigma_{high}^2}u_t$, so inflation's variance is z. The MA prefers the policy described by (1) and (2) to the policy that induces IA, as can be seen by comparing the policies under the same pair $\{\sigma_t^2, u_t\}$. If $\sigma_t^2 = \sigma_{low}^2$ the bounded policy and the fully responsive policy produce the same inflation and the PS remains ignorant under both policies, so the loss is the same. If $\sigma_t^2 = \sigma_{high}^2$ the PS remains ignorant under the bounded policy and $\pi_t = -\sqrt{z/\sigma_{high}^2}u_t$, and under fully responsive policy the PS acquires information, so the MA sets $\pi_t = 0$. The expected loss with $\pi_t = 0$ is

$$\sigma_{high}^2 + (k-1)^2\overline{y}^2 \tag{10}$$

and the expected loss under the bounded policy is

$$(a+b^2)z - 2b\sqrt{z}\sqrt{\sigma_{high}^2} + \sigma_{high}^2 + (k-1)^2\overline{y}^2. \tag{11}$$

(11) is smaller than (10) if

$$z < 4\left[\frac{b}{a+b^2}\right]^2 \sigma_{high}^2, \tag{12}$$

and the assumption $z < [b/(a+b^2)]^2\sigma_{high}^2$ implies (12). ■

Notice the proof solves the MA's intertemporal problem (9) by considering only a generic single period; this is appropriate because there are no intertemporal tradeoffs.

Proposition 2 tells us that if $[b/(a+b^2)]^2\sigma_{low}^2 \leq z < [b/(a+b^2)]^2\sigma_{high}^2$ the MA commits to (absolutely) low inflation to prevent IA when $\sigma_t^2 = \sigma_{high}^2$; this preserves some of its ability to smooth output. If $[b/(a+b^2)]^2\sigma_{low}^2 > z$ the PS would always acquire information under the fully responsive policy, so the MA would commit to bounded inflation in every period. If $[b/(a+b^2)]^2\sigma_{high}^2 \leq z$ the PS would never acquire information under fully responsive policy so the MA would commit to fully responsive inflation in every period. In this case committed policy entails the same responsiveness of inflation to output shocks as under discretion, but mean inflation is zero instead of $b(k-1)\overline{y}/a$.

3.4 Welfare

Suppose σ_t^2 is small enough that the no-IA constraint in the MA's problem is not binding. Then, as with discretion, a small increase (in the comparative statics sense) in σ_t^2 makes both the PS and the MA worse off *ex ante.* However, if the constraint is binding, a small increase in σ_t^2 does not affect the PS's *ex ante* welfare because the MA limits inflation's variance to z. It makes the MA worse off *ex ante* because inflation's variance is constrained to be the same and output's variance is higher.

As with discretion, higher z is worse for the PS, but better for the MA.

4 Many PS agents with different costs of IA

Suppose there is a continuum $[0, 1]$ of PS agents and they differ in their IA costs, with agent n having IA cost $z_n = \gamma + \delta n$ with $\gamma \geq 0$, $\delta \geq 0$.[5] Now i_t denotes the fraction of agents who are informed in period t. I assume all PS agents affect aggregate outcomes equally and that all agents know the current value of i_t as they form their inflation forecasts,[6] so the AS curve is

$$y_t = \overline{y} + b\pi_t - b\left\{i_t E[\pi_t|I_t \cup i_t \cup u_t] + (1 - i_t)E[\pi_t|I_t \cup i_t]\right\} + u_t. \tag{13}$$

4.1 Discretionary policy

In this section I first describe the Nash equilibria of discretionary policy. Intra-period timing is as follows: First each PS agent decides whether to become informed, then the shock is realized, then the MA optimizes myopically, taking u_t and i_t as given. In REE, $E[\pi_t|I_t \cup i_t \cup$

[5] In what follows I assume $\delta > 0$; the case of $\delta = 0$ is taken up in a separate section below.

[6] This is similar to the Walrasian auctioneer; it is as if all PS agents are together and each one submits a tentative IA choice. Agents change their tentative choices in response to each other's choices, until they coordinate on an equilibrium value of i_t. Equilibrium is non-trivial because, as will become apparent below, an agent's expected payoff of becoming informed depends on the proportion of other agents who are informed.

$u_t] = \pi_t$, so

$$\pi_t = \frac{b}{a}\overline{y}(k-1) - \frac{b}{a+(1-i_t)b^2}u_t \tag{14}$$

and

$$y_t = \overline{y} + \frac{a}{a+(1-i_t)b^2}u_t. \tag{15}$$

Inflation's variance conditional on $I_t \cup i_t$ is

$$\underset{\pi}{\mathrm{var}}(i_t, \sigma_t^2) = \left[\frac{b}{a+(1-i_t)b^2}\right]^2 \sigma_t^2. \tag{16}$$

With many PS agents a new welfare issue arises: Because inflation's variance is an increasing function of i_t, informed agents impose a negative externality on uninformed ones. I return to this point in the section on welfare below.

A Nash equilibrium for date t has two components.[7] One is the PS agents, who move first, coordinating on an equilibrium value of i_t. The second part is the MA's optimal behavior given i_t, equation (14). So describing the set of Nash equilibria requires first describing the PS's optimization problem.

Private agent n's loss is

$$l_t^n = Z_{n,t} + (E[\pi_t|\Omega_{n,t}] - \pi_t)^2 \tag{17}$$

where $Z_{n,t} \in \{0, z_n\}$ is the cost n pays for IA, $Z_{n,t} = 0$ meaning n remains ignorant at t, and $\Omega_{n,t}$ is n's date-t information set. Note that while their IA costs differ, the expected costs of remaining ignorant (given some value of i_t) are the same for all private agents. Agent n minimizes its expectation of (17) conditional on $I_t \cup i_t$, acquiring information iff the associated loss, z_n, is less than the expected loss under ignorance, $\mathrm{var}_\pi(i_t, \sigma_t^2)$.

Definition 2 *A Nash equilibrium for date t consists of MA policy (14) and an IA choice for each private agent n such that*

$$n \textit{ is informed iff } z_n < \underset{\pi}{\mathrm{var}}(i_t, \sigma_t^2) \tag{18}$$

[7] There is also an infinity of overall Nash equilibria in the full intertemporal game since the PS could coordinate on switching between date t equilibria as a function of t (if there are multiple date t equilibria). I assume they do not switch unless a change between σ_{low}^2 and σ_{high}^2 compels it.

where i_t is the implied fraction of the PS that is informed.

The number and nature of date t Nash equilibria vary with parameters and whether $\sigma_t^2 = \sigma_{low}^2$ or $\sigma_t^2 = \sigma_{high}^2$ and I now turn to describing the various cases.

