Comment on

'Certainty Equivalence and Model Uncertainty'
by Lars Hansen and Thomas Sargent

at the FRB conference
'Models and Monetary Policy: Research in the Tradition of
Dale Henderson, Richard Porter and Peter Tinsley'

By Volker W. Wieland*

Goethe University of Frankfurt and Center for Financial Studies

1. The Three Musketeers on Control, Model Uncertainty and Monetary Policy

As a former Federal Reserve Board economist it is a great pleasure for me to contribute to this 'Festschrift' honoring the achievements of FRB researchers Dale Henderson, Richard Porter and Peter Tinsley by commenting on research in their tradition, in this case a contribution by Lars Hansen and Thomas Sargent, two of the leading macroeconomists of our time. Soon after I joined the Fed, Dick Porter, at the time the boss of my boss, approached me not to assign work but to discuss potential applications of my thesis research on learning and control to monetary policy. One of his first remarks was 'You should also talk to Peter Tinsley about this' and then introduced me to Peter. Soon afterwards I also discussed my findings with Dale Henderson.

I quickly realized that these three experienced researchers were unique in their ability and willingness to provide advice and guidance to a young economist embarking on his first research papers. Many helpful and sometimes intense discussions followed in the five years I spent at the Fed. Thus, even though Dale, Dick and Peter worked in different divisions, did not co-author papers and did not seem especially close to each other, the picture that immediately came to my mind when preparing this discussion was the one of the three musketeers. Just like the three musketeers of Alexandre Dumas—Athos, Portos and Aramis—were renowned for their superior sword-fighting ability in the service of his majesty the King of France, Dick, Dale and Peter stood out at the Fed as masters of the weapons of the researcher—intellect, curiosity and imagination—and put them to good use in the service of science and, of course, the Chairman and the Governors. So I am not surprised that their achievements are honored jointly in this Festschrift and I was pleased to see Jeff Fuhrer pick up on Dumas at the conference and suggest a D'Artagnan for our three musketeers in the late George Moore.

The paper by Hansen and Sargent on 'Certainty Equivalence and Model Uncertainty' is well placed in this Festschrift. Hansen and Sargent elegantly show that the separation principle, which applies to stochastic control problems with quadratic objectives and linear transition laws, can also be extended to robust control problems of a linear-quadratic nature. In stochastic control this principle, which is often referred to as 'certainty-equivalence', allows the separation of control and estimation (or filtering). In other words, the stochastic control problem with unknown realizations (from known probability functions) can be solved by optimizing first under perfect foresight and then replacing unknown future values with optimal forecasts. Robust control at first glance seems to elude such simplicity. A robust controller is concerned with model misspecfication and considers alternative models without assigning probabilities like a Bayesian decision maker. Robustness is achieved by acting as if nature would choose the worst of the available models and guarding against that possibility. Hansen and Sargent show that a version of the separation principle prevails in a linear-quadratic version of this game. The first step remains to optimize under the assumption that future values are perfectly foreseen. In the estimation step, however, the robust controller distorts beliefs relative to the approximating model to achieve robustness. The distorted law of motion needed to reveal this version of certainty-equivalence arises when nature acts as Stackelberg leader.

The issues Hansen and Sargent touch upon in this paper resonate with the work of Tinsley, Henderson and Porter. Linear-quadratic stochastic control has been used to model economic phenomena for almost fifty years. It has been widely applied to study the behavior of forward-looking economic agents and continues to play an important role in academic research as well as practical monetary policy analysis. For example, central bank staff now routinely compute optimal policy rules for linear macroeconometric models with rational expectations under quadratic policy objectives and make use of such rules as benchmarks for policy prescriptions. This progress owes much to the three FRB musketeers. For example, Tinsley and his co-authors were frontrunners in applying stochastic control and filtering techniques to the analysis of monetary policy (Kalchbrenner, Tinsley and others 1975, 1976, 1977). They introduced feedback control to the design of monetary policy and carefully analysed how all available information may be filtered for use in policy design. Aware of the challenges of practical policymaking, our three honorees were concerned early on with issues of risk, parameter and model uncertainty. Dale Henderson, for example, already investigated optimal macroeconomic policy under conditions of risk in his very first published paper, (Henderson and Turnovsky 1972). At the Fed Tinsley and von zur Muehlen (1981) proposed a maximum probability approach to short-run policy to address key uncertainties. It is not surprising that in this stimulating environment the potential benefits of robust control under model uncertainty were already introduced to monetary policy by von zur Muehlen (1982) much before its recent gain in prominence in macroeconomics. The concern for robustness of policy strategies in the face of likely model specification was also behind the focus on simple rules for policy. In this context it is worth pointing out that Henderson extensively investigated the value of simple interest rate rules, now typically referred to as Taylor rules, in a paper published in the same conference volume as Taylor's well-known contribution, (Henderson and McKibbin (1993)). The most-cited alternative to simple interest rate rules are simple money growth rules. To date, proponents of this approach still tend to rely on the well-known 'P-star' model developed by Hallman, Porter and Small (1991).

