

6/7/02

08:46

Stochastic Control for Economic Models:
Past, Present and the Paths Ahead¹

by

David A. Kendrick

Department of Economics
The University of Texas
Austin, Texas

¹ Prepared for presentation as a keynote address at the Ninth International Conference of the Society of Computational Economics at Aix-en-Provence in June of 2002.

Abstract

On the thirtieth anniversary of a meeting between engineers and economists at Princeton University to discuss the potential application of stochastic control methods to economics, this paper provides a review of past, present and paths ahead. The events surrounding some of the key developments of the past are described along with a discussion of the present state of the research and of paths for future research. The paper is structured around the primary methods for solving quadratic-linear economic stochastic control models, namely open loop, classical control, handcrafted feedback rules, optimal feedback, min-max control, optimal feedback rules with parameter uncertainty and dual control. The paper also includes a discussion of software developments in economic stochastic control modeling.

1. Introduction²

In May of 1972 a meeting of economists and control engineers was arranged at Princeton University by three economists (Edwin Kuh, Gregory Chow and M. Ishaq Nadiri) and a control engineer (Michael Athans). The meeting which was attended by about 40 economists and 20 engineers was to explore the possibility that the application of stochastic control techniques, which had been developed in engineering, would prove to be useful in economics as well, Athans and Chow (1972). At our meeting now of the Society of Computational Economics in Aix-en-Provence, thirty years later, it seems a useful time to look back at the developments of the past, to examine some of the present work and to discuss the paths ahead.

Stochastic control seems at first to be a simple idea of nudging a stochastic process in one direction or another as it flows through time. However, the developments in economics in this field have been multifaceted and it is therefore useful to have a framework for the discussion. The framework to be used here is outlined in Kendrick (2002), which provides a classification system for the stochastic control models that are used in economics. Before outlining elements of that framework it is useful to introduce some of the mathematical notation to be employed. The notation (and most of the discussion) will be limited to systems with quadratic criteria and linear systems equations; however general nonlinear models will be discussed at some places in the paper.³

² I am indebted to Hans Amman, Ray Fair, Ric Herbert, Charles C. Holt, P. Ruben Mercado, Alfred Norman, Berc Rustem, Robert Tetlow, Peter Tinsley, Marco Tucci, Stephen Turnovsky, Peter von zur Muehlen and Volker Wieland for comments and suggestions on earlier drafts of this paper. Responsibility for any remaining errors is my own.

³ A related topic in dynamic economics, which this paper does not cover, is nonlinear dynamics and chaos theory. For discussion of this subject see Chiarella (1988), Brock and Dechert (1991) and Dechert and Hommes (2000).

2. Mathematics of the Basic Quadratic-Linear Optimal Control Models

The criterion function for most quadratic-linear control models in economics is in the form of a tracking function. It contains desired paths for both state and control variables but does not contain cross-terms between states and controls. A distinction is made between terms for the last time period, N , and all other time periods, indexed with k . The criterion function is written

$$J = E \left\{ \frac{1}{2} (x_N - \tilde{x}_N)' W_N (x_N - \tilde{x}_N) + \frac{1}{2} \sum_{k=0}^{N-1} \left[(x_k - \tilde{x}_k)' W_k (x_k - \tilde{x}_k) + (u_k - \tilde{u}_k)' \Lambda_k (u_k - \tilde{u}_k) \right] \right\} \quad (2.1)$$

where

x_k = state vector - an n vector

\tilde{x}_k = desired state vector - an n vector

u_k = control vector - an m vector

\tilde{u}_k = desired control vector - an m vector

W_N = symmetric state variable penalty matrix at terminal period, N

W_k = symmetric state variable penalty matrix for period k

Λ_k = symmetric control variable penalty matrix for period k

E = an expectations operator over the uncertain elements in the model

Sometimes the criterion function is not written in the tracking form, as above, but rather as the more general quadratic form shown below. It contains quadratic and linear terms in the state and control vectors as well as cross-terms between the states and controls, viz.

$$J = E \left\{ \frac{1}{2} x_N' W_N x_N + w_N' x_N + \sum_{k=0}^{N-1} \left[\frac{1}{2} x_k' W_k x_k + w_k' x_k + x_k' F_k u_k + \frac{1}{2} u_k' \Lambda_k u_k + \lambda_k' u_k \right] \right\} \quad (2.2)$$

where

w_N = linear state coefficient vector for period N

w_k = linear state coefficient vector for period k

λ_k = linear control coefficient vector for period k

F_k = coefficient matrix for the $x-u$ cross term

Quadratic-linear computer codes are usually based on the more general form in Eq. (2.2) so quadratic tracking functions in Eq. (2.1) are transformed in the computer codes to the quadratic form before the model is solved. (see Kendrick (1981) pp. 7-8).

The systems equations for the stochastic control model may be written as

$$x_{k+1} = A_k x_k + B_k u_k + C_k z_k + \xi_k \quad k = 0, \dots, N-1 \quad (2.3)$$

where

- x_k = the state vector in period k with n elements
- u_k = the control vector in period k with m elements
- z_k = the exogenous vector in period k with ℓ elements
- A_k = state vector coefficient matrix in period k
- B_k = control vector coefficient matrix in period k
- C_k = exogenous vector coefficient matrix in period k
- ξ_k = vector of additive noise terms in period k

In many cases the z_k vector consists of a single variable that is one in all time periods and the elements of the only column of the C_k matrix are the intercept terms in the system equations. In the additive noise case the only uncertainty is in the ξ_k term and the A_k , B_k and C_k matrices usually contain constant parameters.

In contrast, in the parameter uncertainty case a subset of the parameters in these matrices have true constant values but these true values are unknown to the policy maker who knows only the means and variances of the parameter estimates. This situation is represented in the notation by creating a vector θ_k that contains the uncertain parameters. For example, if the model had three state variables and two control variables the A_k matrix would be 3x3, the B_k matrix would 3x2 and the C_k matrix would be 3x1. If only a few of the coefficients were to be treated as uncertain, viz. the coefficients a_{11} , a_{23} , b_{22} and c_{31} , then the θ_k vector of uncertain parameters would be

$$\theta_k = \begin{bmatrix} a_{11} \\ a_{23} \\ b_{22} \\ c_{31} \end{bmatrix}_k \quad (2.4)$$

We will return to the use of these θ_k vectors later. For now, consider the mathematics for some of the solution methods for stochastic control models.

One approach to solving stochastic control models is to use a handcrafted feedback rule. These rules are crafted without the benefit of optimization methods, but they may nonetheless be carefully crafted based on experience with simulating alternate rules. For the case of macroeconomic applications, these rules are sometimes called “Taylor Rules” viz. Taylor (1999).

In contrast, optimal feedback rules are those computed with optimization methods. These rules can be computed (1) for models with additive noise terms and (2) for models with both additive noise terms and parameter uncertainty. The methods used in the two cases yield feedback rules that are similar and differ only in the use of the expectations operator over the uncertain parameters. In the first of these two methods the optimal feedback rule is written (see Kendrick (1981) Ch. 2) as

$$u_k = G_k x_k + g_k \quad (2.5)$$

where

G_k = the feedback gain matrix

g_k = the vector of feedback parameters

with

$$G_k = -[B'_k K_{k+1} B_k + \Lambda'_k]^{-1} [F'_k + B'_k K_{k+1} A_k] \quad (2.6)$$

where

K_k = the Riccati matrix

and the g_k vector is computed in a similar manner.