4.1.1 Nash Equilibrium values of i_t

Plainly an equilibrium i_t has all agents $n < i_t$ being informed (e.g., if $i_t = 0.75$ then it is agents 0 to 0.75 who are informed, not some other set of agents with measure 0.75). If in period t there is one $i_t \in (0,1)$ such that $z_{i_t} = \text{var}_\pi(i_t, \sigma_t^2)$ I denote it by i_t^*; if there are two such i_t's I denote them by $i_{t,1}^*$ and $i_{t,2}^*$ with $i_{t,1}^* < i_{t,2}^*$. Then the set of Nash equilibrium values of i_t has a simple description: If there is one such intersection i_t^* it is an equilibrium value of i_t, and if there are two intersections then both, $i_{t,1}^*$ and $i_{t,2}^*$, are equilibrium values of i_t. Also, if $\text{var}_\pi(0, \sigma_t^2) \leq z_0$ then $i_t = 0$ is an equilibrium and if $\text{var}_\pi(1, \sigma_t^2) \geq z_1$ then $i_t = 1$ is an equilibrium. The positive measure cases[8] and their equilibria are:

1. $\text{var}_\pi(i_t, \sigma_t^2)$ may lie above z_{i_t} for all i_t (see Figure 1). In this case the equilibrium has $i_t = 1$.

2. $\text{var}_\pi(i_t, \sigma_t^2)$ may lie below z_{i_t} for all $i_t \in [0,1]$. In this case $i_t = 0$.

3. There may be one intersection i_t^*, with $\text{var}_\pi(i_t, \sigma_t^2)$ lying above z_{i_t} for $i_t \in [0, i_t^*)$ and below z_{i_t} for $i_t \in (i_t^*, 1]$. In this case $i_t = i_t^*$.

4. There may be one intersection at i_t^*, with $\text{var}_\pi(i_t, \sigma_t^2)$ lying below z_{i_t} for $i_t \in [0, i_t^*)$ and above z_{i_t} for $i_t > i_t^*$. In this case the equilibrium values of i_t are 0, i_t^*, and 1.

5. There may be two intersections, so that $\text{var}_\pi(i_t, \sigma_t^2)$ lies above z_{i_t} for $i_t \in [0, i_{t,1}^*)$, below z_{i_t} for $i_t \in (i_{t,1}^*, i_{t,2}^*)$, and above z_{i_t} for $i_t > i_{t,2}^*$. The equilibrium values of i_t are $i_{t,1}^*$, $i_{t,2}^*$, and 1.

[8] I ignore knife-edge cases, e.g., $\text{var}_\pi(i_t, \sigma_t^2)$ being tangent to the line $\gamma + \delta i_t$.

Figure 1 illustrates some possible cases.

FIGURE 1 ABOUT HERE.

Given the "Walrasian auctioneer" approach to determining i_t, the intra-period dynamics are simple; from the starting value of i_t they move to the right wherever $\text{var}_\pi(i_t, \sigma_t^2)$ lies above z_{i_t} and move to the left wherever $\text{var}_\pi(i_t, \sigma_t^2)$ lies below z_{i_t}. If the intra-period dynamics move toward an equilibrium from nearby points that equilibrium is stable and if the dynamics move away it is unstable. One might wish to use stability as an equilibrium selection criterion; in case 4, i_t^* is unstable and in case 5, $i_{t,2}^*$ is unstable.

4.1.2 Comparative Statics

For the topic of information acquisition the main parameters of interest are σ_t^2, γ and δ. The way σ_t^2 affects endogenous variables is more complicated than with a representative agent because i_t is not confined to the set $\{0,1\}$. Equilibrium i_t in the interior of $[0,1]$ are given by solutions to

$$\left[\frac{b}{a+(1-i_t)b^2}\right]^2 \sigma_t^2 = \gamma + \delta i_t. \tag{19}$$

To derive the comparative statics effects of σ_{low}^2 or σ_{high}^2, given that the equilibrium remains an interior one, first find the effect on the equilibrium fraction of informed agents.[9] Totally differentiating (19) gives

$$\frac{\partial i_t}{\partial \sigma_t^2} = \frac{b^2[a+(1-i_t)b^2]}{\delta\left[a+(1-i_t)b^2\right]^3 - 2b^4\sigma_t^2}. \tag{20}$$

This can be signed for any particular interior equilibrium using the equilibrium's stability or instability. For stable interior equilibria, $\text{var}_\pi(\cdot)$ intersects $z(\cdot)$ from above and it can be shown that this implies $\partial i_t/\partial\sigma_t^2$ is positive. In this kind of equilibrium, then, the intrinsic noisiness of the economy, σ_t^2, affects IA in an intuitively sensible way; the worse the noise,

[9] The comparative statics $\frac{\partial \omega}{\partial \sigma_{low}^2}$ and $\frac{\partial \omega}{\partial \sigma_{high}^2}$ for some endogenous variable ω are also the effects on ω of a real-time switch between σ_{low}^2 and σ_{high}^2. This is due to the assumption that PS agents do not switch between equilibria if it is avoidable.

the greater the fraction of agents who acquire information. Also, that $\partial i_t / \partial \sigma_t^2$ is positive implies

$$\frac{\partial \operatorname{var}_\pi(i_t, \sigma_t^2)}{\partial \sigma_t^2} = \delta \frac{\partial i_t}{\partial \sigma_t^2} \tag{21}$$

is positive. In such cases, e.g., case 3, inflation's variance is increasing in σ_t^2 for two reasons. One is the direct effect, due to the fact that the MA sets inflation to offset the shock; the other reason is the indirect effect due to the effect of σ_t^2 on the equilibrium i_t. Both of these effects can be seen by inspection of (14). Making a similar argument for unstable equilibria, and applying similar arguments to $\operatorname{var}_y(i_t, \sigma_t^2)$, output's variance conditional on $I_t \cup i_t$, we have the next proposition.

Proposition 3 *For stable interior equilibria* $\frac{\partial i_t}{\partial \sigma_t^2} > 0$, $\frac{\partial \operatorname{var}_\pi(i_t,\sigma_t^2)}{\partial \sigma_t^2} > 0$, *and* $\frac{\partial \operatorname{var}_y(i_t,\sigma_t^2)}{\partial \sigma_t^2} > 0$. *For unstable interior equilibria* $\frac{\partial i_t}{\partial \sigma_t^2} < 0$, $\frac{\partial \operatorname{var}_\pi(i_t,\sigma_t^2)}{\partial \sigma_t^2} < 0$, *and* $\frac{\partial \operatorname{var}_y(i_t,\sigma_t^2)}{\partial \sigma_t^2} < 0$.

