Robust Control

An Entry for the New Palgrave, 2nd Edition

Noah Williams

Department of Economics, Princeton University

E-mail: noahw@princeton.edu

1. OVERVIEW

Robust control considers the design of decision or control rules that fare well across a range of alternative models. Thus robust control is inherently about model uncertainty, particularly focusing on the implications of model uncertainty for decisions. Robust control originated in the 1980s in the control theory branch of the engineering and applied mathematics literature, and it is now perhaps the dominant approach in control theory. Robust control gained a foothold in economics in the late 1990s and has seen increasing numbers of economic applications in the past few years.[1]

The basic issues in robust control arise from adding more detail to the opening sentence above – that a decision rule performs well across alternative models. To begin, define a model as a specification of a probability distribution over outcomes of interest to the decision maker, which is influenced by a decision or control variable. Then model uncertainty simply means that the decision maker faces subjective uncertainty about the specification of this probability distribution. A first key issue in robust control then is to specify the class of alternative models which the decision maker entertains. As we discuss below, there are many approaches to doing so, with the most common cases taking a benchmark *nominal model* as a starting point and considering perturbations of this model. How to specify and measure the magnitude of the perturbations are key practical considerations.

With the model set specified, the next issue is how to choose a decision rule and thus what it means for a rule to "perform well" across models. In Bayesian analysis, the deci-

[1] For related surveys see Hansen and Sargent (2001) and Backus, Routledge, and Zin (2005). For a more comprehensive view of the leading approach to robust control in economics, see Hansen and Sargent (2007).

sion maker forms a prior over models and proceeds as usual to maximize expected utility (or minimize expected loss). Just as we defined a model as a probability distribution, a Bayesian views model uncertainty as simply a hierarchical probability distribution with one layer consisting of shocks and variables to be integrated over, and another layer averaging over models. In contrast, most robust control applications focus on minimizing the worst case loss over the set of possible models (a minimax problem in terms of losses, or max-min expected utility). Stochastic robust control problems thus distinguish sharply between shocks which are averaged over, and models which are not. The robust control approach thus presumes that decision makers are either unable or unwilling to form a prior over the forms of model misspecification. Of course decision makers must be able to specify the set of models as discussed above, but typically this involves bounding the set of possibilities in some way rather than fully specifying each alternative. Finally, there are some approaches which seek a middle ground between the average case and the worst case, for example by maximizing expected utility subject to a bound on the worst case loss. These have been less prominent both in control theory (Limebeer, Anderson, and Hendel (1994) is one example) and in economics (Tornell (2003) is one exception), and thus will not be discussed further. For the remainder of the article robust control will mean a minimax approach.

2. MORE DETAIL

2.1. Robustness and Worst Case Analysis

Broadly speaking, the control theory literature has adopted the worst-case philosophy out of concerns for stability. A basic desiderata for robust control in practice is that the system remain stable in the face of perturbations, and since instability may be equated with infinite loss, minimizing the worst case outcomes will insure stability (when possible). Moreover many engineering applications have specific performance objectives which must be maintained, and a cost function penalizing deviations is not clearly specified. However in dealing with economic agents rather than controlled machines, decision theoretic criteria naturally come into play. In this sphere, robust control is closely related to the notions of Knightian uncertainty, ambiguity, and uncertainty aversion, which are all roughly equivalent (although sometimes differing in formalization).