In contrast, when there is parameter uncertain the rule is called “optimal feedback with parameter uncertainty” (see Kendrick (1981), Ch. 6) and is written as

$$u_k = G_k^\dagger x_k + g_k^\dagger \quad (2.7)$$

with

$$G_k^\dagger = -[E\{B'_k K_{k+1} B_k\} + \Lambda'_k]^{-1} [F'_k + E\{B'_k K_{k+1} A_k\}] \quad (2.8)$$

and the g_k vector computed in a similar manner. In Eq. (2.8) the E is the expectations operator that is taken over the uncertainty in the parameter estimates in the matrices A_k , B_k and C_k . Also the expectations operator plays a similar role in the computation of the Riccati matrices, K_k . It is significant that the only difference between the optimal feedback rule in Eq.(2.5) and the optimal feedback rule with parameter uncertainty in Eq.(2.7) is the expectations operator used in the calculation of the feedback gain matrix, G_k . The same is also true of the computations for the vector of feedback parameters, g_k .

With this notation in hand it is relative straightforward to describe the main attributes of the stochastic control models which are used in economics.

3. Attributes of Stochastic Control Models in Economics

The two principal attributes used in Kendrick (2002) to classify economics stochastic control models are

- stochastic elements
- solution method

Models with no stochastic elements are called deterministic. Among the models with stochastic elements, the least complex are those with a single uncertain vector, namely additive noise terms, ξ_k , in the systems equations (2.3). More complicated models have uncertain parameters as in Eq. (2.4), measurement errors, uncertain initial state vectors and time-varying parameters.

This paper will use the second attribute, solution method, as the basic organizing principle, but with interlaced discussion of the stochastic elements. The solution methods to be used here are⁴

⁴ With a few exceptions game theory methods are not covered in this paper. For a classic discussion of game theory methods in dynamic models see Basar and Olsder (1999). For a recent example see Vallee, Deissenberg and Basar (1999).

- open loop control
- classical control
- handcrafted feedback rules
- optimal feedback rules
- min-max control
- optimal feedback with parameter uncertainty
- dual control

Open loop control is used in deterministic models in which there is no uncertainty. In this method the model is solved for all time periods at once and there is no use of feedback. Classical control was developed by Nyquist, Bode and other engineers with proportional, integral and differential (PID) feedback rules to control dynamic systems mostly using continuous time systems. As discussed above, handcrafted feedback rules are simple feedback rules, mostly developed in discrete time systems, in which the feedback gain matrix and the vector of feedback parameters are adjusted by hand. They are commonly use in models that have additive uncertainty in the system equations. Optimal feedback methods include the development of Riccati equations to compute rules like Eq. (2.5). Min-Max methods seek to minimize part of the criterion while maximizing another part. Robust control is an example of min-max control. Next up in complexity are models with uncertain parameters that are solved using expectations operators in the Riccati equations and the feedback rules like those in Eq. (2.7). All of the methods except dual control are *passive* learning methods. Passive learning means that parameter estimates are updated each time period; however in choosing the control variables each period no consideration is given to the impact of this choice on learning the parameter values. In contrast, dual control uses *active* learning methods to purposefully perturb the systems equations early in time in order to improve parameter estimates that can be used to enhance performance later in time.

The literature on economic control theory models of these types is very large, so I have chosen here to site only illustrative studies – and in many cases studies that have interesting stories. Thus I begin with, Arthur Bryson, an aeronautical engineer.

4. Open Loop Control

Professor Bryson was offering a course in the control theory in 1966 that caught the attention of a small group of economics graduate students and young faculty members at Harvard. Bryson was one of one those great teachers who could make even hard concepts easy; however, the course was at first rough going for the economists. The notation and the concepts were unfamiliar and the examples were all about airplanes and rockets (Bryson and Ho (1969)). Nonetheless Rod Dobell, Hayne Leland, Stephen Turnovsky, Chris Dougherty, Lance Taylor and I persisted and before long I found myself in excited conversations with my colleagues. I realized that this was the kind of mathematics I had been looking for while I had been an economics graduate student at M.I.T. It was all about dynamics, nonlinearities and uncertainty, which seemed a natural way to formulate many economic problems. Moreover, the results were not steady state solutions obtained with closed form analytic methods but rather were numerical algorithms and Fortran codes to compute the entire time path.

At that time growth theory was much discussed but was mostly limited to steady state solutions of one and two sector models. Lance Taylor and I realized that with numerical control theory methods we could develop models with any number of sectors and compute the entire paths of the development of the economy. Also, we could analyze how the shape of those paths changed as initial or terminal conditions changed. There was numerical work on growth models at the time but it was mostly limited to linear programming. We speculated that we could use nonlinear consumption, production and investment functions and still calculate the solution in a relatively short period of time by taking advantage of the dynamics and using the costate as well as the state equations and the optimality conditions in a conjugate gradient algorithm. This all worked better than we expected and we were able to develop a nonlinear four-sector growth model of the Korean economy in Kendrick and Taylor (1970).

At about the same time two control engineers who had shifted their interest to economics - David Livesey at Cambridge University and Robert Pindyck at M.I.T. - developed macroeconomic control theory models. Livesey (1971) used a continuous time nonlinear model and Pindyck (1972, 1973a, 1973b) used a discrete time model with a quadratic tracking function like Eq. (2.1) and linear systems equations like (2.3) but without the additive noise terms. Pindyck's model had ten equations, some with substantial lags, so

in augmented form it became a model with twenty-eight state variables and was thus the first application of control methods to a sizeable macroeconomic model.

Five years later Ray Fair (1978) developed one of the most innovative applications of control methods to macroeconomics when he created a model that could be used to compare the economic performance of a number of American presidents. I remember hearing of this before I saw the actual results and thinking to myself that President Kennedy, who was widely admired, would certainly do better than President Nixon, who had left office under a cloud. Not so! Ray Fair, true scholar that he is, let the record speak for itself and Nixon came out better than Kennedy. The reason was that Nixon had been battling inflation and that was difficult to do in the model. In contrast, Kennedy had been attempting to decrease unemployment and that was easier to achieve in the model.

All of this work was open loop control because it was deterministic since there was no uncertainty in the models and no explicit use of feedback rules. However, there was an earlier line of work that would, along with these open loop models, lay the foundations for stochastic control models in economics. This was the use of classical control and of handcrafted feedback rules for both deterministic and stochastic models.

5. Classical Control

The use of classical control in economics has an interesting provenance. The story begins in Java at the end of the Second World War when a young New Zealand engineer by the name of A. W. H. Phillips was released from a Japanese prison camp. If you have seen the movie, *The Bridge on the River Kwai*, you can imagine what his experience must have been like in that camp. After his release Phillips went to England and first took up the study of sociology and then later of economics at the London School of Economics.