Proposition 3 says the shock's variance can affect inflation's and output's variances in a counterintuitive way, but only in equilibria which are implausible on stability grounds. In stable equilibria, the higher is the shock's variance, the higher is the equilibrium measure of informed agents, and thus the smaller is monetary policy's effect on output. Because of this, the MA sets a higher responsiveness of inflation to the output shock, so inflation's variance is higher. Output's variance is also higher though, because the higher i_t reduces policy's effectiveness. In unstable equilibria, in contrast, an increase in σ_t^2 lowers the measure of informed agents, making policy more effective and thus enabling the MA to simultaneously reduce inflation's variance and output's variance.

One can consider the effects of a discrete difference in σ_{low}^2 or in σ_{high}^2, as well as the infinitesimal differences of derivatives. This allows the analysis of the comparative statics associated with a change in the cases. In particular, i_t may not be in an interior equilibrium for two different values of σ_{low}^2 (or σ_{high}^2), e.g., if for a small value of σ_{low}^2 case 2 pertains and so $i_t = 0$, while for a high value of σ_{low}^2 case 1 pertains and so $i_t = 1$. Then (19) and (21) do not apply, but it is still possible to show the following, the proof of which is deferred to the Appendix:

Proposition 4 *If i_t is not in an interior equilibrium for both of two distinct values of σ^2_{low} (or two distinct values of σ^2_{high}), then $\frac{\Delta i_t}{\Delta \sigma_t^2} \geq 0$, $\frac{\Delta \operatorname{var}_\pi(i_t,\sigma_t^2)}{\Delta \sigma_t^2} > 0$, and $\frac{\Delta \operatorname{var}_y(i_t,\sigma_t^2)}{\Delta \sigma_t^2} > 0$.*

The first inequality is weak because of the possibility that $i_t = 0$ or $i_t = 1$ for both values of σ_t^2. In this case, there is still the direct effect of σ_t^2 on $\operatorname{var}_\pi(\cdot)$ and $\operatorname{var}_y(\cdot)$.

The results thus far can be summarized as follows: i_t changes weakly, and $\operatorname{var}_\pi(i_t, \sigma_t^2)$, and $\operatorname{var}_y(i_t, \sigma_t^2)$ change strictly, in the same direction as σ_t^2 unless the PS agents play an unstable interior equilibrium for both considered values of σ_t^2, in which case the relationships are reversed.

I turn next to comparative statics for the IA cost parameters γ and δ.

Proposition 5 *For stable interior equilibria $\frac{\partial i_t}{\partial \gamma} < 0$ and $\frac{\partial i_t}{\partial \delta} < 0$. It follows that $\frac{\partial \operatorname{var}_\pi(i_t,\sigma_t^2)}{\partial \gamma}$, $\frac{\partial \operatorname{var}_y(i_t,\sigma_t^2)}{\partial \gamma}$, $\frac{\partial \operatorname{var}_\pi(i_t,\sigma_t^2)}{\partial \delta}$, and $\frac{\partial \operatorname{var}_y(i_t,\sigma_t^2)}{\partial \delta}$ are negative. For unstable interior equilibria all six derivatives are positive. For $i_t = 0$ or $i_t = 1$, all six derivatives are zero.*

"Sensible" comparative statics—higher IA costs leading to fewer agents being informed—obtain only in stable interior equilibria.

Finally, there is the possibility $\delta = 0$, i.e., all agents have the same IA cost. This does not reproduce the representative agent version of the model because in that version, perforce, either the entire PS acquires information or none does. The set of equilibrium i_t is simple; there are only three positive-measure cases. One case has $\operatorname{var}_\pi(i_t, \sigma_t^2)$ lying above γ, in which case everyone will be informed, another has $\operatorname{var}_\pi(i_t, \sigma_t^2)$ lying below γ, in which case no one will be informed, and there is a case in which $\operatorname{var}_\pi(i_t, \sigma_t^2)$ intersects z_n once, from below. This case has three equilibria, i_t equal to 0, i_t^*, or 1; only the corner equilibria are stable. One interesting aspect of this case is immediate from (21):

Proposition 6 *If $\delta = 0$, then in the interior equilibrium $\frac{\partial \operatorname{var}_\pi(i_t,\sigma_t^2)}{\partial \sigma_t^2} = 0$.*

Proposition 6 holds because a change in σ_t^2 induces a change in i_t that exactly offsets it. It also implies that a change in σ_t^2 induces no change in the PS's expected losses. Using (19) with $\delta = 0$, one can show the following:

Proposition 7 *If* $\delta = 0$*, then in the interior equilibrium* $\frac{\partial i_t}{\partial \gamma} > 0$*,* $\frac{\partial \operatorname{var}_\pi(i_t, \sigma_t^2)}{\partial \gamma} > 0$*, and* $\frac{\partial \operatorname{var}_y(i_t, \sigma_t^2)}{\partial \gamma} > 0$*.*

Again the counterintuitive result—a rise in IA costs causing more agents to be informed—can occur, but only in an unstable equilibrium.

4.1.3 The effectiveness of monetary policy

The effectiveness of monetary policy, $\partial y / \partial \pi$, is given by

$$\frac{\partial y}{\partial \pi} = b(1 - i_t),$$

which varies between 0 and b. Higher σ_t^2 can lead to lower i_t and therefore more effective policy, but only in unstable interior equilibria; in stable interior equilibria higher σ_t^2 induces higher i_t and therefore less effective policy. The effect of σ_t^2 on policy's effectiveness is discrete if a change in σ_{low}^2 or σ_{high}^2 leads to a change between an interior equilibrium and a corner equilibrium. For example, a small value of σ_{low}^2 may imply case 5 with $i_t = i_{t,1}^*$, while a larger value of σ_{low}^2 may imply case 1 with $i_t = 1$, so the value of σ_{low}^2 makes the difference between policy being somewhat effective and completely ineffective. The dynamic analog of this example would be that a switch from σ_{low}^2 to σ_{high}^2 makes policy suddenly change from somewhat effective to completely ineffective. I next turn to the econometric implications of such switching.

An econometrician estimating a 1960s-era tradeoff of the form $y_t = \alpha + \widehat{b}\pi_t + \widetilde{u}_t$ would seem to detect regime changes in the value of $\widehat{b}$; these would be especially noticeable when a change in σ_t^2 induced a jump between corner equilibria or between a corner equilibrium and an interior equilibrium. Also, the variance of the residuals would not be not constant but autoregressive, and would be positively correlated with the estimated value of $\widehat{b}$ in unstable interior equilibria (due to private agents' IA behavior). Outside of such equilibria, the variance of the residuals would be negatively correlated with the estimated $\widehat{b}$.

