

Robust Control Made Simple:

Lecture Notes¹

Lars E.O. Svensson

First draft: March 2000. This version: February 2007

1 Introduction

Some of the recent literature on robust control in economics (for instance, Giannoni [3], Hansen and Sargent [4] and [5], Onatski [6], Onatski and Stock [7], Sargent's [10] comment on Ball, [1] and Stock's [12] comment on Rudebusch and Svensson [9]) is somewhat technical and difficult to see through. This note attempts to use a simple example of optimal monetary policy to convey the gist of robust control.

Let $m \in M$ denote a particular model of the transmission mechanism of monetary policy, where M denotes a feasible set of models. Let $f \in F$ denote a particular monetary policy, where F is the feasible set of policies. Let V be an (expected) loss function, so $V(f, m)$ denotes the (expected) loss of policy f in model m .

A Bayesian approach to optimal policy under model uncertainty, as in Brainard [2], assigns a prior probability measure, Φ , on the feasible set of models M . The expected loss for given policy is then

$$E_M V(f, m) \equiv \int_{m \in M} V(f; m) d\Phi(m).$$

The optimal policy f^* minimizes the expected loss,

$$f^* = \arg \min_{f \in F} E_M V(f; m).$$

The optimal policy will be a function of M , Φ and V , $f^*(M, \Phi, V)$.

A robust-control approach to optimal policy under model uncertainty, as in Hansen and Sargent [4] and Onatski and Stock [7], does not assign any prior probability measure on the feasible set of models. Instead, it focuses on the maximum loss for any given policy f ,²

$$\max_{m \in M} V(f; m).$$

¹ I thank Christopher Sims and Michael Woodford for comments, and Annika Andreasson for secretarial and editorial assistance.

² If $\max_{m \in M} V(f; m)$ does not exist, it is replaced by the least upper bound, $\sup_{m \in M} V(f; m)$.

The optimal policy \hat{f} is then the policy that minimizes the maximum loss,

$$\hat{f} = \arg \min_{f \in F} \max_{m \in M} V(f; m).$$

Under robust control, the worst possible model for a given policy, $\hat{m}(f)$, defined by

$$\hat{m}(f) = \arg \max_{m \in M} V(f; m),$$

is crucial, and the optimal policy will fulfill

$$\hat{f} = \arg \min_{f \in F} V(f; \hat{m}(f)).$$

The optimal policy will be a function of M and V , $\hat{f}(M, V)$.

I use a simple model to illustrate these concepts.

2 Optimal policy with known model

Consider a simple backward-looking Phillips curve,

$$\pi_{t+1} = \gamma\pi_t + \alpha x_t + \varepsilon_{t+1}, \tag{2.1}$$

where π_t and x_t is inflation and the output gap, respectively, in period t , the coefficients α and γ are positive, and ε_t is an iid “cost-push” shock with $E[\varepsilon_t] = 0$ and $\text{Var}[\varepsilon_t] = \sigma_\varepsilon^2$.

Let the period loss function be

$$L_t = \pi_t^2 + \lambda x_t^2, \tag{2.2}$$

where $\lambda \geq 0$ is the relative weight on output-gap stability, and let the social loss function be

$$E[L_t] = E[\pi_t^2] + \lambda E[x_t^2] = \text{Var}[\pi_t] + \lambda \text{Var}[x_t], \tag{2.3}$$

where the last equality follows if we restrict policy to result in $E[\pi_t] = E[x_t] = 0$.

For simplicity, we consider the output gap as the central bank’s control variable and note that inflation is an endogenous predetermined variable. We first want to find the optimal policy when α and γ are known. The optimal policy will be a feedback on the only predetermined variable, π_t ,

$$x_t = -f\pi_t, \tag{2.4}$$

where the response coefficient f remains to be determined. Inflation is then given by

$$\pi_{t+1} = (\gamma - \alpha f)\pi_t + \varepsilon_t.$$

We realize that the optimal policy will fulfill $0 \leq f \leq \frac{\gamma}{\alpha}$.