2. A Separation Principle for Robust Control

This paper is part of a larger research effort by Hansen and Sargent to introduce robust behavior modelled in terms of robust control as a behavioral assumption in economics. This research agenda has the potential to explain economic behavior by positing a preference for robustness when the standard expected utility maximization paradigm fails to produce explanations. The upcoming monograph 'Misspecification in Recursive Macroeconomic Theory' by Hansen and Sargent promises to become an influential piece of work providing a highly useful toolbox for further research in this area. In the past, both authors have also contributed importantly to macroeconomics by popularizing the use of linear-quadratic stochastic control techniques. For those techniques, the certainty-equivalence principle implies a very useful simplification of dynamic stochastic analysis. Optimal rules can be derived in two steps where step one is to optimize for a given set of expecations (perfect foresight) and step two is to form expectations optimally. Thus, linear-quadratic stochastic control constitutes an easy-to-use approach for analyzing optimizing economic behavior in the presence of uncertainty. Divergences between empirical implications of linear-quadratic model economies and real-world economic data identify important starting points for further research.

The present paper focusses on the extension of the certainty-equivalence principle from stochastic to robust control of linear dynamic economies with quadratic objective. To review the contribution and possible limitations of this paper I follow the notation of Hansen and Sargent (HS) and use y_t to denote the vector of state variables with an exogenous component z_t, an endogenous component x_t and the following transition laws:

$$z_{t+1} = f(z_t, \epsilon_{t+1}) \tag{1}$$

$$x_{t+1} = g(x_t, z_t, u_t) \tag{2}$$

ϵ_{t+1} refers to an i.i.d. sequence of random vectors with mean-zero normal distribution. f and g are known linear functions. The decision maker chooses the control variable u_t to maximize the expected value of the quadratic objective function r:

$$E[\sum_{t=0}^{\infty} \beta^t r(y_t, u_t)|\mathbf{y}^0] \tag{3}$$

where $\beta \in (0, 1)$. It is well known that the solution under certainty about future variables is a linear rule. Certainty equivalence implies that the deterministic solution is identical to the solution to the stochastic control problem with future variables replaced by their rational forecasts. The optimal rule then feeds back on initial information on x_t and z_t, $u_t = h(x_t, z_t)$.

Hansen and Sargent's robust controller fears that (1) is misspecified and measures the difference from the true law of motion for the uncertain state z by the distortions w,

$$z_{t+1} = f(z_t, \epsilon_{t+1} + w_{t+1}). \tag{4}$$

An upper bound on this difference is defined by η_0 according to:

$$\hat{E}[\sum_{t=0}^{\infty} \beta^t w'_{t+1} w_{t+1} \Big| y_0 \leq \eta_0] \tag{5}$$

where $\hat{E}$ concerns the distribution generated by the distorted law of motion (4). A robust decision rule is derived from the Markov perfect equilibrium of the following two-player zero-sum game:

$$\min_{\mathbf{w}_1} \max_{\mathbf{u}_0} \hat{E}[\sum_{t=0}^{\infty} \beta^t \left\{ r(y_t, u_t) + \theta \beta w'_{t+1} w_{t+1} \right\}] \Big| y_0 \tag{6}$$