He began the construction of a series of connected glass tube models where the flow represented GNP and the feedback system represented the use of monetary and fiscal policy. He most likely conferred about this with Arnold Tustin, a control theorist at Imperial College. As mentioned above, engineers called these PID controllers for proportional, integral or differential. Thus the feedback rule may depend on some combination of proportional, integral and differential values of the state variables.

Mostly these methods are used in continuous time models. I understand that one or more of Phillip's hydraulic models still exist and I hope that they will be preserved carefully. One of Phillips papers about this work is Phillips (1954)⁵ and Tustin's views are summarized in Tustin (1953).

6. Handcrafted Feedback Rules

In the 1970's another British-trained control engineer, Anthony Healy, who was teaching at the University of Texas at the time, took an interest in economic models and applied the use of feedback rules to a well-known model that had been developed at the St. Louis Federal Reserve Bank. He and his student wrote down the feedback rule in the form of Eq. (2.5)

$$u_k = G_k x_k + g_k$$

and then did simulations of the model. Between each simulation they adjusted the parameters in the feedback gain matrix, G , and the vector of feedback parameters, g , while searching for those values which gave more desirable performance, viz Healey and Summers (1974). Thus their approach built on the use of feedback rules from PID controllers but some of their work moved in the direction of handcrafted feedback rules. An individual using this approach does not use Bode-Nyquist diagrams to design a classical controller but rather writes down a feedback rules like Eq. (2.5) and tries various values in the feedback gain matrix, G , and the vector of feedback parameters, g , while doing simulations of the systems to evaluate the performance of the various rules. Thus this was a halfway house between classical control and the optimal feedback rules to be discussed later in this paper.

Also around this time there was work at the Federal Reserve Board of Governors on the use of feedback rules in monetary policy, viz. Kalchbrenner and Tinsley (1976). However this approach was apparently not adopted for regular policy use by the FRB at this time. The reason was concern by some that the underlying criterion implied direct policy objectives for inflation and output/employment rather than for variables that were under more direct control of the FRB such as monetary aggregates. This is reflected in

⁵ See Leeson (2000) for Phillips collected works as well as a short biography.

the paper by Tinsley, von zur Muehlen and Fries (1982) that focused on money market models to control monetary aggregates and included volatility tradeoff curves between money and interest rates.

Later Karakitsos and Rustem (1984, 1985) used parameterized feedback rules for nonlinear macroeconomic models. However, widespread interest in handcrafted feedback rules did not develop until recently and this has been primarily due to the work of John Taylor (1993b, 1999). Interestingly enough, John wrote his dissertation about thirty years ago on the most complex of the solution methods we will discuss here, namely dual control, but he has become famous for his work on the least complex method, namely the handcrafted rules which are now widely called “Taylor” rules. A good example of this kind of rule is in Henderson and McKibbin (1993).

There is an intermediate ground between purely handcrafted rules and the optimal rules that are discussed later in the paper. In this approach the modeler chooses to include only a subset of the states in the feedback rule but then optimizes the choice of the parameters in that feedback rule. For example in a model which might have states lagged several periods the modeler will include only the current states and none of the lagged states in the feedback rule; however a gradient method or some other optimization procedure is used to determine the parameters of the feedback gain matrix, G , and the vector of feedback parameters, g . Examples of such rules are Currie and Levine (1993, Ch. 6), Williams (1999) and Tetlow and von zur Muehlen (1999).

Two other recent examples of work on handcrafted feedback rules are those by Tetlow (2000) and Levin, Wieland and Williams (1999). The latter paper takes a step in the direction of dealing with uncertainty about model specification by identifying optimized simple rules that are robust in the sense that they perform well across a range of diverse but plausible macroeconomic models.

The future of handcrafted rules is that they will most likely get better and better with experience until they provide performance which is almost as good as that obtained with optimal feedback rules, while maintaining somewhat simpler forms. In models with long lags it seems likely that a substantial part of the improvement will come from incorporating lags into the handcrafted feedback rules just as they are automatically incorporated into the optimal feedback rule when the state vector is augmented to include the lagged state variables.

Before moving too rapidly forward though let me go back in time once again to pick up the tread of the story in the development of optimal feedback rules.

7. Optimal Feedback Rules

In the 1950's a young but very distinguished-to-be group of economists (Charles Holt, Franco Modigliani, John Muth and Herbert Simon), who were later to win a bevy of Nobel prizes, came together at Carnegie-Mellon University in Pittsburgh. Holt had come from an engineering background at M.I.T. and Simon's father had a background in mechanical engineering methods. The four developed control methods and applied them to microeconomics by computing variables for production, inventories and the labor force in a firm. Their solutions were in the form of linear decision rules where production, for example, at a point in time was made a linear function of past inventory levels.

The foursome was eager not only to develop the theory and mathematics of this subject but also to demonstrate how their ideas could be put to work in an actual enterprise. So they searched around in Pittsburgh until they found a paint factory that was willing to supply them data. The result was one of the earliest uses of control methods in economics, namely Holt, Modigliani, Muth and Simon (1960).

Another product of this work - though it was published earlier in time - was Herbert Simon's (1956) certainty equivalence theorem that was a precursor to the work in stochastic control theory with quadratic linear systems and additive noise terms. In related work on the other side of the Atlantic at this time Henri Theil (1957) who was working in methods of economic planning for the Dutch government using quadratic linear dynamic models.

After the completion of the paint-factory work, Holt (1962) turned his attention to the use of linear decision rules in macroeconomic models. He developed a model that was quadratic in the criterion function and linear in the systems equations to analyze fiscal and monetary policy.

All of these developments with optimal linear decision rules can be thought of today as optimal feedback rules which are computed using dynamic programming methods on quadratic-linear systems that yield Riccati equations, which are used to obtain the key components of the feedback gain matrix G_k and vector g_k as discussed above.⁶ This approach is sometimes called “modern control” to distinguish it from “classical control”.

One source for a large set of recent models that can be solved for optimal feedback rules is the Duali software (Amman and Kendrick (1999c)). This set contains additive noise versions of many well-known models such as those of Pindyck (1973a), Abel (1975), MacRae (1972) and Hall and Taylor (1993) [in Mercado and Kendrick (1999)]. All of these can be solved for Riccati equations to determine optimal feedback rules. For another example of recent work applying optimal feedback rule to macroeconomics see Woodford (1999).

An innovative recent paper in this area advances across the attributes of stochastic elements to cases that involve measurement error and lags. The study by Coenen, Levin and Wieland (2001) considers the case in which there are both forward variables and lags in the systems equations. The lagged terms can be modeled in the control framework by augmenting the state vector to include those terms (see Kendrick (1981) Ch. 2). Thus in a model with three-period lags the augmented state would be

$$z_k = \begin{bmatrix} x_k \\ x_{k-1} \\ x_{k-2} \\ x_{k-3} \end{bmatrix} \quad (7.1)$$

and the feedback rule would be

$$u_k = G_k z_k + g_k \quad (7.2)$$

⁶ For a discussion of a dynamic programming approach to solving nonlinear models see Rust (1996).

Therefore the control variable would be determined by the feedback rule as a function of the current state vector as well as by that vector lagged once, twice and thrice. Of course, due to revision of the data the states with the longest lags most likely have the smallest measurement error. Thus the optimal controller must consider the fact that one would like to feedback most heavily on the current state; however it is the noisiest. Therefore the feedback rule must strike a balance between depending on the most recent state of the economy and the state with the least measurement error.