The variances of inflation and the output gap then fulfill

$$\text{Var}[x_t] = f^2 \text{Var}[\pi_t], \quad (2.5)$$

$$\text{Var}[\pi_t] = \frac{1}{1 - (\gamma - \alpha f)^2} \sigma_\varepsilon^2, \quad (2.6)$$

and the social loss will equal

$$\mathbb{E}[L_t] = \frac{1 + \lambda f^2}{1 - (\gamma - \alpha f)^2} \sigma_\varepsilon^2 \equiv V(f; \alpha, \gamma; \lambda). \quad (2.7)$$

The optimal policy $f = f^*(\alpha, \gamma; \lambda)$ is then given by

$$f^*(\alpha, \gamma; \lambda) \equiv \arg \min_f V(f; \alpha, \gamma; \lambda).$$

It will fulfill the first-order condition

$$V_f(f^*; \alpha, \gamma; \lambda) \equiv 2 \frac{\lambda f^* [1 - (\gamma - \alpha f^*)^2] - \alpha (1 + \lambda f^{*2}) (\gamma - \alpha f^*)}{[1 - (\gamma - \alpha f^*)^2]^2} = 0 \quad (2.8)$$

and the second-order condition

$$V_{ff}(f^*; \alpha, \gamma) > 0.$$

It is easy to show (see appendix A) that the partial derivatives of the optimal policy with respect to α , γ and λ fulfill

$$f_\alpha^*(\alpha, \gamma; \lambda) < 0,$$

$$f_\gamma^*(\alpha, \gamma; \lambda) > 0,$$

$$f_\lambda^*(\alpha, \gamma; \lambda) < 0.$$

Furthermore, we have

$$f^*(\alpha, \gamma; 0) = \frac{\gamma}{\alpha} \text{ and } f^*(\alpha, \gamma; \infty) \equiv \lim_{\lambda \rightarrow \infty} f^*(\alpha, \gamma; \lambda) = 0.$$

3 Optimal policy under robust control

Let us now introduce some model uncertainty. Assume that we know the supports of the coefficients α and γ , but not their (constant) true values.³ (Thus, the true model is assumed to remain the same over time. Appendix B examines the case when the true model can change over time.) Let

³ Söderström [11] and Svensson [14] examine the consequences of Brainard-type Bayesian uncertainty in the model used in Svensson [13], which has an aggregate-demand equation added to the Phillips curve (2.1).

the supports be given by $\alpha \in [\alpha_1, \alpha_2]$ and $\gamma \in [\gamma_1, \gamma_2]$, where $0 < \alpha_1 < \alpha_2$ and $0 < \gamma_1 < \gamma_2$. Furthermore, I restrict the bounds of the supports to fulfill

$$\alpha_2 \gamma_2 - \alpha_1 \gamma_1 < \alpha_1 + \alpha_2. \quad (3.1)$$

As explained below, this condition is sufficient for the existence of a policy that results in bounded inflation and output-gap variance for all possible models. The condition can be written

$$\gamma_2 < 1 + (1 + \gamma_1) \frac{\alpha_1}{\alpha_2}$$

and implies that γ_2 must not be too large and is fulfilled for any $\gamma_2 \leq 1$.

Thus, the models can be indexed by $m = (\alpha, \gamma)$ and the feasible set of models M is given by

$$M \equiv \{m = (\alpha, \gamma) \mid \alpha_1 \leq \alpha \leq \alpha_2, \gamma_1 \leq \gamma \leq \gamma_2\}.$$

Let the feasible set of policies F be given by (2.4) with $f \geq 0$, that is, $F = \{f \mid f \geq 0\}$. Selecting the optimal policy under robust control means committing initially to a policy $\hat{f} \in F$ so as to minimize the social loss for the worst possible model; the latter being the one that results in maximum social loss for a given policy. That is,

$$\hat{f} = \arg \min_{f \in F} \max_{(\alpha, \gamma) \in M} V(f; \alpha, \gamma; \lambda).$$

The restriction $f \geq 0$ will not be binding, as we shall see below.