HS call this a *multiplier problem.* In the manuscript of their upcoming monograph they discuss the equivalence of the multiplier problem to the min-max problem $\min_{\mathbf{w}_1} \max_{\mathbf{u}_0} \hat{E}[\sum_{t=0}^{\infty} \beta^t r(y_t, u_t)] \Big| y_0$ subject to the constraint (5). They call the latter specification a *constraint problem.*

In the Markov perfect equilibrium the players choose sequentially and simultaneously each period, taking the other player's decision rule as given. The distortions w feed back on the endogenous state x and require the maximizing player to design a rule that accounts for the possibility of mis-specified dynamics. The Markov perfect equilibrium is defined by the decision rules of the two players, $u_t = H(x_t, z_t)$ and $w_{t+1} = W(x_t, z_t)$, with $H(.)$ representing the robust rule. This rule promises a lower maximal rate at which the objective can deteriorate with increases in misspecification as measured by the term $w'_{t+1} w_{t+1}$.

From the Markov perfect equilibrium it is not apparent that a version of the separation principle applies to this game. An important contribution of the HS paper in this volume is to show that the separation principle applies to the Stackelberg version of this game. Due to the zero-sum nature of this game the solutions under the Stackelberg and Markov timing protocols are identical. In the Stackelberg case, nature acts as Stackelberg leader and chooses a plan for setting the future w's taking into account the best response of the robust controller, who chooses the $u's$ sequentially and regards the w's as an exogenous process. Thus, at time zero the Stackelberg leader determines transition laws for the w's and a larger set of state variables that in turn depend on the best response of the robust controller. Aware of these transition laws (albeit not the future outcomes of these variables) the robust controller can use them for forecasting. These forecasts are distorted relative to the optimal forecasts in the stochastic control problem. HS go on to show that the Stackelberg equilibrium exhibits certainty equivalence and the robust rule can be determined by two separate steps: first, compute the optimal rule for the non-stochastic problem where the future states z are known; secondly, derive expectations of future states from the distorted transition laws implied by the Stackelberg leader's plan and plug them in the deterministic decision rule.

This finding produces several benefits. It provides a new perspective on robust decision making that is particularly appealing for economics given the importance of linear-quadratic models and the separation priniple in standard economic analysis. Furthermore, it allows to use recursive methods for computing robust rules and it provides a model for rationalizing the actions of a robust controller from a Bayesian perspective. Hansen and Sargent deliver powerful tools for analyzing misspecification and exploring the implications of a preference for robustness. For the remainder of this comment I will discuss some limitations of HS' approach to model uncertainty and robustness in more detail.

3. Limitations of HS Robust Control

3.1. Limited Degree of Model Uncertainty or "When Certainty-Equivalence Fails"

The degree of model uncertainty considered by HS in this paper remains rather limited. The transition laws $f(.)$ and $g(.)$ for the exogenous and endogenous state variables are assumed to be known to the decision maker. Taking the case of a linear transition law $g(.) =$

$g_1 x_t + g_2 z_t + g_3 u_t$, this means the parameters $g_{i,i=\{1,2,3\}}$ are treated as known. g_3, for example, measures the effect of the decisions u on state x, that is policy effectiveness. Assuming knowledge of such parameters seems rather unattractive when the objective is to analyze the implications of uncertainty on economic decision making. Parameter uncertainty has been studied extensively in stochastic control in the context of optimal monetary policy as well as consumption and investment dynamics. It is understood that the separation principle for stochastic control will fail under multiplicative parameter uncertainty even with linear models and quadratic objective.

Of course, one can consider a min-max game with uncertainty concerning model structure. In such a game the robust controller tries to guard against averse outcomes regarding nature's choice of parameters governing the transition laws g and f. For example, the robust controller may suspect distortions concerning the policy effectiveness parameter $g_3 + w_{g3}$ and derive a robust rule. The robustness of monetary policy in the presence of structured model uncertainty has been studied by Onatski and Stock (2000), Giannoni (2000), Tetlow and von zur Muehlen (2001) and Zakovic, Rustem and Wieland (2002, 2004a, 2004b). The conclusions differ somewhat across studies. However, Tetlow and von zur Muehlen (2001) as well as Zakovic et al. (2004b) find that a policy that is robust to structured model uncertainty implies a more cautious response to inflation and output fluctuations than in the absence of uncertainty.

