Uncertainty, Policy Ineffectiveness and Long Stagnation of the Macroeconomy

A Discussion by Laura Veldkamp
New York University, Stern School of Business

August 7, 2004

When reform has no direct aggregate effects, but only changes agents' incentives to act, uncertainty can weaken its effect. Because agents learn about policy change by observing their signals, increasing signal noise makes it harder for agents to know that reform has occurred. When policy changes are difficult to observe, they change agents' actions less and have less aggregate effect. This is the main point illustrated in the model by Aoki and Yoshikawa.

1 Key Features of the Model

To highlight the model's mechanism, I will examine the origin of the transition rates put forward in equations (4) and (5). What are the decisions agents are making that cause this fraction of them to choose high or low output each period? In the process, we will revisit the question of what particular kind of uncertainty the paper is talking about, and which assumptions about the environment are crucial.

A large number of agents choose one of two actions: high output $y^\star$ or low output y. Each agent's choice about which action to take depends on a signal they observe about the profitability of each action and the actions that other agents are taking. The payoff to $y^\star$ is $V_1(\bar{y})$; the payoff to y is $V_2(\bar{y})$, where $\bar{y}$ is the average output per capita. Since $\bar{y}$ depends on the output decisions of other agents, it represents a production externality. Agents here are playing a form of coordination game. If others choose high output, then choosing high output has a higher expected payoff.

The agent chooses $y^\star$ iff $E_i[V_1(\bar{y})] > E_i[V_2(\bar{y})]$. However, $V_1(\bar{y})$ and $V_2(\bar{y})$ are not known. Only noisy signals of their values are observed:

$$\hat{V}_{1i}(\bar{y}) = V_1(\bar{y}) + \epsilon_{1i}$$

$$\hat{V}_{2i}(\bar{y}) = V_2(\bar{y}) + \epsilon_{2i}$$

where $\epsilon_{2i} - \epsilon_{1i}$ is distributed i.i.d. with a Gibbs distribution.[1]

What role does government policy have in this setting? The government can affect the relative payoffs from choosing high versus low output. In other words, it can shift $g(\bar{y}) = V_1(\bar{y}) - V_2(\bar{y})$. This policy has no direct effect on output. It simply changes agents' incentives by shifting the distribution of their signals. The assumption that there is no direct effect of policy on output is important. It means that agents must learn about the policy, through their signals, for the policy to have any effect. It is this learning process that is inhibited by uncertainty.

[1]The particular distribution used here is not crucial. The Gibbs distribution just allows aggregate transition rates to be succinctly specified.

Given the signal distributions and decision rules, individual and aggregate transition rates from y to $y^\star$ and vice-versa are specified. The fraction of the population that chooses high output $y^\star$ is

$$\bar{x} = P(\hat{V}_{1i}(\bar{y}) > \hat{V}_{2i}(\bar{y})) = P(\epsilon_{2i} - \epsilon_{1i} < V_1(\bar{y}) - V_2(\bar{y})) \\ = [1 + e^{-\beta g(\bar{y})}]^{-1}$$

where $g(\bar{y}) = V_1(\bar{y}) - V_2(\bar{y})$ and the second line comes from the cumulative Gibbs distribution function.

From the expression for $\bar{x}$, we can see why it is that policies are ineffective when uncertainty is high. First, realize that β is inversely related to the variance of the signal noise. If beta is zero, then the cumulative distribution is 1/2 everywhere, meaning that probability is spread uniformly over the real line. This is a distribution with infinite variance. If beta is infinite, then the cumulative distribution switches from 0 to 1 at zero, meaning that all probability mass is concentrated at zero. A distribution with all probability mass at one point has zero variance. Next, note that as β approaches zero, movements in $g(\bar{y})$ have less and less effect on $\bar{x}$. This is the sense in which uncertainty (low β), makes policy (changes in $g(\bar{y})$ ineffective.

Clearly a low value of β increase the variance of the signal noise. However, it is not clear that this actually represents high uncertainty. There is fundamental uncertainty about $(V_1(\bar{y}))$ and $(V_2(\bar{y}))$, and strategic uncertainty about $\bar{y}$. As β falls, strategic uncertainty diminishes. The reason is clear in the extreme case where signal variance is infinite. When signals are completely uninformative, agents ignore them. If agents ignore their private signals about fundamentals, then there is no strategic uncertainty. Common knowledge is effectively restored; $\bar{y}$ is known. In the less extreme case where variance becomes very large, for any finite cutoff signal, 1/2 of agents will be above that signal and 1/2 below. If the form of V is known, and only $\bar{x}$ is unknown, uncertainty disappears in the limit. You should ignore your signal and just play a simultaneous move game.

Information in this model comes from two sources, signals and knowledge about other players' strategies. Uncertainty is non-monotonic in β because as highly variable signals are making agents more uncertain, they are also becoming less uncertain about the aggregate actions of other agents.

Of course, the policy ineffectiveness conclusion still holds. This is because government policy works through signals, not though knowledge of other agents' action. If signals are disregarded, policy is ineffective. But, the effect is really an effect that works through uninformative signals, not uncertainty, per se.

2 Recasting the Model in the Language of Global Games

Aoki and Yoshikawa's model has as key features agents who observe heterogeneous signals about aggregate actions, and positive externalities. This puts their model squarely in the class of models referred to as 'global games.'[2] Looking at their mechanism within a typical global games framework simplifies the setup considerably and sheds some light on the importance of prior beliefs.