Starting with the observations of the classic Ellsberg (1961) paradox – that (some) decisionmakers prefer environments with known odds to those with uncertain probabilities, there has been a broad literature in decision theory which has weakened the Savage axioms to incorporate preferences which display such aversion to uncertainty or ambiguity. The most widely used characterization is due to Gilboa and Schmeidler (1989), who axiomatized ambiguity preferences with multiple priors. Decisionmaking with multiple priors can be represented as max-min expected utility: maximizing the utility with respect to the least favorable prior from a convex set of priors. More recently, Epstein and Schneider (2003) have extended the static environment of Gilboa and Schmeidler to a dynamic context, where the set of priors is updated over time. Hansen, Sargent, Turmuhambetova, and Williams (2006) formally established the links between robust control and ambiguity aversion, showing that the model set of robust control as discussed above can be thought of as a particular specification of Gilboa and Schmeidler's set of priors. Moreover, although the ambiguity preferences are characterized by posing particular counterfactuals which require multiple priors, once the least favorable prior is chosen, behavior could be rationalized as Bayesian with that prior. Thus from a Bayesian viewpoint Sims (2001) views robust control as a means of generating priors, which then naturally leads to questioning whether the worst case prior accurately reflects actual beliefs and preferences.[2] Finally, in many cases robust or ambiguity averse preferences are similar to enhanced risk aversion, and in some cases they are observationally equivalent. This insight dates to Jacobson (1973) and Whittle (1981) in the control theory literature, and the relations between robust control and a particular specification of Kreps and Porteus (1978)/Epstein and Zin (1989)/Duffie and Epstein (1992) recursive utility with enhanced risk aversion have been shown by Anderson, Hansen, and Sargent (2003), Hansen, Sargent, Turmuhambetova, and Williams (2006), and Skiadas (2003).

[2]See also Svensson (2001). Hansen, Sargent, Turmuhambetova, and Williams (2006) show how to back out the Bayesian prior which rationalizes robust decisionmaking.

2.2. Control Theory Background

Since many of the ideas and inspiration for robust control in economics come from control theory, we give here just a broad outline of its development. More detail and different perspectives can be found in the books by Zhou, Doyle, and Glover (1996), Başar and Bernhard (1995), and Burl (1999). Throughout the late 1960s and early 1970s optimal control came into its own, largely through the work of Kalman on linear quadratic (LQ) control and filtering. While this approach remains widely used today throughout economics, starting in the late 1970s and early 1980s the control theory literature started to change as theory and practice showed some of the limitations of the LQ approach. Although LQ control with full observation (the so-called linear quadratic regulator or LQR) was known to be robust to some types of model perturbations, Doyle (1978) showed that there are no such assurances in the case of partial observation (the so-called linear-quadratic-Gaussian or LQG case, which is an LQR control matched with a Kalman filter). Doyle's paper title and abstract are classic in the literature – title: "Guaranteed Margins for LQG Regulators", abstract: "There are none."

Spurred by this and related work, control theorists started to move away from LQ control to look for a more robust approach. Zames (1981) was influential in the development of H_∞ control as a more robust alternative to LQ control. Loosely speaking, in LQ control the quadratic cost means that performance is measured with a 2-norm across frequencies. By contrast, H_∞ uses an $\infty-$norm that looks at the peak of the losses across frequencies. It is also interpretable as the maximal magnification of the disturbances to outputs of interest. While the early robust control literature used a frequency domain approach, in the late 1980s Doyle and others developed state space formulations (see Doyle, Glover, Khargonekar, and Francis (1989) for example) which gave explicit solutions and allowed for alterative formalizations. For example, the H_∞ approach was given alternative justifications in terms of penalizing disturbances from the nominal model, which can be implemented as a dynamic game between a decision maker seeking to minimize losses and a malevolent agent seeking to maximize loss. (See Başar and Bernhard (1995) for a development of this approach.) Finally, the uncertainty sets in the H_∞ approach are unstructured – they represent perturbations of the model which are bounded but have no particular form. The implications of structured perturbations have been studied more recently. Some examples include parametric perturbations, unmodelled dynamics, or uncertainty only about particular channels or connections in a model. Applications with structured uncertainty use the structured singular value (also known an μ) rather than the H_∞ norm as a measure of performance. Although there are some important stability and performance criteria, in general constructing control rules is a more daunting task and the theory is not as fully developed as the unstructured case.