An econometrician incorporating survey data on expectations, estimating $y_t = \alpha + \widehat{b}\,(\pi_t - \pi_t^e) + \widetilde{u}_t$ where π_t^e is the PS's average expectation, would generate a $\widehat{b}$ that would be an unbiased estimate of b. However, the econometrician would observe the following:

(1) the variance of the residuals would not be constant but autoregressive,

(2) the expectations data would sometimes be highly correlated with the residuals (i.e., in sequences of periods in which i_t was high) and sometimes not,

(3) the degree of correlation would itself be positively correlated with the residuals' variance if the PS were playing a stable equilibrium, and negatively correlated if the PS were playing an unstable equilibrium,

(4) there would be a perfect multicolinearity problem between expectations and inflation in sequences of periods in which $i_t = 1$.

4.1.4 Welfare

The welfare measure of interest is agents' welfare during their entire lives, not just from living in any particular variance regime. For this reason, it is not necessary to study the expected losses conditional on σ_t^2; the welfare analysis uses the unconditional expectations of the losses.[10] I first take up the effect of σ_t^2.

Proposition 8 *Suppose $\delta = 0$. In interior equilibria $\frac{\partial E(l_t)}{\partial \sigma_t^2} = 0$ and in corner equilibria $\frac{\partial E(l_t)}{\partial \sigma_t^2} > 0$. If $i_t = 0$ then $\frac{\partial E(l_t^n)}{\partial \sigma_t^2} > 0$ for all n. If $i_t > 0$ then $\frac{\partial E(l_t^n)}{\partial \sigma_t^2} = 0$ for all n.*

Proof. See the Appendix. ■

Proposition 9 *If $\delta > 0$, higher σ_t^2 is Pareto superior to lower σ_t^2 within a given unstable interior equilibrium. Otherwise, higher σ_t^2 is Pareto inferior to lower σ_t^2.*

Proof. See the Appendix. ■

For stable interior equilibria in Proposition 9, an increase in σ_t^2 hurts the MA because it induces an increase in the equilibrium variances of inflation and output. For the PS, an agent n who was informed at the lower value of i_t is not affected by the change; n's loss

[10] This distinction is of little consequence anyway, because for a given σ_{high}^2 (respectively, σ_{low}^2), a change in σ_{low}^2 (respectively, σ_{high}^2) affects the conditional and unconditional expected losses in the same direction.

is still z_n. Agents who are uninformed in both situations are *ex ante* worse off, since their expected forecast error is higher with a higher σ_t^2. Note the deterioration in their expected welfare is not only due to the direct effect of higher shock's variance on inflation's variance, but also due to the indirect effect, the externality caused by other agents' IA behavior. PS agents who were uninformed before and are now informed have a change from an expected loss of $\mathrm{var}_\pi(i_t, \sigma_t^2)$ to an expected loss of z_n. Plainly they're worse off; the fact that n was uninformed previously means n's expected loss under ignorance was less than z_n and now it is z_n. Similar intuitions apply to unstable interior equilibria.

I turn next to the welfare effects of γ and δ. For $i_t = 0$ a small increase in γ or δ does not affect either the MA's or the PS's expected loss. For $i_t = 1$ an increase in γ or δ does not affect the MA's expected loss, but it does affect the PS's loss since the cost of becoming informed rises for all agents.[11]

For stable interior equilibria, recall that $\partial i_t/\partial\gamma < 0$ and $\partial i_t/\partial\delta < 0$. It follows that the MA's expected loss is decreasing in δ and in γ since

$$\frac{\partial E(l_t)}{\partial\gamma} = \frac{2ab^2\left[a+b^2\right]\sigma_t^2}{\left[a+(1-i_t)b^2\right]^3}\frac{\partial i_t}{\partial\gamma}, \tag{22}$$

$$\frac{\partial E(l_t)}{\partial\delta} = \frac{2ab^2\left[a+b^2\right]\sigma_t^2}{\left[a+(1-i_t)b^2\right]^3}\frac{\partial i_t}{\partial\delta}. \tag{23}$$

That is, higher IA costs induce fewer informed agents, which improves the MA's ability to stabilize the economy. For unstable interior equilibria the results are reversed.

For the PS welfare results are more complicated. In stable interior equilibria $\partial i_t/\partial\gamma$ and $\partial i_t/\partial\delta$ are negative. Those who are uninformed for both values of the changing parameter, and those who switch from informed to uninformed, have lower expected losses, because inflation's variance is lower for lower i_t. However, those who are informed for both values of the parameter pay higher IA costs. Therefore the effect of the change on the PS cannot be judged by the Pareto standard.

In unstable equilibria an increase in IA costs is a Pareto deterioration. Agents who are informed for both values of the parameter pay higher IA costs, agents who are uninformed

[11] Except that a change in δ leaves agent 0's loss unchanged.

for both values of the parameter have higher expected forecast errors, and agents who switch from uninformed to informed also have higher losses.

4.2 Committed policy

This section contains results on optimal policy with commitment. I assume the MA can observe only the aggregate measure of agents who are informed, and cannot observe individual agents. Therefore an agent will defect from any value of i_t the MA tries to induce if, in that period, the expected loss of remaining ignorant is greater than that agent's IA cost. This is because each agent, having measure zero, knows his behavior will not affect the aggregate behavior the MA observes.

With commitment the MA faces a tradeoff; making inflation more responsive to shocks stabilizes output more through the effect on uninformed agents, but also may induce fewer uninformed agents. This tradeoff makes it necessary for the MA to consider how its inflation rule affects the equilibrium i_t.[12]

This implies a three-step procedure for solving the MA's problem. The first step is to check whether the constraint imposed by the PS's possible IA binds. This involves checking, for each value of σ_t^2, whether the discretionary variance of inflation for $i_t = 0$ is small enough to prevent IA, $[b/(a+b^2)]^2 \sigma_t^2 \leq \gamma$. If this inequality holds for a given value of σ_t^2 the MA is done; it commits, for periods with that value of σ_t^2, to (14) with $i_t = 0$, but with mean inflation equal to zero.

If $[b/(a+b^2)]^2 \sigma_t^2 > \gamma$ for one or both values of σ_t^2 the MA proceeds to the second step. For this step note the MA can induce any $i_t \in [0,1]$ as an equilibrium by committing to a policy with appropriate variance (unless $\delta = 0$, a possibility taken up below). If the MA desires $i_t = \bar{i}$ it simply commits to a policy rule $\pi(i_t, \sigma_t^2, u_t)$ such that $\operatorname{var}\left[\pi(i_t, \sigma_t^2, u_t)|I_t\right] = \gamma + \delta\bar{i}$, no matter what i_t the PS coordinates on at t.[13] That commitment in fact induces the PS

[12] This tradeoff also occurs under discretionary policy, but the MA doesn't consider the tradeoff because it optimizes each period taking that period's i_t as given.