While HS focus on fully optimal decision rules, the above authors focus on simple rules that respond to a subset of state variables. It is known from stochastic control that certainty-equivalence fails with simple rules. Concerns about model uncertainty have been an important motivation for the popularity of simple rules in the practice of policymaking. The history of monetary policy has long been characterized by debates regarding simple money growth and interest rate rules due to concerns regarding the dynamics of the economy including long and variable lags in the monetary transmission process. The important contributions of Hallman, Porter and Small and Henderson and McKibbin to this debate have already been mentioned above. It is of particular interest for policy applications to constrain mini-max analysis in this manner. In this case, the robust controller commits to following a simple rule and then chooses the response parameters of that rule according to a min-max criterion.

3.2. Limited Degree of Robust Behavior

In the present paper HS focus on the multiplier problem defined by equation (6). In this problem the parameter θ plays a key role in determining the extent of robust behavior. As θ goes to infinity the minimizing player, i.e. nature, will set the distortions w equal to zero. Thus, the preference for robustness disappears. As θ goes towards zero there is a positive lower value $\underline{\theta}$, which HS call a breakdown point. Beyond this point it is fruitless to seek more robustness within this framework because nature is sufficiently unconstrained so that she can push the criterion function to $-\infty$ despite the best efforts of the robust controller. HS provide an interpretation of θ in terms of risk-sensitive preferences as discussed by Whittle (1990). In the forthcoming monograph HS carefully examine the link between $\underline{\theta}$ in the multiplier problem and the upper bound $\bar{\eta}$ in the constraint problem and derive conditions under which both games result in the identical equilibrium.

Clearly, the lower bound $\underline{\theta}$ or the upper bound $\overline{\eta}$ are key parameters in specifiying the degree of robustness. The approach developed in this regard in the upcoming monograph of HS is impressive. Nevertheless, it seems to me that for many applications, including monetary policy, other specifications of robust preferences may seem more intuitive and more flexible to use. For example, borrowing from Rustem and Howe (2002) an alternative approach is to solve the following min-max problem with box constraints:

$$
\begin{aligned}
\min_{w} \quad \max_{u} \quad & \sum_{t=0}^{T} r(y_t, u_t, w_t), \\
s.t. \quad & \underline{w}_t \leq w_t \leq \overline{w}_t, \\
and \quad & \text{transition law for } y_t
\end{aligned}
\tag{7}
$$

The box constraints on w can be asymmetric and time varying. Thus, they offer highest flexibility in specifying the worst cases a policymaker would like to consider. To derive the robust policy no assumptions regarding the formation of expectations and measurability are needed. Furthermore, the policy derived in this manner can guarantee a lower bound on the gain (or upper bound on the loss) across the whole range of distortions w to be considered. Thus, it provides full insurance for this range. If desired one can constrain the flexibility offered by the box constraints by considering the cost of such insurance. An intuitive measure of this cost would be the implied loss in expected performance that is incurred compared to the optimal stochastic control rule. A policymaker's willingness to sacrifice expected performance in order to purchase insurance cover for a certain range of bad outcomes would determine the desired range to be covered. Zakovic et al. (2004b) provide an example of this type of analysis for a simple model of monetary policy.

As shown by Rustem and Howe (2002) the above min-max problem allows multiple solutions. Their algorithm guarantees cover across all solutions. Multiple solutions may arise if the objective is concave in w. If the objective is convex in w and concave in u it has the saddle point property. The role of the parameter θ in HS is to convexify the objective in the w dimension.

3.3. Learning and The Failure of the Separation Principle

Hansen and Sargent treat model uncertainty as given and assume that nothing can be learned. Such an assumption seems rather unnatural when considering uncertainty regarding model structure such as the parameters of the transition laws $g(.)$ and $f(.)$ discussed above. In practice, a decision maker will attempt to learn more about parameters that govern the impact of his actions on performance. To illustrate this incentive to learn I turn to the linear-quadratic permanent income model of consumption used by HS for illustrative purposes. They consider a consumer who receives the exogenous endowment process z_{2t}, which follows a first-order autoregression. The consumer needs to allocate the endowment between consumption c_t and savings x_t and wants to maximize

$$\begin{aligned} -E_0 \quad & \sum_{t=0}^{\infty} \beta^t (c_t - z_1)^2, \beta \in (0,1) \\ s.t. \quad & x_{t+1} + c_t = R x_t + z_{2t} \\ and \quad & z_{2,t+1} = \mu_d (1-\rho) + \rho z_{2t} + c_d(\epsilon_{t+1} + w_{t+1}) \end{aligned} \tag{8}$$