2.3. The Hansen-Sargent Approach

In the economics literature, the most prominent and influential approach to robust control is due to Hansen and Sargent (and their co-authors), which is summarized in their monograph Hansen and Sargent (2007). This approach starts with a nominal model and uses entropy as a distance measure to calibrate the model uncertainty set. More specifically, the model set consists of those models whose relative entropy or Kullback-Leibler distance from the nominal model is bounded by a specified value. Note that this puts no structure on the uncertainty, but only restricts the alternative models to those which are difficult to distinguish statistically from the nominal model. In practice, a Lagrange multiplier theorem is typically used to convert the entropy constraint into a penalty on perturbations from the model. Then the solution of the control problem is found via a dynamic game implementation: the agent maximizes utility by his choice of control, while an evil agent minimizes utility by his choice of perturbation, while being penalized by the entropy of the deviations. Relative to the control theory literature such as Başar and Bernhard (1995), the main differences are that all models are stochastic, while control theory largely uses deterministic models. One exception is Petersen, James, and Dupuis (2000) who use a similar approach to consider uncertain stochastic systems. In addition, discounting is not typically considered in control theory, while it is natural in economics. In full information problems discounting has relatively little effect, but it raises important issues in problems with partial information (see Hansen and Sargent (2005a) and Hansen and Sargent (2005b)). Finally, the Hansen-Sargent approach naturally extends beyond the LQ setting laid out in Hansen, Sargent, and Tallarini (1999), with some examples in Anderson, Hansen, and Sargent (2003), Cagetti, Hansen, Sargent, and Williams (2002), and Maenhout (2004).

To be more concrete, consider an LQ example where x_t is the state, i_t is the agent's control, and ε_t is an i.i.d. Gaussian shock. The nominal model is:

$$x_{t+1} = Ax_t + Bi_t + C\varepsilon_{t+1}, \tag{1}$$

and the agent's intertemporal preferences are:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(x_t' Q x_t + i_t' R i_t \right) \tag{2}$$

where $0 < \beta < 1$ and Q and R are negative definite matrices. The approach of Hansen and Sargent perturbs the nominal model with an additional "misspecfication shock" w_{t+1} which is allowed to be correlated with the state x_t:

$$x_{t+1} = Ax_t + Bi_t + C(\varepsilon_{t+1} + w_{t+1}). \tag{3}$$

The shock w_{t+1} is used to represent alternative models. These models are made to be close to the nominal model in an entropy sense by imposing the bound:

$$E_0 \sum_{t=0}^{\infty} \beta^t w_{t+1}' w_{t+1} \leq \eta \tag{4}$$

for some constant $\eta \geq 0$. The agent then maximizes (2) with respect to the worst case perturbed model (3) from the set (4). Using a Lagrange multiplier theorem, the constraint set can be converted to a penalty and the decision problem can by solved recursively by solving the Bellman equation for a two player zero sum game:

$$V(x) = \max_i \min_w \left\{ x'Qx + i'Ri + \beta\theta w'w + \beta E\left[V(Ax + Bi + C(\varepsilon + w)) | x \right] \right\} \tag{5}$$

where $\theta > 0$ is a Lagrange multiplier on the constraint (4) and the expectation is over the Gaussian shock ε. Often this multiplier formulation is taken as the starting point, for example Maccheroni, Marinacci, and Rustichini (2006) characterize preferences of this form, with θ governing the degree of robustness. As $\theta \to \infty$ the penalization becomes so great that only the nominal model remains (thus $\eta \to 0$), and the decision rule is less robust. Conversely, there is typically a minimal value of θ beyond which the value is $V(x) = -\infty$. This gives the most robust decision rules, allowing for the largest uncertainty set.

2.4. Adding Structure to the Uncertainty Set

The approach discussed above uses unstructured uncertainty, and has been well developed and extended in different dimensions. We now discuss some alterative approaches which put more structure on the uncertainty set. There are many reasons to do so. It may be that some of the models that are close to the nominal model in a statistical sense may not be plausible economically. Alternatively, the decision makers may have a discrete set of models in mind, and bounding them all in one uncertainty set may include extraneous implausible models. Perhaps most substantively, the decision maker may be more confident some aspects of the model relative to others. Some examples of this include knowing the model up to the values of parameters, or being more certain about the dynamics of certain variables in the model. Not taking into account the particular structure may give a misleading impression of the actual uncertainty the decision makers face.