[13] In this section $\operatorname{var}\left[\pi(i_t, \sigma_t^2, u_t)|I_t\right]$ is used to denote the conditional variance of inflation; this new notation emphasizes that under commitment the variance is not necessarily $\operatorname{var}_\pi(i_t, \sigma_t^2)$, the variance under

to coordinate on $i_t = \bar{i}$. Therefore, the MA's second step is to find, for each $\bar{i} \in [0,1]$, the loss-minimizing policy rule such that $\operatorname{var}\left[\pi(i_t, \sigma_t^2, u_t)|I_t\right] = \gamma + \delta\bar{i}$. The third step is to compare the implied expected losses for all $\bar{i}$; the $\bar{i}$ implying the lowest expected loss is the optimum.

It can be shown by arguments similar to those for the representative agent case that the solution to the second-step problem is the policy function[14]

$$\pi(i_t, \sigma_t^2, u_t) = -\sqrt{\frac{\gamma + \delta\bar{i}}{\sigma_t^2}}u_t. \tag{24}$$

(It is instructive to compare this with the corresponding policy for the representative agent PS, $\pi(u_t) = -\sqrt{z/\sigma_{high}^2}u_t$.) In the third step the MA substitutes (24) into its expected loss function and minimizes over $\bar{i}$ subject to $0 \leq \bar{i} \leq 1$. Suppressing the term $\sigma_t^2 + (1-k)^2\overline{y}^2$, maximizing the negative of the loss gives an optimization problem that has at least one solution, since the objective function is a continuous function of the choice variable and is defined over a compact feasible set of that choice variable. The Lagrangian is

$$\max_{\bar{i}} L = 2b\sqrt{\sigma_t^2}(1-\bar{i})\sqrt{\gamma + \delta\bar{i}} - \left[a + b^2(1-\bar{i})^2\right]\left[\gamma + \delta\bar{i}\right] + \mu_1\bar{i} - \mu_2\bar{i} \tag{25}$$

where μ_1 and μ_2 are the Lagrangian multipliers on the constraints $\bar{i} \geq 0$ and $\bar{i} \leq 1$ respectively. It is easy to show $\bar{i} = 1$ is not a solution, so the relevant necessary conditions for a solution are

$$\begin{aligned} 0 &\leq \bar{i} < 1 \qquad (26) \\ \mu_1\bar{i} &= 0 \end{aligned}$$

$$\begin{aligned} \frac{\partial L}{\partial \bar{i}} &= 0 = b\delta(1-\bar{i})\frac{\sqrt{\sigma_t^2}}{[\gamma + \delta\bar{i}]^{1/2}} + 2b^2(1-\bar{i})[\gamma + \delta\bar{i}] \qquad (27) \\ &\quad -2b\sqrt{\sigma_t^2}\sqrt{\gamma + \delta\bar{i}} - \delta b^2(1-\bar{i})^2 - \delta a + \mu_1. \end{aligned}$$

discretionary policy.

[14] Note that if $\bar{i} = 0$ this policy implies an inflation variance of γ, which also the variance implied by the inequality $\left[\frac{b}{a+b^2}\right]^2 \sigma_t^2 \leq \gamma$ when it holds with equality. That is, the MA's two subproblems overlap where PS agent 0 is just indifferent about becoming informed.

It is evident from (27) that the second-order condition for a global optimum will not be satisfied. Furthermore, (27) is not analytically solvable because it is isomorphic to a quintic equation in the variable $\sqrt{\gamma + \delta \bar{i}}$.

However, we can draw the following conclusions. The MA never allows all PS agents to be informed; this is intuitively sensible because in that case monetary policy cannot stabilize output at all. In some regions of the parameter space the MA will conduct policy in a way consistent with some PS agents engaging in IA. In other regions the MA will conduct policy such that there is no IA. In particular, if $\gamma > 0$ and σ_t^2 is small enough, the MA can set inflation as in (5), $\pi(\cdot) = [-b/(a+b^2)]\, u_t$, without triggering any IA.

4.3 Special case: $\gamma = 0$

Suppose $\gamma = 0$, so that agent $n = 0$ incurs no cost of becoming informed and agents $n \approx 0$ incur arbitrarily small costs. For example, these agents may be at the heart of the financial system and so have easy access to economic data. In this case the MA could induce complete ignorance only by having inflation be completely unresponsive to shocks; it is easy to show then that complete ignorance is not optimal.

Proposition 10 *If $\gamma = 0$ and $\delta > 0$ the MA's optimum has $\bar{i} \in (0, 1)$ in every period.*

Proof. For $\gamma = 0$ the step 1 inequality will obviously not be satisfied and the constraint for achieving $\bar{i} = 0$ is $\operatorname{var}[\pi(i_t, \sigma_t^2, u_t)|I_t] \leq 0$, which is identical to $\operatorname{var}[\pi(i_t, \sigma_t^2, u_t)|I_t] = \gamma + \delta \bar{i}$, the constraint that is used in step two above. Therefore the appropriate first-order condition is (27) with $\gamma = 0$:

$$\frac{\partial L}{\partial \bar{i}} = b\sigma_t \sqrt{\delta} \left[\frac{1}{\bar{i}^{1/2}} - 3\bar{i}^{1/2} \right] + \delta b^2 \left[4\bar{i} - 3\bar{i}^2 \right] - \delta a - \delta b^2 + \mu_1.$$

$\lim_{\bar{i} \downarrow 0} \frac{\partial L}{\partial \bar{i}} = \infty$, so $\bar{i} = 0$ is not a solution. As noted above, $\bar{i} = 1$ is not a solution, and there is at least one solution. It follows that the solution is interior. ■

4.4 Special case: $\delta = 0$

If $\delta = 0$, that is, all PS agents have the same cost γ of acquiring information, the MA has two options. By committing to set inflation's variance equal to or less than γ, it can induce $\bar{i} = 0$. If inflation's variance is greater than γ, the equilibrium value of $\bar{i}$ is one. Plainly the MA would never choose to induce $\bar{i} = 1$, so it chooses the optimal policy that induces $\mathrm{var}\left[\pi(i_t, \sigma_t^2, u_t)|I_t\right] \leq \gamma$. Therefore the MA commits to (5) if the variance associated with that policy is no greater than γ, i.e., $\left[b/(a+b^2)\right]^2 \sigma_t^2 \leq \gamma$. Otherwise the MA must restrain inflation's response to the shocks so that its variance is γ, using (24) with $\delta = 0$,

$$\pi(\bar{i}, \sigma_t^2, u_t) = -\sqrt{\frac{\gamma}{\sigma_t^2}} u_t. \tag{28}$$

That this induces $\bar{i} = 0$ is implied by the assumption that an agent indifferent about becoming informed does not do so. If $\delta = \gamma = 0$, all PS agents will be informed unless inflation's variance is zero. In this case the MA can never achieve any stabilization effects, so it simply commits to setting inflation to zero in all periods.