The constant z_1 is a bliss level of consumption. $R > 1$ is a time-invariant gross rate of return on financial assets x_t held at the beginning of period t, and $|\rho| < 1$ describes the persistence of the endowment process. w_{t+1} is a distortion to the mean of the endowment process that represents possible model misspecification. HS show how concern over model-misspecification introduces a type of precautionary savings.

HS assume that the effect of consumption on utility is known with certainty. At first glance this assumption may seem appropriate. After all, one should know what one likes even if one may never be sure how to afford it. On second thought however, one could think of many examples where the impact of consumption on utility is uncertain. To name one, there exist quite different views on the impact of cigarette consumption on health and thus on utility. Such uncertainty arises with much of what we consume and quite often with regard to the quantity of consumption. To capture this uncertainty one could add an additional state x_2 to the above utility function to measure sickness or ill health and a corresponding transition law relating x_2 to the quantity of consumption:

$$x_{2,t} = \phi_1 x_{2,t-1} + \phi_2 c_t + \epsilon_{2,t+1} \tag{9}$$

Here the impact of consumption on health status is measured by the parameter ϕ_2. A consumer concerned with the relationship between consumption and health will attempt to learn about ϕ_2. One approach would be to vary consumption and thereby improve estimates of this parameter.

Combining learning about structured uncertainty with worst-case analysis is an unexplored area which deserves attention. In the area of stochastic control there exists a literature on learning with regard to unknown parameters. In fact, the above example for health and consumption is reminiscent of the seminal contribution of Grossman, Kihlstrom and Mirman (1977) on learning and consumption. It is well known that the separation principle fails for such learning problems and that control and estimation have to be addressed jointly. As a result, analytical solutions of optimal decision rules are typically not available and theoretical research has largely focused on long-run asymptotic behavior of beliefs. While quantitative approaches to stochastic control with learning (also called adaptive control) are more promising, the involved computational problems are complex due to the presence of nonconvexities (see for example Amman and Kendrick (IER 1995)). For this reason, optimal learning has not yet been widely applied. However, by now there has been sufficient progress concerning computer speed and numerical methods to render a class of learning problems that is general enough for interesting applications readily computable.

I will conclude my comment with an advertisement intended to encourage further application. I will present a generic learning and control problem that can be solved with numerical dynamic programming methods in a few minutes on a standard laptop. This problem has been analyzed in detail by Beck and Wieland (2002). The numerical algorithm is available in Fortran code from the following website: www.volkerwieland.com/software.htm. The decision maker is faced with the following linear stochastic process:

$$x_t = x_{t-1} + \phi u_t + \epsilon_t \quad \text{where } \epsilon_t \sim N(0, \sigma_\epsilon) \tag{10}$$

Beliefs regarding the unknown parameter ϕ are summarized by the mean p_t and the variance v_t. Given a choice for the control u_t and a realization of the shock ϵ_t a new observation x_t becomes available and the decision maker updates his beliefs regarding ϕ according to

$$\begin{aligned} p_t &= p_{t-1} + v_{t-1}(u_t)F^{-1}(x_t - x_{t-1} + p_{t-1}u_t) \qquad (11) \\ v_t &= v_{t-1} - v_{t-1}(u_t)F^{-1}(u_t)v_{t-1} \\ \text{where} \quad & F = (u_t)v_{t-1}(u_t) + \sigma_\epsilon. \end{aligned}$$

The decision maker then maximizes the following objective

$$\underset{u}{\text{Max}} \quad E\left[\sum_{t=0}^{\infty} \beta^t \left(-(x_t - x^*)^2\right) \mid (x_0, u_0, p_0, v_0)\right] \tag{12}$$

$$\text{s.t. transition laws for } x_t, p_t \text{ and } v_t$$

There are three state variables, x_t, p_t and v_t. The transition laws for the beliefs p_t and v_t are nonlinear. Consequently, the separation principle does not apply. This problem can be solved by numerical dynamic programming techniques which are subject to the curse of dimensionality. Computation time will depend on the number of grid points chosen for the three state variables and the Gaussian quadrature with respect to the random shock ϵ.