There are many ways of building in structured uncertainty, and the distinctions between cases are not always clear. For example, consider the same nominal model (1) as above, but suppose that instead of the unstructured perturbations (3) the uncertainty is instead solely in the values of the parameters A and B. Thus we can represent the parametric perturbed models as:

$$x_{t+1} = (A + \hat{A})x_t + (B + \hat{B})i_t + C\varepsilon_{t+1} \tag{6}$$

for some matrices $\hat{A}$ and $\hat{B}$. Of course it's possible to re-write (6) as a version of (3) with:

$$w_{t+1} = \hat{A}x_t + \hat{B}i_t, \tag{7}$$

so in principle parametric perturbations are just a special case of the unstructured uncertainty. However what makes a substantive difference is how uncertainty is measured, that is whether we restrict w_{t+1} as in (4) or whether we restrict the parameters $\hat{A}$ and $\hat{B}$, say by bounding them in a confidence ellipsoid around the nominal model. Moreover, as (7) makes clear the differences between the uncertainty measurements will depend on the actual control rule i_t in place. Onatski and Williams (2003) provide an example of a simple estimated model where the uncertainty specifications matter dramatically for outcomes. In particular, the optimal policy for the largest possible unstructured uncertainty set (i.e. for the minimal value of θ) leads to instability for relatively small parametric perturbations. Thus the particular structure and measurement of uncertainty can have important implications for decisions.[3]

Some economic applications with structured uncertainty include the following:

- The simplest cases are uncertainty sets with discrete possible models. Some examples include: Levin and J. Williams (2003) who consider both Bayesian and minimax approaches, Cogley and Sargent (2005) and Svensson and Williams (2006) who focus on a Bayesian approach, and the recent work of Hansen and Sargent (2006) who have built this type of structure into their robust approach.

- Another common form is parameter uncertainty within a fully specified model. Brainard (1967) is the classic reference from a Bayesian perspective with many references in this line, while Giannoni (2002) and Chamberlain (2000) consider minimax approaches.

- Somewhat more broad are cases with different parametric model specifications. For example this includes uncertainty about dynamics (lags and leads), variables which may enter, uncertainty about data quality, and other features which are built into parametric extensions of the nominal model. Examples include the model error modeling approach of Onatski and Williams (2003) and the empirical specifications of Brock, Durlauf, and West (2003).

- Finally, the model sets may be nonparametric but structured in particular ways. For example, Onatski and Stock (2002) consider different structured types of uncertainty such as linear time-invariant perturbations, nonlinear time-varying perturbations, and perturbations which only enter particular parts of the model. Other examples include

[3]Petersen, James, and Dupuis (2000) modify the unstructured approach described above to deal with structured uncertainty by to separating the entropy penalty for unstructured perturbations from a different penalization for structured perturbations.

nonparametric specifications of uncertainty which differs across frequencies as in Onatski and Williams (2003) and Brock and Durlauf (2005).

References

Anderson, E., L. P. Hansen, and T. Sargent (2003). A quartet of semi-groups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association 1*, 68–123.

Backus, D. K., B. R. Routledge, and S. E. Zin (2005). Exotic preferences for macroeconomists. In M. Gertler and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2004*, pp. 319–390. Cambridge: MIT Press.

Başar, T. and P. Bernhard (1995). *H_∞-Optimal Control and Related Minimax Design Problems*. Boston: Birkhauser.

Brainard, W. (1967). Uncertainty and the effectiveness of policy. *American Economic Review 57*, 411–425.

Brock, W. A. and S. N. Durlauf (2005). Local robustness analysis: Theory and application. *Journal of Economic Dynamics and Control 29*, 2067–2092.

Brock, W. A., S. N. Durlauf, and K. D. West (2003). Policy evaluation in uncertain economic environments. *Brookings Papers on Economic Activity*, 235–301.

Burl, J. B. (1999). *Linear Optimal Control: $\mathcal{H}_2$ and $\mathcal{H}_\infty$ Methods*. Menlo Park, California: Addison-Wesley.

Cagetti, M., L. P. Hansen, T. J. Sargent, and N. Williams (2002). Roubstness and pricing with uncertain growth. *Review of Financial Studies 15*, 363–404.

Chamberlain, G. (2000). Econometric applications of maxmin expected utility theory. *Journal of Applied Econometrics 15*, 625–644.

Cogley, T. and T. J. Sargent (2005). The conquest of U.S. inflation: Learning and robustness to model uncertainty. *Review of Economic Dynamics 8*, 528–563.

Doyle, J. C. (1978). Guaranteed margins for LQG regulators. *IEEE Transactions on Autmatic Control 23*, 756–757.