8. Min-Max Control

The min-max solution method has been used for some years in economic modeling, viz von zur Muehlen (1982), Rustem (1992, 1994, 1998) and Rustem and Howe (2002) but has gained recent prominence in the form of robust control, viz Hansen and Sargent (2001). In the previous methods that we have discussed the criterion function is either minimized or maximized. In contrast, in min-max control the criterion is minimized with respect to one set of control variables and maximized with respect to a second set. In robust control this second set of control variables may also be thought of as additive noise terms that are used to represent the degree of uncertainty in the specification of the model. Alternatively this second set of control variables is sometimes thought of in a game theory context as the force of nature that attempts to overcome the best intent of the policy maker. Thus the systems equations for a robust control model may be written by using Eq. (2.3) and adding a second set of noise terms so that this equation becomes

$$x_{k+1} = A_k x_k + B_k u_k + C_1 w_k + C_2 z_k + \zeta_k \quad (8.1)$$

where

w_k = an n vector of noise/control terms, i.e. the second set of control variables

In this case one can augment the control vector to

$$\hat{u}_k = \begin{bmatrix} u_k \\ w_k \end{bmatrix} \quad (8.2)$$

and rewrite Eq. (8.1) as

$$x_{k+1} = A x_k + \hat{B} \hat{u}_k + C_2 z_k + \varepsilon_k \quad (8.3)$$

where

$$\hat{B} = [B \quad C_1] \quad (8.4)$$

Also the criterion function in Eq. (2.1) can be used with modification of the Λ matrix i.e.

$$J = E \left\{ \frac{1}{2} (x_N - \tilde{x}_N)' W_N (x_N - \tilde{x}_N) + \sum_{k=1}^{N-1} \frac{1}{2} \left[(x_k - \tilde{x}_k)' W_k (x_k - \tilde{x}_k) + (\hat{u}_k - \tilde{u}_k)' \hat{\Lambda}_k (\hat{u}_k - \tilde{u}_k) \right] \right\} \quad (8.5)$$

where

$$\hat{\Lambda} = \begin{bmatrix} \Lambda & 0 \\ 0 & -\nu I \end{bmatrix} \quad (8.6)$$

and ν is a scalar that represents the degree of uncertainty aversion in the robust control model. Also I is the identity matrix with the same row and column dimensions as the w vector. Since ν has a negative sign the criterion embodies min-max control. The code attempts to minimize the criterion with respect to all control variables. However, because of the minus sign on the ν , the function is effectively minimized with respect to all other arguments and maximized with respect to the argument affected by the ν , which because of the identity matrix is the sum of the w terms.

Also it seems straightforward to combine ordinary expected value controller like “optimal feedback with parameter uncertainty” with min-max control. This is so because, in the limiting case, when ν approaches zero the model reverts to an expected value controller while, as ν becomes large, the degree of uncertainty in the robust control part of the model grows larger.

Marco Tucci and I experimented with a robust control model in the spirit of the Hansen, Sargent and Tallarini (1999) model last year. In that model as ν gets larger, which represents increases in uncertainty in their framework, the w variables become much more vigorous. This in turn causes the original control vector u to be much more vigorous in an attempt to offset these noise-like terms. In contrast, in expected value controllers as the uncertainty in the B matrix of parameters multiplied by the controls increases one expects in general from the Brainard (1967) results that the control may become more cautious. So robust control methods and expected value controller appear to give opposite advice in the face of increases in uncertainty. Each kind of advice is probably appropriate in different cases. For example, as uncertainty about a bridge’s structural members grows it seems sensible to make those structures stronger. On the other hand, as uncertainty about the response of the economy to changes in the money supply grows it would seem to be sensible to be more cautious.

Last year at the Yale conference of the Society of Computational Economics, Rustem, Wieland and Zakovic (2001) presented a paper on macroeconomics and robust control. In addition, Rustem and Howe (2002) include in their book a discussion of min-max control methods in which there is a multiplicative set of noise terms. This leads to general nonlinear formulations of the min-max problem. Also, Giannoni (2002) uses min-max on the values of parameters rather than on additive noise terms and Tetlow and von zur Muehlen (2001a, 2001b) use a variety of robust control methods and provide a comparison across these methods.

For the future one can expect that there will be work to sort out the situations in which expected value controllers or min-max controllers have comparative advantage.

Also, for the future we might expect to see more use of an alternative approach to model uncertainty that focuses on the use of competing models. This approach, which was

pioneered by Kalchbrenner and Tinsley (1977), focuses on pooling models with different frequencies of observation and the design of policies for the pooled models.

9. Optimal Feedback with Parameter Uncertainty

When econometricians estimate macroeconomic models they customarily remove those variables whose coefficients are not significantly different from zero. However, when the resulting model is simulated, no distinction is made between those variables whose coefficients are barely significant and those that are strongly significant. Thus the information in the parameter covariance matrix is left behind in the transition from the estimated model to the simulated model. Stochastic control methods recover and use this information. As discussed above, the difference in the feedback rule between models with and without parameter uncertainty is shown in Eqs. (2.5) and (2.7), i.e.

$$u_k = G_k x_k + g_k \quad (9.1)$$

with

$$G_k = -[B'_k K_{k+1} B_k + \Lambda'_k]^{-1} [F'_k + B'_k K_{k+1} A_k] \quad (9.2)$$

for the models without parameter uncertainty and

$$u_k = G_k^\dagger x_k + g_k^\dagger \quad (9.3)$$

with

$$G_k^\dagger = -[E\{B'_k K_{k+1} B_k\} + \Lambda'_k]^{-1} [F'_k + E\{B'_k K_{k+1} A_k\}] \quad (9.4)$$

for the model with parameter uncertainty.

The difference between these two feedback rules is only in the expectations operators on the $B'KB$ and the $B'KA$ terms. However, this small difference may be very important. If a variable has a coefficient in the B matrix with a relatively small t-test then the variance for that coefficient will be large and the $E\{B'KB\}$ term will be large. Therefore the corresponding elements in the feedback gain matrix G will be small. Thus the policy encapsulated by the feedback rule may rely less heavily on variables whose coefficients are very uncertainty than on those whose coefficients are highly certain. This is the basis for the general notion that the control from the second feedback rule will

be more cautious than the one from the first rule. However, this caution relationship is more complicated than it would at first appear and may be violated even in some simple cases depending on the number of states and controls and the timing of the uncertainty, viz. Craine (1979). For extensions and recent discussion of these issues see Mercado and Kendrick (2000) and Mercado (2001).

Two of the ways that the mathematics for feedback rules came into economics was via the book of Aoki (1967) and the article of Farison, Graham and Shelton (1967). Also, the first application of this kind of thinking in macroeconomics was by Brainard (1967). This was followed by a large number of studies including those of Shupp (1972) (1976), Henderson and Turnovsky (1972), Chow (1973, 1975), Kendrick and Majors (1974), Tinsley, Craine and Havenner (1974), Turnovsky (1973)(1975) (1977), and Craine, Havenner and Tinsley (1976). Many of these papers use quadratic-linear models. For a recent paper on a nonlinear model with parameter uncertainty see Fair (2002b). Also, most of these paper are applications to short run macroeconomic stabilization and use discrete time models. For work on continuous time models of stabilization see Turnovsky (1973) and of growth theory models (both continuous and discrete time) see Turnovsky (2000).