3.1 The worst possible model

From (2.7) follows that, for any given f , the worst possible model is the one that maximizes $|\gamma - \alpha f|$. Let $\hat{m}(f)$ denote the worst possible model for any given $f \in F$, that is,

$$\hat{m}(f) \equiv \arg \max_{(\alpha, \gamma) \in F} |\gamma - \alpha f|.$$

We have

$$|\gamma - \alpha f| = \begin{cases} \gamma - \alpha f & \text{if } \gamma - \alpha f \geq 0, \\ \alpha f - \gamma & \text{if } \gamma - \alpha f < 0. \end{cases}$$

Because $f \geq 0$, we realize that the maximum of $|\gamma - \alpha f|$ for a given f is the maximum of $\gamma_2 - \alpha_1 f$ and $\alpha_2 f - \gamma_1$,

$$\max_{(\alpha, \gamma)} |\gamma - \alpha f| = \max(\gamma_2 - \alpha_1 f, \alpha_2 f - \gamma_1).$$

Thus, $\hat{m}(f)$ is either (α_1, γ_2) or (α_2, γ_1) , depending on whether $\gamma_2 - \alpha_1 f \geq \alpha_2 f - \gamma_1$, that is, whether $f \leq \bar{f}$, where

$$\bar{f} \equiv \frac{\gamma_1 + \gamma_2}{\alpha_1 + \alpha_2} > 0. \quad (3.2)$$

It follows that $\hat{m}(f)$ is given by

$$\hat{m}(f) = \begin{cases} (\alpha_1, \gamma_2) & \text{if } f \leq \bar{f}, \\ (\alpha_2, \gamma_1) & \text{if } f \geq \bar{f}. \end{cases} \quad (3.3)$$

We note that the worst possible model is on the boundary of feasible set of models, and hence depends crucially on the assumed feasible set of models. We also note that the worst possible model is unique if $f \neq \bar{f}$ but can take either of the two values in (3.3) if $f = \bar{f}$.

Now we can see why I imposed the condition (3.1); this condition insures that

$$\max_{(\alpha, \gamma) \in M} |\gamma - \alpha \bar{f}| < 1,$$

which is sufficient for the existence of a policy that results in finite inflation and output-gap variance for all feasible models.

3.2 Strict inflation targeting

The simplest case is strict inflation targeting, $\lambda = 0$, in which case

$$V(f; \alpha, \gamma, 0) \equiv \frac{1}{1 - (\gamma - \alpha f)^2} \sigma_\varepsilon^2.$$

The optimal policy under strict inflation targeting and robust control is then given by

$$\hat{f} = \arg \min_{f \in F} \max_{(\alpha, \gamma) \in M} |\gamma - \alpha f| = \arg \min_{f \in F} \max(\gamma_2 - \alpha_1 f, \alpha_2 f - \gamma_1).$$

In order to minimize $\max(\gamma_2 - \alpha_1 f, \alpha_2 f - \gamma_1)$, it is best to select $f = \hat{f}$ where \hat{f} fulfills

$$\gamma_2 - \alpha_1 \hat{f} = \alpha_2 \hat{f} - \gamma_1,$$

that is,

$$\hat{f} = \bar{f},$$

where \bar{f} is given by (3.2).