4.5 Welfare

I consider the effects of σ_t^2, δ, and γ. While the comparative statics can be counterintuitive with discretion, that happens only in unstable equilibria, which the MA can rule out with commitment. Plainly if the constraint imposed by the PS's possible IA binds, higher γ makes every PS agent worse off and makes the MA better off, because higher γ allows the MA to increase inflation's variance. For δ the same analysis applies. Higher σ_t^2 makes everyone, PS and MA, worse off, whether or not the constraint binds. This is because even if it does not bind, there is still the direct effect of the shock's variance on economic volatility, which all players dislike. If γ is large enough that the constraint does not bind, a small change in γ or δ do not affect either MA's or the PS's expected losses. A small increase in σ_t^2 increases both the MA's and the PS's expected losses. For $\delta = 0$, if the constraint binds an increase in σ_t^2 does not affect the PS's loss, because the MA is compelled to keep inflation's variance equal to γ (see (28)).

5 Conclusion

This paper has presented results for optimal information acquisition in a setting in which information acquisition affects the effectiveness of monetary policy. Under discretionary policy, the endogenous IA causes a time-varying inflation-output tradeoff, and in fact it tends to weaken monetary policy exactly when monetary policy would be most useful, in high-volatility periods. An increased variance of the shock can cause less IA and can be a Pareto improvement, though such counterintuitive results occur only in unstable equilibria. Under commitment the MA can "manage" policy's effectiveness by committing to low-variance policy rules that limit the PS's incentive to become informed. However, the required restraint also limits the MA's ability to stabilize the economy. Finally, the model also brings out a novel externality: informed agents, by increasing the MA's optimal inflation variance, impose a negative externality on uninformed ones.

6 Literature Cited

Altissimo, F., and G.L. Violante (2001). "The nonlinear dynamics of output and unemployment in the U.S." *Journal of Applied Econometrics* 16, 461-486.

Apergis, Nicholas, and Stephen M. Miller (2004). "Macroeconomic rationality and Lucas's misperceptions model: Further Evidence from Forty-one Countries." *Journal of Economics and Business* 56, 227-241.

Barro, R.J., and D. Gordon (1983). "A positive theory of monetary policy in a natural rate model." *Journal of Political Economy* 91, 589-610.

Blanchard, Olivier J., and Stanley Fischer (1992). Lectures on Macroeconomics. Cambridge, MA.: MIT Press.

Branch, W.A., G.W. Evans, J. Carlson, and B. McGough (2005). "Monetary policy, endogenous inattention, and the volatility tradeoff." Working paper, University of California at Irvine.

Canzoneri, M.B. (1985). "Monetary policy games and the role of private information."

The American Economic Review 75, 1056-1070.

Evans, G.W., and G. Ramey (1992). "Expectation calculation and macroeconomic dynamics." *The American Economic Review* 82, 207-224.

Fischer, Stanley (1977). "Long-term contracts, rational expectations, and the optimal money supply rule." *Journal of Political Economy* 85, 191-205.

Hahm, S. (1987). "Information acquisition in an incomplete information model of the business cycle." *Journal of Monetary Economics* 20, 123-140.

Ho, K.Y., and A.K.C. Tsui (2002). "Conditional volatility asymmetry of business cycles: evidence from four OECD countries." Mimeo, Department of Economics, National University of Singapore.

Lucas, R.E. Jr. (1972). "Expectations and the Neutrality of Money." *The Journal of Economic Theory* 4, 103-124.

Lucas, R.E. Jr. (1973). "Some international evidence on output-inflation tradeoffs." *The American Economic Review* 63, 326-334.

Pike, Ralph W., 1984, Optimization for Engineering Systems. New York: Van Nostrand Reinhold Company.

Romer, Christina D., and David R. Romer (2000). "Federal Reserve information and the behavior of interest rates." *The American Economic Review* 90, 429-457.

Sims, Christopher A., and Tao Zha (2006). "Were there regime switches in US monetary policy?" *The American Economic Review* 96, 54-81.

Walsh, C.E. (1995). "Is New Zealand's reserve bank act of 1989 an optimal central bank contract?" *Journal of Money, Credit, and Banking* 27, 1179-1191.

Woodford, Michael (2002). "Imperfect Common Knowledge and the Effects of Monetary Policy." In P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton University Press, Princeton.

FIGURE 1. Some possible cases with many agents and discretionary policy. The horizontal axis is $n \in [0, 1]$. The dashed line is the IA cost function. The three convex curves are three possible MA best-response functions, giving inflation's variance as a function of $n = i_t$.

7 Appendix

7.1 Proof of Proposition 2

The MA can solve the intertemporal problem by independently solving each period's problem because there are no state variables in one period that are affected by the MA's actions in previous periods. Because the MA can usefully condition date t inflation on u_{t-1} (the PS knows u_{t-1} at the start of date t), the generic problem it solves is $\min E[l_t|I_t]$. Substituting the AS equation and the policy function $\pi_t = \pi(u_t)$ into the MA loss function gives

$$l_t = a\pi(u_t)^2 + \{b\pi(u_t) - bE[\pi(u_t)|I_t] + u_t + (1-k)\overline{y}\}^2 \tag{29}$$

and taking expectations from before u_t is known—since this is a committed policy—gives the expected loss the MA minimizes:

$$E[l_t|I_t] = \sigma_t^2 + (1-k)^2\overline{y}^2 + (a+b^2)E[\pi^2(u_t)|I_t] + 2bE[u_t\pi(u_t)|I_t] - b^2[E[\pi(u_t)|I_t]]^2. \tag{30}$$

The problem, suppressing the term $\sigma_t^2 + (1-k)^2\overline{y}^2$ and the time subscript on u_t, is

$$\min_{\pi(u)\in C^2} E[l_t] \quad = \quad (a+b^2)\int_{-\infty}^{\infty} \pi(u)^2 f_t(u)du + 2b\int_{-\infty}^{\infty} u\pi(u)f_t(u)du \tag{31}$$

$$-b^2\left[\int_{-\infty}^{\infty}\pi(u)f_t(u)du\right]^2.$$

The constraint that the private sector not acquire information, in slack variable form, is

$$\int_{-\infty}^{\infty}\pi(u)^2 f_t(u)du-\left[\int_{-\infty}^{\infty}\pi(u)f_t(u)du\right]^2+s^2=z \tag{32}$$

where s is the complementary slackness variable. The MA chooses a C^2 function $\pi(u)$ to minimize (31) subject to (32). The Lagrangian is

$$\begin{aligned} L &= (a+b^2)\int_{-\infty}^{\infty}\pi(u)^2 f_t(u)du+2b\int_{-\infty}^{\infty}u\pi(u)f_t(u)du \\ &\quad -b^2\left[\int_{-\infty}^{\infty}\pi(u)f_t(u)du\right]^2-\lambda\left[\int_{-\infty}^{\infty}\pi(u)^2 f_t(u)du-\left[\int_{-\infty}^{\infty}\pi(u)f_t(u)du\right]^2\right]. \end{aligned} \tag{33}$$