Doyle, J. C., K. Glover, P. P. Khargonekar, and B. A. Francis (1989). State-space solutions to standard $\mathcal{H}_2$ and $\mathcal{H}_\infty$ control problems. *IEEE Transactions on Autmatic Control 34*, 831–847.

Duffie, D. and L. G. Epstein (1992). Stochastic differential utility. *Econometrica 60*(2), 353–394.

Ellsberg, D. (1961). Risk, ambiguity and the Savage axioms. *Quarterly Journal of Economics 75*, 643669.

Epstein, L. and M. Schneider (2003). Recursive multiple-priors. *Journal of Economic-Theory 113*, 1–31.

Epstein, L. and S. Zin (1989). Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*, 937–969.

Giannoni, M. (2002). Does model uncertainty justify caution? Robust optimal monetary policy in a forward-looking model. *Macroeconomic Dynamics 6*, 111–144.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*, 141–153.

Hansen, L. P., T. Sargent, and T. Tallarini (1999). Robust permanent income and pricing. *Review of Economic Studies 66*, 873–907.

Hansen, L. P. and T. J. Sargent (2001). Acknowledging misspecification in macroeconomic theory. *Review of Economic Dynamics 4*, 519–535.

Hansen, L. P. and T. J. Sargent (2005a). Robust estimation and control under commitment. *Journal of Economic Theory 124*, 258–301.

Hansen, L. P. and T. J. Sargent (2005b). Robust estimation and control without commitment. unpublished.

Hansen, L. P. and T. J. Sargent (2006). Fragile beliefs and the price of model uncertainty. unpublished.

Hansen, L. P. and T. J. Sargent (2007). *Robustness*. Princeton, New Jersey (forthcoming): Princeton University Press.

Hansen, L. P., T. J. Sargent, G. A. Turmuhambetova, and N. Williams (2006). Robust control and model misspecification. *Journal of Economic Theory 128*, 45–90.

Jacobson, D. H. (1973). Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Transactions for Automatic Control AC-18*, 1124–131.

Kreps, D. M. and E. L. Porteus (1978). Temporal resolution of uncertainty and dynamic choice. *Econometrica 46*, 185–200.

Levin, A. T. and J. Williams (2003). Robust monetary policy with competing reference models. *Journal of Monetary Economics 50*, 945–975.

Limebeer, D. J. N., B. D. O. Anderson, and B. Hendel (1994). A Nash game approach to mixed H_2/H_∞ control. *IEEE Transactions on Automatic Control 39*, 69–82.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Dyanmic variational preferences. *Journal of Economic Theory 128*, 4–44.

Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *Review of Financial Studies 17*, 951–983.

Onatski, A. and J. H. Stock (2002). Robust monetary policy under model uncertainty in a small model of the US economy. *Macroeconomic Dynamics 6*, 85–110.

Onatski, A. and N. Williams (2003). Modeling Model Uncertainty. *Journal of the European Economic Association 1*, 1087–1122.

Petersen, I. R., M. R. James, and P. Dupuis (2000). Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Transactions on Automatic Control 45*, 398–412.

Sims, C. A. (2001). Pitfalls of a minimax approach to model uncertainty. *American Economic Review 91*, 51–54.

Skiadas, C. (2003). Robust control and recursive utility. *Finance and Stochastics 7*, 475–489.

Svensson, L. E. (2001). Robust control made simple. Working paper, Princeton University.

Svensson, L. E. O. and N. Williams (2006). Monetary policy with model uncertainty: Distribution foreacast targeting. Working paper, Princeton University.

Tornell, A. (2003). Exchange rate anomalies under model misspecification: A mixed optimal/robust approach. Working paper, UCLA.

Whittle, P. (1981). Risk sensitive linear quadratic Gaussian control. *Advances in Applied Probability 13*, 764–777.

Zames, G. (1981). Feedback and optimal sensitivity: model reference transformations, multiplicative seminorms, and approximate inverses. *IEEE Transactions on Automatic Control 26*, 301–320.

Zhou, K., J. Doyle, and K. Glover (1996). *Robust and Optimal Control*. Upper Saddle River, New Jersey: Prentice Hall.