There was even a substantial study of this sort underway at the Federal Reserve Bank in Minneapolis. Then something happened to derail this important work to analyze the effects of uncertainty in economics. The influence of Robert Lucas's (1976) development of the ideas labeled "rational expectations" became great including the notion that announcements of policy rules would result in changes in behavior and therefore invalidation of the feedback rules.

Another aspect of Lucas's rational expectations ideas is embodied in models that include forward variables as shown below in an example with two-period leads, i.e.

$$x_{k+1} = A_k x_k + B_k u_k + C_k z_k + D_1 x_{k+1|k}^e + D_2 x_{k+2|k}^e + \zeta_k \quad k = 0, \dots, N-1 \quad (9.5)$$

where

$x_{k+1|k}^e$ = the expected value of the state variable at period $k+1$ as projected from period k

D_1 = forward looking variable parameter matrix for the $k+1|k$ variables

D_2 = forward looking variables parameter matrix for the $k+2|k$ variables

Thus the state variable in period k is a function of the expected value of future state variables as well as of the lagged state and the control variable.

Since the idea of forward variables like those above was of obvious importance in dynamics and control, we invited Lucas to give a talk at the Society of Economic Dynamics and Control (SEDC) conference in Cambridge, Mass in 1975. He declined to come, but Finn Kydland did accept our invitation. Finn's talk at that meeting was well attended and listened to carefully. The reaction of the control engineer standing next to me was typical – “no problem – you just have to treat it like a game theory problem”. However, it was a problem in the minds of many! As a result, work on control theory models in general and stochastic control models in particular went into rapid decline and remained that way for a substantial time.

In my judgment it was a terrible case of “throwing the baby out with the bath water”. The work on uncertainty (other than additive noise terms) in macroeconomic policy mostly stopped and then slowly was replaced with methods of solving models with rational expectations and with game theory approaches.

At the time that Lucas's ideas began to take hold, I felt that the notion of using forward variables in macroeconomic models as in Eq. (9.5) was a potentially useful advance; however, to turn our backs on uncertainty other than additive noise terms and toward rational expectations and game theory was in my judgment a great loss. There was a time when it was very difficult to publish any macroeconomic models – no matter the technique or other advances involved – so long as that model did not contain rational expectations. While it seems reasonable to me that econometric surveys will show that

some forward variables play a significant role in equations like investment and wage and price formation, I believe that the jury is still out on the strength of these effects and think that there was a substantial over-reaction by economists when these ideas first became popular.

Eventually the new methods by Blanchard and Khan (1980), Fair and Taylor (1983), Anderson and Moore (1985), Oudiz and Sachs (1985), Fisher, Holly and Hughes Hallett (1986), Sims (1996), Juillard (1996), Zdrozny and Chen (1999), Binder and Pesaran (2000) and others for solving rational expectations models took hold. At first these methods were applied to deterministic models and to models with additive uncertainty, viz. Currie and Levine (1993a). Later it was shown that these methods could also be applied to models with uncertain parameters, viz Amman and Kendrick (2000). In the meantime computers had become much faster and it now became apparent that one could apply stochastic control methods with parameter uncertainty to relative large models like the Korean model used by Lee (1998).

The increased speed of computers had also meant that it was possible to do large numbers of Monte Carlo runs on macroeconomic models and to address carefully the following kind of question.

Is macroeconomic performance likely to be better on average if policy makers explicitly consider parameter variances and covariances when designing policy rules?

Amman and Kendrick (1999b) were able to do 10,000 Monte Carlo runs on a small and simple macroeconomic model to address this question. Their results can be encapsulated in a plot of pairs of criterion values from optimal feedback and optimal feedback with parameter solutions across Monte Carlo runs. An example of such a plot, which we called an “asymmetric roman candle”, is shown in Fig. 9.1 below.

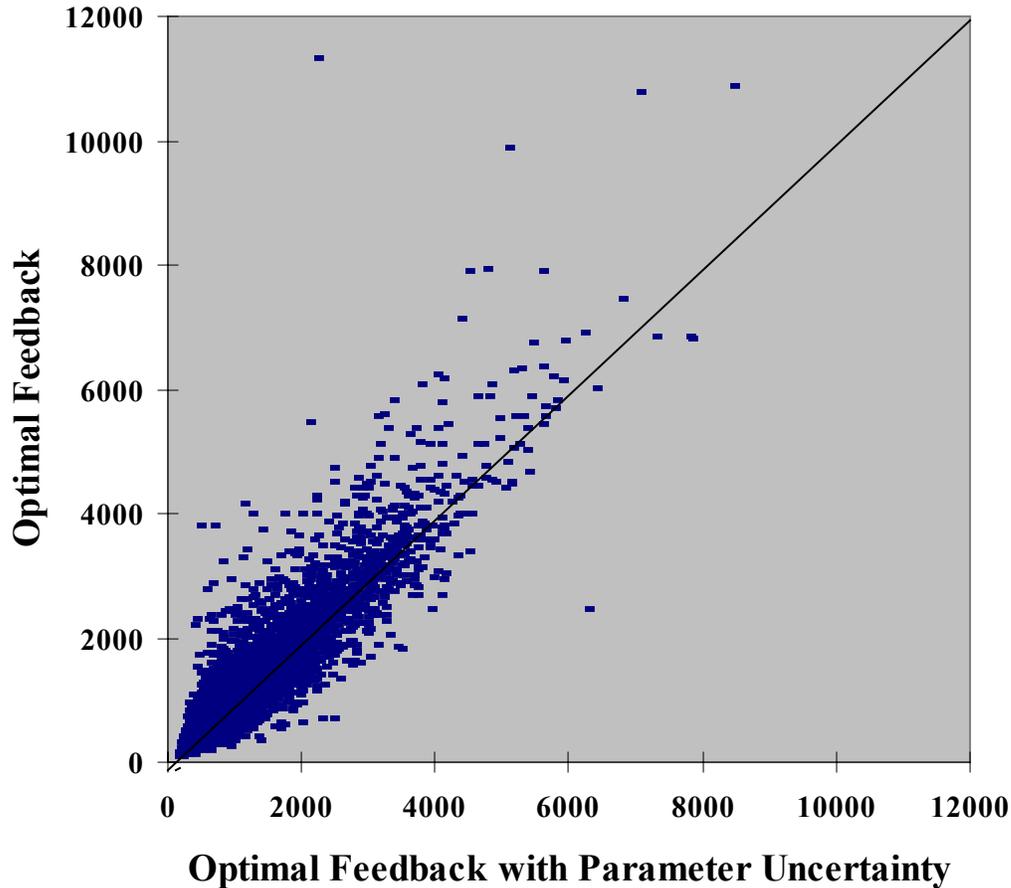


Figure 9.1 Comparison of Criterion Values Across Monte Carlo Runs

A roman candle is a popular kind of firework for children that spurts a fountain of sparks every few seconds. In the figure above if the sparks are symmetric about the 45 degree line, then the optimal feedback (OF) rule like Eq. (9.1) and the optimal feedback rule with parameter uncertainty (OFP) like Eq. (9.3) give basically identical results. However, as the figure shows, Hans and I found that the points are asymmetrically distributed. Thus for a majority of the Monte Carlo runs the criterion values for the optimal feedback solution is greater, and thus worse, than the criterion value of the corresponding solution for optimal feedback with parameter uncertainty. This shows that, with this model, the conservative path for policy makers is to include the variance and covariance information when computing the feedback rule and thus to have a lower criterion value.