Following the argument of Pike (1984, Chapter 8), define $\overline{\pi}(u)\equiv\pi(u)+\alpha n(u)$ where $\pi(u)$ is the optimal function, α is an arbitrary parameter and $n(u)$ is an arbitrary C^2 function with $n(-\infty)=n(\infty)=0$. Define

$$\begin{aligned} \Phi(\alpha) &= (a+b^2-\lambda)\int_{-\infty}^{\infty}[\overline{\pi}(u)]^2 f_t(u)du+2b\int_{-\infty}^{\infty}u\overline{\pi}(u)f_t(u)du \\ &\quad +(\lambda-b^2)\left[\int_{-\infty}^{\infty}\overline{\pi}(u)f_t(u)du\right]^2. \end{aligned} \tag{34}$$

The derivative with respect to α is

$$\begin{aligned} \frac{\partial\Phi(\alpha)}{\partial\alpha} &= 2(a+b^2-\lambda)\int_{-\infty}^{\infty}\overline{\pi}(u)n(u)f_t(u)du \\ &\quad +2b\int_{-\infty}^{\infty}un(u)f_t(u)du+2(\lambda-b^2)\int_{-\infty}^{\infty}\overline{\pi}(u)f_t(u)du\int_{-\infty}^{\infty}n(u)f_t(u)du. \end{aligned} \tag{35}$$

Candidate optima are identified by setting $\frac{\partial\Phi(\alpha)}{\partial\alpha}$ equal to zero; also, since by the definition of $\pi(u)$ as the optimal function, $\frac{\partial\Phi(\alpha)}{\partial\alpha}=0$ must be satisfied with $\alpha=0$,

$$\int_{-\infty}^{\infty}\left\{2(a+b^2-\lambda)\pi(u)+2bu+2(\lambda-b^2)\left[\int_{-\infty}^{\infty}\pi(u)f_t(u)du\right]\right\}n(u)f_t(u)du=0. \tag{36}$$

The fundamental lemma of the calculus of variations implies the term in braces is zero, so

$$\pi(u)=\frac{b^2-\lambda}{a+b^2-\lambda}\int_{-\infty}^{\infty}\pi(u)f_t(u)du-\frac{b}{a+b^2-\lambda}u. \tag{37}$$

Multiplying both sides of (37) by $f_t(u)$ and then integrating both sides yields $\int_{-\infty}^{\infty} \pi(u) f_t(u) du = 0$. Putting this back into (37) yields a general solution $\pi(u) = -\frac{b}{a+b^2-\lambda} u$, and one particular solution has $\lambda = 0$, which gives the solution when the constraint doesn't bind, $\pi(u) = -\frac{b}{a+b^2} u$. If the constraint holds with equality, so $s = 0$ and λ is not (necessarily) zero: Substituting $\pi(u) = -\frac{b}{a+b^2-\lambda} u$ into the constraint (32) gives

$$\lambda = a + b^2 \pm b\sqrt{\frac{\sigma_t^2}{z}}. \tag{38}$$

Substituting both possibilities for λ into $\pi(u) = -\frac{b}{a+b^2-\lambda} u$ implies the solution when the constraint binds is

$$\pi(u) = -u\sqrt{\frac{z}{\sigma_t^2}}. \tag{39}$$

Thus the optimal policy is: In periods in which the constraint doesn't bind,

$$\pi(u) = -\frac{b}{a+b^2} u \tag{40}$$

and in periods in which the constraint binds,

$$\pi(u) = -u\sqrt{\frac{z}{\sigma_t^2}}. \tag{41}$$

The fact that this is the only solution to the first-order necessary conditions implies the uniqueness of the Nash equilibrium asserted in the Proposition. The second order condition for minimization is satisfied because the problem is to minimize a quadratic form.

7.2 Proof of Proposition 4

If i_t does not start and end in the interior of $[0,1]$, then (19) and (21) do not apply and we have the more general equations

$$\Delta \underset{\pi}{\mathrm{var}}(i_t, \sigma_t^2) = \frac{b^2}{[a + (1-i_t)b^2]^2} \Delta \sigma_t^2 + \frac{2b^4 \sigma_t^2}{[a + (1-i_t)b^2]^3} \Delta i_t \tag{42}$$

$$\Delta \underset{y}{\mathrm{var}}(i_t, \sigma_t^2) = \frac{a^2}{[a + (1-i_t)b^2]^2} \Delta \sigma_t^2 + \frac{2a^2 b^2 \sigma_t^2}{[a + (1-i_t)b^2]^3} \Delta i_t. \tag{43}$$

One possibility is that the equilibrium is a corner, $i_t = 0$ or $i_t = 1$, which is not affected by a change in σ^2_{low} (or σ^2_{high}), so that $\Delta i_t / \Delta \sigma^2_t = 0$. It can be shown that for both $i_t = 0$ and $i_t = 1$, $\frac{\Delta \operatorname{var}_\pi(i_t,\sigma^2_t)}{\Delta \sigma^2_t}$ and $\frac{\Delta \operatorname{var}_y(i_t,\sigma^2_t)}{\Delta \sigma^2_t}$ are positive. Another possibility is that i_t jumps between an interior equilibrium, i^*, and a corner, so that the change in i_t is either $\pm i^*$ or $\pm(1 - i^*)$. Finally, i_t might jump back and forth between zero and one. In all of these situations that are not governed by (20), it can be shown that $\operatorname{sign}(\Delta i_t) = \operatorname{sign}(\Delta \sigma^2_t)$ (if $\Delta i_t \neq 0$), so output's and inflation's variances change in the same direction as the shock's variance. ■

7.3 Proof of Proposition 5

For interior equilibria, differentiate (19) with respect to γ:

$$\frac{\partial i_t}{\partial \gamma} = \frac{\left[a + (1 - i_t)b^2\right]^3}{2b^4\sigma^2_t - \delta\left[a + (1 - i_t)b^2\right]^3}$$

and with respect to δ:

$$\frac{\partial i_t}{\partial \delta} = i_t \frac{\partial i_t}{\partial \gamma}.$$

working with these two derivatives, one can easily show Proposition 5.

7.4 Proof of Proposition 7

$\delta = 0$ implies equation (19) is a quadratic in i_t with only the "minus" solution potentially in the unit interval. The effect of a change in γ on i_t is

$$\frac{\partial i_t}{\partial \gamma} = \frac{\left[a + (1 - i_t)b^2\right]^3}{2\sigma^2_t b^4} > 0.$$

Therefore, using (16) and (15), $\frac{\partial \operatorname{var}_\pi(i_t,\sigma^2_t)}{\partial \gamma} > 0$ and $\frac{\partial \operatorname{var}_y(i_t,\sigma^2_t)}{\partial \gamma} > 0$.