For this example, the grid has 16 points for p, 21 for v and 26 for x. Expectations with respect to the normally distributed random variable are compute by means of 15 point Gaussian quadrature. The algorithm is coded in FORTRAN and executed on a laptop with

Mobile Pentium III 1000 MHZ. Convergence is achieved in less than 8 minutes, that is, less time than the presentation of this comment at the conference. The implied optimal decision rule is reported in Figure 1, which contains 16 panels. Each panel shows the optimal rule (dotted line) for u as a function of the state x for a given belief (p, v). Each dot corresponds to a grid point. Moving from left to right the panels are associated with greater uncertainty as measured by the variance v. Moving from top to bottom the panels are associated with greater expected values p. Policies are rather smoothly approximated except for an occasional discontinuity at $x = 0$. The figure compares the optimal rule under learning to a certainty-equivalent rule (solid lines) that disregards uncertainty. For low degrees of uncertainty or for medium to large deviations of x from zero the optimal rule implies less responsiveness of the control u to deviations of x than the certainty-equivalent rule. In other words, parameter uncertainty and learning induce a precautionary motive to decisions. For high degrees of uncertainty and very small deviations of the state x from the target value of zero however, the optimal rule implies a higher responsiveness than the certainty-equivalent rule and at zero even active perturbations. At those points learning introduces a degree of experimentation in order to improve information regarding the unknown parameter.

REFERENCES

Amman, H. and D. Kendrick (1995), "Nonconvexities in Stochastic Control Models," *International Economic Review*, Vol. 36, 2, 455-475.

Beck, G. and V. Wieland (2002), "Learning and Control in a Changing Economic Environment," Journal of Economic Dynamics and Control, 26 (9-10).

Giannoni, M. (2002), "Does model uncertainty justify caution ? Robust optimal monetary policy in a forward-looking model ",*Macroeconomic Dynamics*, 6(1), 111-144.

Grossman, S., R. Kihlstrom and L. Mirman (1977), A Bayesian Approach to the Production of Information and Learning by Doing, Review of Economic Studies, 44 (3), 533-547.

Hallman, J. , R. Porter and D. Small (1991), "Is the Price Level Tied to the M2 Monetary Aggregate in the Long Run?," *American Economic Review* , 81, 841-858.

Hansen, L. and T. Sargent (forthcoming), "Misspecification in Recursive Macroeconomic Theory," Manuscript.

Henderson, D. and W.J. McKibbin, (1993), "A comparison of some basic monetary policy regimes for open economies: implications of different degrees of instrument adjustment and wage persistence" *Carnegie-Rochester Conference Series on Public Policy*, 39, 221-318.

Henderson, D. and S. Turnovsky, (1972), "Optimal Macroeconomic Policy Adjustment under conditions of risk," *Journal of Economic Theory*,4 , 58-72 .

Kalchbrenner, J. and P. Tinsley (1975), "On the Use of Optimal Control in the Design of Monetary Policy," FRB Special Studies Paper No 76.

Kalchbrenner, J. and P. Tinsley (1976), "On the Use of Feedback Control in the Design of Aggregate Monetary Policy," *American Economic Review* , 66, Vol. 7, 349-55.

Kalchbrenner, J. and P. Tinsley with J. Berry and B. Garrett, (1977), On Filtering Auxiliary Information in Short-run Monetary Policy, in 'Optimal Policies, Control Theory and Technology Exports', Carnegie-Rochester Conference Series on Public Policy, 7, 39-84.

Onatski, A. and J. Stock, (2002), "Robust monetary policy under model uncertainty in a small model of the U.S. economy",*Macroeconomic Dynamics*,6 (1).

Rustem, B. and A. Howe (2002), Algorithms for Worst-Case Design with Applications to Risk Management, Princenton University Press.