What is the future in this part of stochastic control research in economics? I think that the question of whether or not the sparks above will be symmetrically or asymmetrically distributed in most economic models will take some time to determine. Lee's (1998) results on the middle-sized model of Korea reinforce the asymmetric outcome; however Mercado and Kendrick (1999) did some experiments on a version of Hall and Taylor's (1993) popular textbook macroeconomic model and found the opposite result. For theoretical reasons, it seems clear that in the long run it will be wise to include parameter variances and covariances when computing macroeconomic policy rules; however, we have a long way to go before we know the magnitude of this difference in a wide variety of models with different sizes and specifications.

There is a second area in which I think that future work would be valuable. As discussed above, the reaction at the Cambridge SEDC conference to Finn Kydland's talk in 1975 was that game theory could be employed along with stochastic control methods to capture the changes in the behavior of individuals in response to changes in policies. However, there is a simpler approach that I believe holds great promise. This is to use the Kalman filter as Hans and I have done in our Amman and Kendrick (2002) paper. The idea is to use time-varying parameters in the models - and note that this is not only time varying parameter estimates but also that the true values of the parameters are changing over time. Then when the policy rule is changed, some of the behavioral parameters will also change. However, these changes can be tracked by using the Kalman filter so that when a new feedback rule is computed in the next time period it will be based on recent behavior. In this setting the policy maker will always be a step behind the changes in behavior of the economic agents - but only one step behind. So if the changes in behavior are relatively small, this policy procedure may be close enough.

So far in our studies Hans and I have found that the Kalman filter approach does well in comparison to a hypothetical approach in which the policy maker has "insight" as to exactly how a change in policy will affect behavior. However, results will most likely differ across models with different specifications, so we will need to learn what types of models require full-blown game theory methods and what types of models need only the simpler Kalman filter approach.

Another approach to dealing with some aspects of rational expectations is to employ structural rather than reduced form models. A recent critical analysis of small structural macro models can be found in Kozicki and Tinsley (2002). Alternatively one may use

models with time varying parameters in the Swamy and Tinsley (1980) tradition. The use of time varying parameters will be discussed more in the dual control context below.

10. Dual Control

In all the types of stochastic control discussed so far the control is used for a single purpose, namely to move the system in desired directions. In contrast, in dual control the policy variables are used for two different purposes. The first is to move the system in desired directions and the second is to perturb the system in ways that enhance learning of the uncertain parameters.

As was discussed above, dual control is sometimes called “active learning” to distinguish it from “passive learning”. The choice of control variables in the optimal feedback with parameter uncertainty method discussed above is passive learning because the choice of the control variables in each period does not consider the impact of that choice on the improvement of the parameter estimates. In contrast, in dual control the control variables are used actively to gain information.⁷

The early work in economics on dual control was done by Prescott (1972), Taylor (1974) and MacRae (1972, 1975). Later Fred Norman and I teamed up with two control engineers, Edison Tse and Yaakov Bar Shalom, to apply the Tse-Bar Shalom approach (1973) to econometric models. Fred created a first-order dual control without measurement errors in Norman (1976) and then later addressed the problem of computational complexity of alternative stochastic control formulations in Norman and Jung (1977) and Norman (1981, 1994). I used the second order dual control approach of Tse and Bar-Shalom with measurement error in Kendrick (1981, 1982).

Fred and I decided early in the project that the best way to be sure we understood the engineering control methods was to derive all the mathematics afresh and to write the computer code anew from scratch. Then, as a further caution, we decided that each of us would do the computer coding separately so that we could check results with one another even though we were using slightly different algorithms. My memory - though Fred's

⁷ For a different approach to learning with a focus on learning the states rather than the parameters see Herbert (1998).

may differ - is that this went well for a time, but then at an advanced stage of the project one of us phoned the other to compare the latest results. To our chagrin, the results were different. We got off the phone and began the tedious process of tracking down the error. After a time neither of us had found a problem. Then after a longer time, with more determined digging, neither of us found an error. A bug could not be found in either code.

The problem was in the value of the cost-to-go for the first time period in a model with a single control variable. The gradient methods that the two of us were using were converging to different solutions. Eventually we hit on the idea of using a grid search rather than a gradient procedure and we plotted the cost-to-go value for a range of control variable settings. This yielded a curve roughly like the one shown in Fig. 10.1 below.

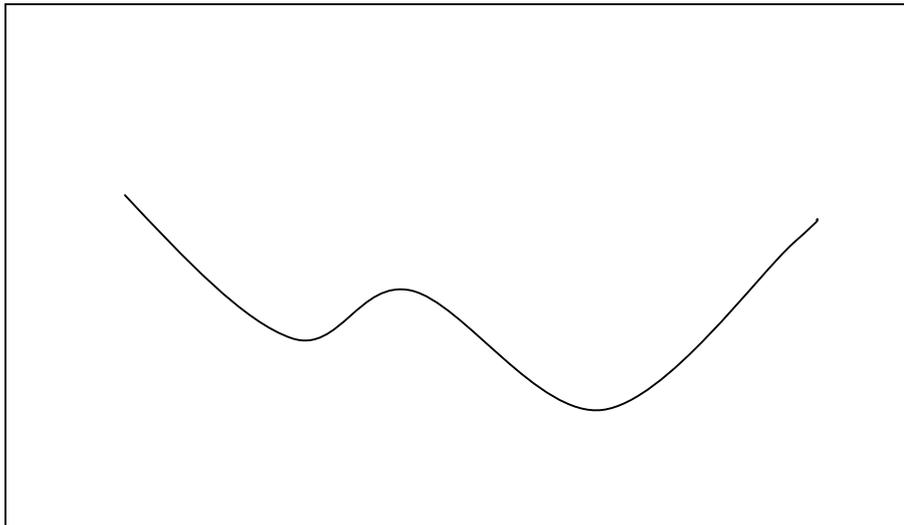


Figure 10.1 The Optimal-Cost-To-Go Plotted Against the Value of the First-Period Control

The function was non-convex and we had each found one of the local optima! Thus both codes were most likely correct.

After a time of rejoicing that we had found the problem, we were sobered by the realization that we might regularly be facing problems in which the criterion function was nonconvex so that simple gradient methods could not be relied upon to find the global optimum. However, we could not be sure that we had not both made some kind of an

error that caused the non-convexities, Norman, Norman and Palash (1979) and Kendrick (1978). It was only some years later in work by Bruce Mizrach (1991) and by Hans and myself in Amman and Kendrick (1995) that we begin to understand that non-convex criterion functions could be endemic to dual control and might be expected to occur in many cases.