7.5 Proof of Proposition 8

First consider the effect on the MA's expected loss, $\frac{\partial E(l_t)}{\partial \sigma^2_t}$. It has been shown that in the interior equilibrium $\frac{\partial \operatorname{var}_\pi(i_t,\sigma^2_t)}{\partial \sigma^2_t} = 0$ (Proposition 6). Using (15), we have

$$\frac{\partial \operatorname{var}_y}{\partial \sigma^2_t} = \left[\frac{a}{a + (1 - i_t)b^2}\right]^2 + \frac{2a^2b^2\sigma^2_t}{\left[a + (1 - i_t)b^2\right]^3}\frac{\partial i_t}{\partial \sigma^2_t}. \tag{44}$$

In corner equilibria, then,

$$\frac{\partial \operatorname{var}_y}{\partial \sigma_t^2} = \left[\frac{a}{a+(1-i_t)b^2}\right]^2$$

and in interior equilibria, substituting (20) into (44) gives

$$\frac{\partial \operatorname{var}_y}{\partial \sigma_t^2} = 0.$$

Therefore $\frac{\partial E(l_t)}{\partial \sigma_t^2} > 0$ in corner equilibria and $\frac{\partial E(l_t)}{\partial \sigma_t^2} = 0$ in interior equilibria.

For the PS, the result for $i_t = 0$ follows from the fact that $\frac{\partial \operatorname{var}_\pi(i_t,\sigma_t^2)}{\partial \sigma_t^2} > 0$ in that case. The result for $i_t \in (0,1)$ follows from Proposition 6, and the result for $i_t = 1$ follows because n's loss in that case is always z_n.

7.6 Proof of Proposition 9

An uninformed agent n's expected loss is $E(l_t^n) = \operatorname{var}_\pi(i_t, \sigma_t^2)$ and the MA's loss in REE is

$$l_t = \left[\frac{a+b^2}{a}\right]\overline{y}^2(k-1)^2 - \frac{2\left[a+b^2\right]\overline{y}(k-1)}{a+(1-i_t)b^2}u_t + \frac{a\left[a+b^2\right]}{\left[a+(1-i_t)b^2\right]^2}u_t^2.$$

The MA's unconditional expected loss is

$$\begin{aligned} E(l_t) \;&=\; \left[\frac{a+b^2}{a}\right]\overline{y}^2(k-1)^2 \\ &\quad + a\left[a+b^2\right]\left\{P_{low}E_{low}\left[\frac{\sigma_{low}^2}{\left[a+(1-i_t)b^2\right]^2}\right] + (1-P_{low})E_{high}\left[\frac{\sigma_{high}^2}{\left[a+(1-i_t)b^2\right]^2}\right]\right\} \end{aligned}$$

where P_{low} is the unconditional proportion of the time that $\sigma_t^2 = \sigma_{low}^2$, E_{low} denotes expectations conditional on $\sigma_t^2 = \sigma_{low}^2$, and E_{high} denotes expectations conditional on $\sigma_t^2 = \sigma_{high}^2$. The expectation operators on the RHS are actually unnecessary, since everything they refer to is known, conditional on knowing the variance (and the i_t the PS plays for a given variance). So

$$\frac{\partial E(l_t)}{\partial \sigma_{low}^2} = a\left[a+b^2\right]P_{low}\left\{\frac{2b^2\sigma_{low}^2}{\left[a+(1-i_t(\sigma_{low}^2))b^2\right]^3}\frac{\partial i_t(\sigma_{low}^2)}{\partial \sigma_{low}^2} + \frac{1}{\left[a+(1-i_t(\sigma_{low}^2))b^2\right]^2}\right\} \tag{45}$$

$$\frac{\partial E(l_t)}{\partial \sigma_{high}^2} = a\left[a+b^2\right](1-P_{low})\left\{\frac{2b^2\sigma_{high}^2}{\left[a+(1-i_t(\sigma_{high}^2))b^2\right]^3}\frac{\partial i_t(\sigma_{high}^2)}{\partial \sigma_{high}^2} + \frac{1}{\left[a+(1-i_t(\sigma_{high}^2))b^2\right]^2}\right\} \tag{46}$$

There are several cases if $\delta > 0$:

(1) If $i_t = 0$ initially and a change in σ_t^2 does not affect i_t, then $\frac{\partial E(l_t)}{\partial \sigma_t^2} > 0$, and because $\frac{\Delta \operatorname{var}_\pi(i_t, \sigma_t^2)}{\Delta \sigma_t^2} > 0$ (see Proposition 4), it follows that $\frac{\partial E(l_t^n)}{\partial \sigma_t^2} > 0$ for all $n \in [0, 1]$.

(2) If $i_t = 1$ initially and a change in σ_t^2 does not affect i_t, then $\frac{\partial E(l_t)}{\partial \sigma_t^2} > 0$ and $\frac{\partial E(l_t^n)}{\partial \sigma_t^2} = 0$ for all $n \in [0, 1]$.

(3) For stable interior equilibrium values of i_t, recall $\frac{\partial i_t}{\partial \sigma_t^2} > 0$ (Proposition 3), and this implies $\frac{\partial E(l_t)}{\partial \sigma_{low}^2} > 0$ and $\frac{\partial E(l_t)}{\partial \sigma_{high}^2} > 0$. For the PS: An agent n who was informed at the lower value of i is not affected one way or another by the change; n's loss is still z_n. Agents who are uninformed in both situations are *ex ante* worse off, since their expected forecast error is higher with a higher σ_t^2 (Proposition 3). PS agents who were uninformed before and are now informed have a change from an expected loss of $\operatorname{var}_\pi(i_t, \sigma_t^2)$ to an expected loss of z_n. They're worse off; that agent n was uninformed previously means n's loss under ignorance was less than z_n and now it is z_n. A similar analysis applies for situations as in Proposition 4, except that the effect on i_t is $\frac{\Delta i_t}{\Delta \sigma_t^2} \geq 0$ instead of $\frac{\partial i_t}{\partial \sigma_t^2} > 0$.

(4) For unstable interior equilibria, since $\frac{\partial i_t}{\partial \sigma_t^2} < 0$, an increase in σ_t^2 benefits the MA (equation (20), Proposition 3, and equation (46)). PS agents who are informed at both values of i have losses that are unaffected by σ_t^2. Agents who are uninformed in both situations are *ex ante* better off, since their expected forecast error is lower for higher σ_t^2. PS agents who are informed at a lower value of σ_t^2 and are uninformed at a higher value of σ_t^2 have a change from an expected loss of z_n to an expected loss of $\operatorname{var}_\pi(i_t, \sigma_t^2)$. By an argument symmetric to that in the previous case, they're better off.