3.3 Flexible inflation targeting

Let us now consider the case of flexible inflation targeting, $\lambda > 0$, when $V(f; \alpha, \gamma; \lambda)$ is given by (2.7). Since the term λf^2 in the numerator adds to the social loss, we realize that, *for sufficiently large* λ , it is optimal to lower f below \bar{f} , and to trade off some increase in the denominator for a reduction in the numerator. For $f < \bar{f}$, we then have $\hat{m}(f) = (\alpha_1, \gamma_2)$, hence we have

$$\hat{f} = \arg \min_{f \in F} \max_{(\alpha, \gamma) \in M} V(f; \alpha, \gamma; \lambda) = \arg \min_{f \in F} V(f; \alpha_1, \gamma_2; \lambda) = f^*(\alpha_1, \gamma_2; \lambda),$$

where $f^*(\alpha_1, \gamma_2; \lambda)$ denotes the optimal policy if the true model is known to be (α_1, γ_2) . Thus, the optimal policy for an arbitrary λ , $\hat{f}(\lambda)$, is given by

$$\hat{f}(\lambda) = \min[f^*(\alpha_1, \gamma_2; \lambda), \bar{f}].$$

Note that, for small λ , if (α_1, γ_2) were the true model, the optimal policy would imply $f^*(\alpha_1, \gamma_2; \lambda) > \bar{f}$. Indeed, we know that $f^*(\alpha_1, \gamma_2; 0) = \frac{\gamma_2}{\alpha_1} > \bar{f} = \frac{\gamma_1 + \gamma_2}{\alpha_1 + \alpha_2}$. This is because, if (α_1, γ_2) were the true model, one would not need to consider that the worst possible model shifts to (α_2, γ_1) when f rises above \bar{f} .

Put differently, there is a level $\bar{\lambda} > 0$ for the weight on output-gap stabilization for which

$$f^*(\alpha_1, \gamma_2; \bar{\lambda}) = \bar{f}. \quad (3.4)$$

Then we have

$$\hat{f}(\lambda) = \begin{cases} \bar{f} \leq f^*(\alpha_1, \gamma_2; \lambda) & \text{if } \lambda \leq \bar{\lambda}, \\ f^*(\alpha_1, \gamma_2; \lambda) < \bar{f} & \text{if } \lambda > \bar{\lambda}. \end{cases} \quad (3.5)$$

Since $f^*(\alpha_1, \gamma_2; \lambda) \geq 0$ and $\bar{f} > 0$, this also confirms that the restriction $\hat{f} \geq 0$ imposed above is not binding.

In the analysis above, the model is assumed to be constant over the time. The case when the model is varying over time is examined in appendix B. The result is that, under strict inflation targeting, the worst possible model is not unique but includes both the values (α_1, γ_2) and (α_2, γ_1) . Under flexible inflation targeting, if λ is so large as to fulfill $\lambda > \bar{\lambda}$, the worst possible model is (α_1, γ_2) and constant over time. Thus, surprisingly little of the analysis changes if the model is allowed to vary over time.

4 Discussion

Robust control consists of selecting the policy that results in the best outcome for the worst possible model. In this example, the worst possible model ends up being on the boundary of the assumed

feasible set of models. Thus, the assumptions about the feasible set of models are crucial for the outcome. Robust control has been proposed as a non-Bayesian alternative that utilizes less prior assumptions about the model. If the worst possible model somehow ended up in the interior of the feasible set of models, one could perhaps argue that the outcome is less sensitive to the assumptions about the feasible set of models. When the worst possible model is on the boundary of the feasible set of models, this is no longer so. Prior assumptions about the feasible set of models become crucial to the outcome.

Furthermore, if a Bayesian prior probability measure were to be assigned to the feasible set of models, one might find that the probability assigned to the models on the boundary are exceedingly small. Thus, highly unlikely models can come to dominate the outcome of robust control.

For weights on output-gap stabilization below a critical value ($\lambda \leq \bar{\lambda}$), optimal policy is given by \bar{f} in (3.2), which equals the “middle” γ , $\frac{\gamma_1 + \gamma_2}{2}$, divided by “middle” α , $\frac{\alpha_1 + \alpha_2}{2}$. Thus, policy is “intermediate” in this sense, neither particularly “aggressive” nor particularly “cautious” (although still depending on the bounds of the supports). For weights above that critical value ($\lambda > \bar{\lambda}$), policy is given by the optimal policy when the economy has a low α and a high γ . The worst possible model is one where the inherent serial correlation in inflation is the largest and the effect of the output gap on inflation is the smallest. Since the response coefficient is decreasing in α and increasing in γ , this means that the response coefficient is the largest for the different models (when the model is known). In this sense, the policy under robust control ends up being aggressive.