Consider the following alternative setup for the model. There is a continuum of agents indexed by i. Their objective function captures two competing incentives: output y_i should be close to an unknown target (z) and close to the average output of others $\bar{y}$.

$$max_{y_i} E[\bar{y} - (y_i - z)^2 - (y_i - \bar{y})^2]$$

This formulation of utility captures the idea in the original model that there are positive externalities in production, and that there is an unknown fundamental that agents are learning about.

[2]See e.g. Morris and Shin, 1998; Angeletos, Hellwig and Pavan, 2003; Angeletos and Pavan, 2004.

The information structure gives all agents a common prior belief $z \sim N(\hat{z}, \sigma_z^2)$, and endows each agent with a private signal $s_i \sim N(z, \sigma_s^2)$.

Government policy can change z. Just as in the original model, this affects agents incentives to produce, without directly changing the quantity produced. Uncertainty about the fundamental is increasing in the signal variance, σ_s^2.

Taking a first-order condition of the objective yields the following optimal output

$$y_i = \frac{1}{2}E_i[z] + \frac{1}{2}E_i[\bar{y}]$$

Since each agent knows that other agents are solving the same problem and have the same optimal strategy, each can forecast what the median other agent will do. Given this relationship, I can solve for $E_i[\bar{y}]$ by guess and verify.

$$E_i[\bar{y}] = \frac{\sigma_s^2}{\sigma_s^2 + \sigma_z^2}\ \hat{z} + (1 - \frac{\sigma_s^2}{\sigma_s^2 + \sigma_z^2})\ s_i$$

This yields the following solution for each agent's production decision as a function of the common prior belief of all agents and the agent's individual signal.

$$y_i = \frac{\sigma_z^2 s_i + 2\sigma_s^2 \hat{z}}{\sigma_z^2 + 2\sigma_s^2}$$

Now we can ask what the effect of a government policy that changes z is. This is a change in the mean of the s_i distribution. Looking at how y_i reacts to changes in signals, we see that as σ_s^2 rises, government policy that works through private signals becomes less effective. So, this simple model can replicate the finding of Aoki and Yoshikawa.

This setup makes clear that there are actually two sources of information that we should be considering: prior beliefs and signals. If the government could influence prior beliefs, the problem of policy ineffectiveness would be solved. Changing $\hat{z}$ and shifting the distribution of s_i both by Δz would increase every agent's action by Δz, regardless of the variances of priors and signals. This tells us that the solution to the problem of policy ineffectiveness is a rather straightforward one: Announce the reform.

3 Conclusions

What this paper tells us is that when agents learn about policy changes only through private signals, and when policy changes change agents incentives have no direct effect on output, policy change will be less effective when signal noise increases. If the government can communicate its change in policy to agents in some public way, policy would retain its effectiveness. The question this raises is, "What kinds of information frictions prevent governments from credibly communicating policy changes?" We read the same newspapers and watch the same TV broadcasts. My own work (Veldkamp 2004) shows that it is optimal for agents to consume largely the same bundle of signals because of how information markets price those signals. Where is all this inside information about the state of the economy coming from? It would be useful to know more about the realistic counterparts to the information structure they have in mind.

In this setting, signal noise makes policy ineffective, regardless of the state of the business cycle. This raises the question whether uncertainty can prolong booms as well. In a business cycle model, there is good reason to focus on a stagnation as a high-uncertainty time. Work done by Chamley and Gale (1994), Veronesi (1999), and Van Nieuwerburgh and Veldkamp (2003) suggests that there is more uncertainty when emerging from a recession and this can restrain the recovery. These predictions, as well as accompanying empirical evidence suggest that heterogeneous information effects might contribute to business cycle asymmetry.

Finally, the learning is the model is static but the state in the model is dynamic. Agents should observe some utility flow along the way. Since payoffs depend on the true state and others' actions, utility realizations are new sources of information that should be explicitly incorporated into beliefs. If aggregate macro variables perfectly reveal government reform the period after it takes place, stagnation will not last very long at all.

References

[1] Angeletos, George-Marios; Hellwig, Christian; Pavan, Alessandro. "Coordination and Policy Traps." NBER Working Paper No. w9767, June, 2003.

[2] Angeletos, George-Marios; Pavan, Alessandro. "Transparency of Information and Coordination in Economies with Investment Complementarities." *American Economic Review, Papers and Proceedings* May, 2004.

[3] C. Chamley, D. Gale, Information Revelation and Strategic Delay in a Model of Investment, *Econometrica*, 62:5 (1994), 1065-1085.

[4] Morris, Stephen and Shin, Hyun Song. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review*, v.88(3), 1998, pp.587-97.

[5] Van Nieuwerburgh, Stijn and Veldkamp, Laura. "Learning Asymmetries in Real Business Cycles." NYU Stern Working Paper, 2003.

[6] Veronesi, Pietro. "Stock Market Overreactions to Bad News in Good Times: A Rational Expectations Equilibrium Model." *Rev. Finan. Stud.*, 12:5 (1999), 975-1007.

[7] Veldkamp, Laura. "Media Frenzies in Markets for Financial Information." NYU Stern working paper, 2004. pp. 695-732.