This was later confirmed by Volker Wieland (2000a, 2000b, 2001) who found the same kind of nonconvexities using an entirely different approach. Volker's work follows in the line of research on optimal control with learning of parameters developed by John Taylor in Taylor (1974) with later contributions by Easley and Kiefer (1988), Kiefer and Nyarko (1989), Aghion, Bolton, Harris and Jullien (1991) and others.⁸ Volker derives an approximation to the optimal policy function that turns out to be discontinuous with jumps at the points in the state space corresponding to the nonconvexities in Figure 9.1.

In our work around 1995, Hans and I had found that the nonconvexities in the cost-to-go in dual control problems were much more likely to occur in the first few time periods than in the later time periods. The reason for this was that the elements of the covariance matrix of the parameters for the parameter vector θ discussed above, i.e. $\Sigma^{\theta\theta}$, was likely to be large in the first time periods and then to decline as learning occurred. This characteristic prompted Marco Tucci (1998) to develop an innovative approach to finding global optimal solutions. Marco used multiple starting points as a test to see if nonconvexities were present. If they were present, the search was intensified and, if not, the code quickly moved on to the next time period. Since comparison of methods in stochastic control is usually done with Monte Carlo runs, the speed of the code is most important and Marco's approach seems to offer an important advantage. A second and improved version of this approach is described in a more recent paper – Tucci (2002).

Then what about the paths ahead in dual control? Though some progress has been made on dealing with nonconvexities in dual control, much remains to be done. We do not yet know whether nonconvexities will occur in most economic models or only in models with certain kinds of specification or of parameterization. Also, there is considerable active research in the mathematical programming field on ways of finding global optima.

⁸ For a discussion of this line of research see Beck and Wieland (2002).

Hopefully some of these methods will prove to be highly efficient when applied to dual control models.

In addition to nonconvexities, a major area of research in dual control in the future is likely to be models with time-varying parameters. In all stochastic control models where there is passive or active learning, parameters are updated each time period so the *estimates* of the parameters are time varying. However, in almost all stochastic control models used so far the true parameter values have been constant. In contrast, in the early work of Sarris (1973) and the more recent work of Tucci (1989, 1997) the true parameter values are considered to be time varying. Since it seems reasonable that the parameters of human behavior change over time, this is likely to turn out to be a most important direction for research in this field.

A third area for the future comes from the following. One of the key questions with dual control is whether or not these solutions are significantly better than optimal feedback with parameter uncertainty or even than ordinary optimal feedback solutions. I believe that substantial effort will be required to nail down the answer to this question. The mathematics of this subject is not difficult, but it is complex, so checking and rechecking of derivations and different ways of doing the derivations will be required before we can be sure that no mathematical errors remain. Also, the computer coding is even more complex than the mathematics. In econometrics there are multiple computer codes available so they can be checked against one another. In dual control there are, as yet, only a few codes and checking even those against one another is a time-consuming and difficult task.

Against this background, Hans and I have found by doing 10,000 Monte Carlo runs on a set of small macroeconomic models that dual control was better than optimal feedback with parameter uncertainty in most of our models. Also, optimal feedback with parameter uncertainty was better than simple optimal feedback, see Amman and Kendrick (1994, 1997). It could be that some specifications of models will maintain this ordering and others will not – that remains to be seen and is a major path ahead for the research.

Once we get the math and the computer coding straight, a most interesting question can be addressed. This is the question of when it is sensible to purposefully perturb an economic system to gain information. It seems likely that in some microeconomic or

public finance cases it will be useful to do perturbations. For example, when a new product is launched one might do well to learn the coefficients of the demand function by varying prices over time. Similarly, in public finance one might do well to learn optimal levels of taxes or subsidies by perturbing these early in time to gain information.

However, in macroeconomics perturbation will be more problematic. Would it ever be appropriate for the Fed to purposefully perturb interest rates, even in the slightest amount, in order to gain better parameter estimates? Most likely not; however I can remember the response of one control engineer when I raised this question with him some years ago. He replied that while it might be irresponsible to perturb interest rates, it could be even more irresponsible not to do so in situations where policy makers were in danger of using seriously incorrect estimates.

This completes the discussion of the methods for solving economic stochastic control problems. However, before closing it is useful to consider briefly the software systems that are used to solve this class of model.

11. Software

One of the early breakthroughs in software for control theory models was the TROLL system developed at M.I.T. under the leadership of Edwin Kuh. This system was used by Bob Pindyck (1972, 73a, 73b) to solve his quadratic linear control model of the U.S. economy. This system remains in widespread use for estimation, simulation and optimization of economic models.

In spite of the advances of systems like TROLL, old reliable Fortran was, is and most likely will continue to be a language that is widely used by economists when solving stochastic control models.⁹ As described above, Fred Norman and I each wrote our own versions of the Tse and Bar-Shalom dual control algorithm in Fortran. My version was given the name Dual and has been used by several generations of graduate students in developing dissertations. One of those students was Hans Amman, who while a graduate student and young faculty member at the University of Amsterdam, rewrote the Dual code in Fortran 77 and adapted it for use on personal computers. This code, which is

⁹ For example, Ken Judd makes available Fortran code for most of the methods – including those for dynamic optimization - covered in his book on numerical methods, Judd (1998).

called Dualpc (Amman and Kendrick (1999c)), is still in use today for solving dual control models. Later a variant of the Dualpc code was created by Marco Tucci (1989) for models with time-varying parameters. This code was the vessel for Marco's work on global optimization that was discussed above.

Present day stochastic control software can be divided into three groups:

1. compiler-driven systems
 - a. algorithmic systems
 - GAUSS - Breslaw (1994), Lin (2001)
 - MATLAB - Math Works(1999)
 - b. modeling systems
 - GAMS - Brooke, Kendrick and Meeraus (1982)
2. menu-driven systems
 - Duali - Amman and Kendrick (1999c)

Most of the software systems commonly used by economists are compilers with an associated language. However, there is an important division within this group between algorithmic systems and modeling systems.

Algorithmic systems give the user the most freedom to use any algorithm that he or she desires to solve the economic models. For example, I use in my graduate computational economics class a Riccati system developed by Hans Amman in GAUSS. MATLAB is one of the languages used by Gary Anderson (Anderson and Moore (1985)) to make available his highly efficient system for solving large-scale control models with forward variables. Hansen and Sargent (2001) have used MATLAB to provide robust control solution procedures. Hans and I have used MATLAB for stochastic control models with forward variables, viz those in Amman and Kendrick (1999a, 2000).

In contrast, modeling systems like GAMS lower the entry cost tremendously by removing the necessity of writing the code for algorithms. GAMS takes the focus off of writing code for algorithms and puts it on model development. The GAMS system comes with a substantial number of optimizers embedded in the code to solve linear and nonlinear optimization problems. So all the user has to do is to write the code to describe the model and then call one or more of the optimizers. For example, Mercado, Kendrick

and Amman (1998) used the Hall and Taylor (1993) textbook model and the Taylor (1993a) model to illustrate the use of GAMS for solving macroeconomic models.