Finally, under robust control, the commitment to the reaction function is made initially and forever, given only an assumption about the feasible set of models. Clearly, it would be better to observe the economy for some time and learn more about the model before making a commitment to a given policy, especially if the model is constant over time. This leads to optimal learning, as in Wieland [16]–[18], for instance.

A Properties of the optimal policy with known model

Differentiating the first-order condition (2.8), it directly follows that the optimal policy $f^*(\alpha, \gamma; \lambda)$ fulfills

$$\begin{aligned} f_\alpha^*(\alpha, \gamma; \lambda) &\equiv -\frac{V_{f\alpha}(f^*; \alpha, \gamma; \lambda)}{V_{ff}(f^*; \alpha, \gamma; \lambda)} = -2\frac{2\lambda f^{*2}(\gamma - \alpha f^*) + \alpha^2(1 + \lambda f^{*2})f^*}{[1 - (\gamma - \alpha f^*)^2]^2 V_{ff}(f^*; \alpha, \gamma; \lambda)} < 0, \\ f_\gamma^*(\alpha, \gamma; \lambda) &\equiv -\frac{V_{f\gamma}(f^*; \alpha, \gamma; \lambda)}{V_{ff}(f^*; \alpha, \gamma; \lambda)} = -2\frac{-2\lambda f^*(\gamma - \alpha f^*) - \alpha(1 + \lambda f^{*2})}{[1 - (\gamma - \alpha f^*)^2]^2 V_{ff}(f^*; \alpha, \gamma; \lambda)} > 0, \\ f_\lambda^*(\alpha, \gamma; \lambda) &\equiv -\frac{V_{f\lambda}(f^*; \alpha, \gamma; \lambda)}{V_{ff}(f^*; \alpha, \gamma; \lambda)} = -2\frac{f^*[1 - (\gamma - \alpha f^*)^2] - \alpha f^{*2}(\gamma - \alpha f^*)}{[1 - (\gamma - \alpha f^*)^2]^2 V_{ff}(f^*; \alpha, \gamma; \lambda)} \\ &= -2\frac{\alpha(\gamma - \alpha f^*)/\lambda}{[1 - (\gamma - \alpha f^*)^2]^2 V_{ff}(f^*; \alpha, \gamma; \lambda)} < 0, \end{aligned}$$

where the last equality uses the first-order condition, (2.8).

B Time-varying model

Suppose the model may vary from period to period. Let the Phillips curve be

$$\pi_{t+1} = \gamma_{t+1}\pi_t + \alpha_{t+1}x_t + \varepsilon_{t+1}, \quad (\text{B.1})$$

instead of (2.1), where $m_{t+1} \equiv (\alpha_{t+1}, \gamma_{t+1}) \in M$ is the model in period $t+1$. Let $f_t \in F \equiv \{f \mid f \geq 0\}$ be the policy in period t , where

$$x_t = -f_t\pi_t. \quad (\text{B.2})$$

Let the worst possible model in period $t+1$ be the model $m_{t+1} \in M$ that results in the largest value of

$$\mathbf{E}_t L_{t+1} = \mathbf{E}_t \pi_{t+1}^2 + \lambda \mathbf{E}_t x_{t+1}^2. \quad (\text{B.3})$$

Under the policy $f_t \in F$, we have

$$\pi_{t+1} = (\gamma_{t+1} - \alpha_{t+1}f_t)\pi_t + \varepsilon_{t+1}.$$

It follows that, for given m_{t+1} and f_{t+1} ,

$$\mathbf{E}_t L_{t+1} = [(\gamma_{t+1} - \alpha_{t+1}f_t)^2 + \lambda f_{t+1}^2]\pi_t^2 + \sigma_\varepsilon^2.$$

Consider f_{t+1} independent of m_{t+1} . Then, it follows that the worst possible model is given by

$$m_{t+1} = \arg \max_{m \in M} |\gamma - \alpha f_t|.$$

That is,

$$m_{t+1} = \hat{m}(f_t),$$

where $\hat{m}(f)$ is given by (3.2) and (3.3).