Tetlow, R. and P. von zur Muehlen, (2001), "Robust monetary policy with misspecified models: Does model uncertainty always call for attenuated policy?",*Journal of Economic Dynamics and Control*, 25, 911-949.

Tinsley, P. and P. von zur Muehlen, (1976), "A Maximum Probability Approach to Short-Run Policy," Journal of Econometrics, 15, 1981.

von zur Muehlen, P., (1982), “Activist vs. Non-Activist Monetary Policy: Optimal Policy Rules under Extreme Uncertainty, Manuscript, Federal Reserve Board, Washington, DC, April 1982.

Whittle, P., (1990), Risk-Sensitive Optimal Control, New York: Wiley.

Zakovic, S., Rustem, B. and V. Wieland, (2002),”A Continuous Minimax Problem and its Application to Inflation Targeting, in, Georges Zaccour (ed.), Decision and Control in Management Science, Kluwer.

Zakovic, S., Rustem, B. and V. Wieland, (2004a),”Mean-Variance Optimization of Nonlinear Systems and Worst-Case Analysis,” Working Paper, March 2004.

Zakovic, S., Rustem, B. and V. Wieland, (2004b),”Stochastic Optimization and Worst-Case Analysis in Monetary Policy Design,” Presented at Society for Computational Economics Annual Conference, Amsterdam, July 2004.

FOOTNOTES

* Professur für Geldtheorie und -politik, Johann-Wolfgang-Goethe Universität, Mertonstrasse 17, D-60325 Frankfurt am Main, Germany, phone: +49 69 798-25288, e-mail: wieland@wiwi.uni-frankfurt.de, homepage: http://www.volkerwieland.com.

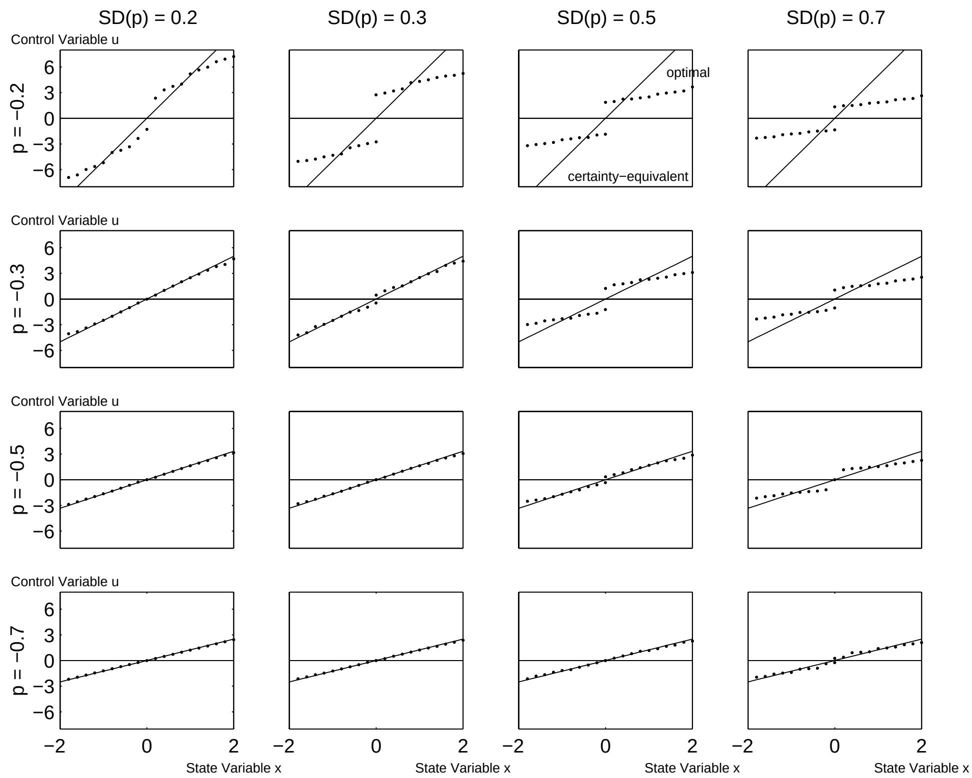

Fig. 1. Approximating an Optimal Learning Rule.