Under strict inflation targeting, $\lambda = 0$, it follows that the robust policy is constant over time and given by

$$\hat{f}_t = \bar{f},$$

where \bar{f} is given by (3.2). Then, $m_{t+1} = \hat{m}(\bar{f})$ is not unique but can vary arbitrarily over time between the values (α_1, γ_2) and (α_2, γ_1) .

Under flexible, inflation targeting, $\lambda > 0$, it follows that the policy is also constant and given by

$$\hat{f}_t = \hat{f}(\lambda),$$

where $\hat{f}(\lambda)$ is given by (3.4) and (3.5). For $\lambda \leq \bar{\lambda}$, $\hat{f}_t = \bar{f}$, and it follows that the worst possible model $m_{t+1} = \hat{m}(\bar{f})$ can vary over time between the values (α_1, γ_2) or (α_2, γ_1) . For $\lambda > \bar{\lambda}$, $\hat{f}_t = f^*(\alpha_1, \gamma_2; \lambda) < \bar{f}$ and the worst possible model is given by $m_{t+1} = (\alpha_1, \gamma_2)$ and is constant over time.

References

- [1] Ball, Laurence (1999), “Policy Rules for Open Economies,” in Taylor [15].
- [2] Brainard, William (1967), “Uncertainty and the Effectiveness of Policy,” *American Economic Review* 57, *Papers and Proceedings*, 411–425.
- [3] Giannoni, Marc P. (1999), “Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model,” Working Paper.
- [4] Hansen, Lars Peter, and Thomas J. Sargent (1998), “Alternative Representation of Discounted Robust Linear Quadratic Control,” Working Paper.
- [5] Hansen, Lars Peter, and Thomas J. Sargent (2000), “Robust Control and Filtering of Forward-Looking Variables,” Working Paper.
- [6] Onatski, Alexei (1998), “On the Robustness of Policy Rules to Different Kinds of Model Uncertainties,” Working Paper.

- [7] Onatski, Alexei, and James H. Stock (2000), “Robust Monetary Policy Under Model Uncertainty in a Small Model of the U.S. Economy,” NBER Working Paper No. 7490.
- [8] Rudebusch, Glenn D. (1999), “Is the Fed Too Timid? Monetary Policy in an Uncertain World,” Working Paper, FRBSF.
- [9] Rudebusch, Glenn D., and Lars E.O. Svensson (1999), “Policy Rules for Inflation Targeting,” in Taylor [15].
- [10] Sargent, Thomas J. (1999), “Comment on Ball (1999),” in Taylor [15].
- [11] Söderström, Ulf (2000), “Monetary Policy with Uncertain Parameters,” ECB Working Paper No. 13.
- [12] Stock, James H. (1999), “Comment on Rudebusch and Svensson (1999),” in Taylor [15].
- [13] Svensson, Lars E.O. (1997), “Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets,” *European Economic Review* 41, 1111–1146.
- [14] Svensson, Lars E.O. (1999), “Inflation Targeting: Some Extensions,” *Scandinavian Journal of Economics* 101, 337–361.
- [15] Taylor, John B., ed. (1999), *Monetary Policy Rules*, Chicago University Press.
- [16] Wieland, Volker (1996) “Monetary Policy, Parameter Uncertainty and Optimal Learning,” Working Paper, Federal Reserve Board. *Journal of Monetary Economics*, forthcoming.
- [17] Wieland, Volker (1998), “Monetary Policy and Uncertainty about the Natural Unemployment Rate,” FEDS Working Paper No. 22.
- [18] Wieland, Volker (2000), “Learning by Doing and the Value of Optimal Experimentation,” *Journal of Economic Dynamics and Control* 24(4).