Finally, consider the menu-driven software systems. Over the years I have been impressed with the high cost in time and energy that must be paid by graduate students to learn the computer codes which are used to solve stochastic control economics models. Worse yet, I have wanted to introduce undergraduates to the use of optimal control methods. Also, I have noticed how dependent academic researchers in this field become on their programmers and how much momentum they lose when a first-rate graduate student programmer completes his or her studies and moves away. One solution to these problems is the development of menu-based stochastic control systems like Duali (Amman and Kendrick (1999c)). With this kind of a system menus and dialog boxes are used to permit the user to select options for the specification and solution of models and the visualization of the results. Such systems enable economists who are very good at microeconomics or macroeconomics, but who have little interest in computer programming, to develop high-quality stochastic control models. On the other hand, menu-driven systems are less flexible when the user wants to employ a new algorithm. Also, menu driven systems are less transparent and it is therefore more difficult to be sure the code is working properly and to make changes if it is not.

Hans Amman and I have developed a working procedure in the last few years that uses both compiler-driven and menu-driven software systems. We prototype our models in MATLAB and if the method proves the test of time, then we program it into the Duali software. This system allows relatively quick coding of new algorithmic ideas in MATLAB followed by the development of low-entry-cost options for others once the ideas are embedded in the Duali code.

The distinction made above of the three types of software is useful for understanding the functionality of various types of software. However, these distinctions are being broken down. The algorithmic systems like MATLAB and GAUSS now make available add-ons that embody different algorithms for solving models. Thus one can use an algorithmic language as a modeling language by purchasing the desired add-on solvers and focusing most of one's attention on model development. Also, MATLAB now provides an add-on that enables one to use it like a menu-driven system. This system, called Simulink, has been used by Ric Herbert and Rod Bell (1997) for models of the Australian economy. In the Simulink system the ability to manipulate symbols that

represents various parts of the model turns complex programming tasks into simple mouse moves and greatly increases programmer efficiency.

What are the paths ahead in software development for stochastic control models? One of these paths is to make the software available over the Internet in the form of applications service providers (ASP). Ray Fair (2002a) has taken this approach with the FAIRMODEL macroeconomic software and has thus made that software available to a worldwide community of users while he maintains the software efficiently on a single machine under a desk in his office at Yale. One program that Ray makes available for downloading at this site is the Fair-Parke program that can be used, for example, for solving optimal control problems for nonlinear models with rational expectations.

At the second meeting of this society at the University of Texas in 1995 Hans and I gave a paper on programming languages in economics (Amman and Kendrick (1999d)), which elicited a wonderful and somewhat heated discussion among proponents of Fortran, Visual Basic and C. This debate was similar to an ongoing discussion that Hans and I had been having with one another about the relative merits of Visual Basic and C. A young Hans argued that untenured professors needed the fast development time of Visual Basic even if it cost them some speed in the execution of the program and that only tenured professors like David could afford the luxury of the slow development time and fast computational speeds of C. Well, Hans is now tenured and as well there are new languages like C# which have the rapid prototype-development speed of Visual Basic and Web-based application frameworks similar to that used by Ray Fair. Moreover these languages can be linked to C applications for fast execution.

So maybe the path ahead for stochastic control software for economists will offer the ease of use and low entry-cost of menu-driven systems, the wide availability and easy access of Web-based applications and the blinding speed needed to solve dual control models with non-convex criterion functions over many Monte Carlo runs.

12. Conclusions

So what have we learned about stochastic control of economic models since the Princeton meeting in 1972 and what are the paths ahead in this research now in 2002?

We have learned ...

that medium sized deterministic linear-quadratic and general nonlinear models with tens to hundreds of state variables can be solved on desktop computers and with enough graphical and analytical feedback for us to track through the complexities in such models

that handcrafted feedback rules can be developed for macroeconomic models which are simpler than the full optimal rules and have solutions which are close but not quite as good as the optimal rules

that optimal feedback methods yield solutions that are usually not as good as optimal feedback with parameter uncertainty methods so that in macroeconomic applications it is cautious and wise to take account of parameter variances and covariances when computing optimal policies

that methods which include measurement errors on state variables provide a way to analyze the tradeoff between feedback rules which rely heavily on recent state variables which have larger measurement errors and more distant-past state variables which have undergone more revision and have lower measurement errors

that it appears that in the face of increasing uncertainty (1) *min-max controllers* will have a comparative advantage on models like those for bridge construction when one wants to be more aggressive and strengthen supporting members and (2) *expected value controllers* will have a comparative advantage when one wants to be more cautious as in monetary policy choice

that models with rational expectations in the form of forward variables can be solved so efficiently, particularly when the model is linear, that the presence of these variables does not add much to the complexity of the solution method

that the learning of parameters and therefore the reduction in variance of these parameter estimates in small quadratic-linear dual control models occurs rapidly – perhaps more rapidly than is realistic

that quadratic-linear dual control models with large uncertainty in parameter estimates at initial time may well have non-convex criterion functions and therefore must be solved with global optimization methods

that modern computers are fast enough that we can even solve quadratic-linear dual control models that include not only time varying parameter estimates but also time varying true parameters

that by using models with time-varying parameters, forward variables and parameter updating procedures like the Kalman filter we can address well the concern that feedback rules will change when policy changes are announced

Against this background then what are some of the key questions in the research agenda for the path ahead?

1. Will the early trend that handcrafted feedback rules prove to be simpler to understand and develop and almost as good as optimal feedback rules hold as models with longer lag structures are analyzed?
2. Will the present result that optimal feedback rules in quadratic-linear models *with* parameter uncertainty is better than optimal feedback rules *without* parameter uncertainty in most models continue to hold as larger and more complex models are used?
3. Will measurement errors prove to be large enough that they should be considered when computing optimal policies – particularly in cases where some states have much larger measurement errors than others?
4. Will adaptive control methods prove to be so much better than passive control methods that they will be widely used in microeconomic, public finance and perhaps even in some macroeconomic settings?
5. Will algorithms and codes for solving nonconvex dual control models become so efficient that one can investigate the previous question even in cases where the degree of parameter uncertainty is so great as to result in nonconvex criterion functions?

6. We now know that we can develop more realistic models by including parameter uncertainty, time-varying parameters, measurement error and forward-variables and can use dual control methods to solve these models. However, each of these facets adds to the complexity of the calculations and the difficulty of sorting through and understanding the results. Therefore which of these attributes will prove to be most important in increasing the realism of our models while at the same time permitting us to obtain reliable solutions?

Though it is difficult to capture the above in a single concluding sentence it seems to me that the most important aspect of the paths ahead is for us to focus on the impact of various types of uncertainty on policy determination, i.e. not only additive shocks but also parameter uncertainty, measurement errors, uncertainty about the initial states and even about model specification. We have made substantial progress since the Princeton meeting, however, the paths ahead of us now look just as challenging and as full of potential for improving economic analysis as they did during that meeting thirty years ago.

References

- Abel, Andrew B. (1975), "A Comparison of Three Control Algorithms to the Monetarist-Fiscalist Debate," *Annals of Economic and Social Measurement*, **4**, 239-252.
- Aghion, P., P. Bolton, C. Harris and B. Jullien (1991), "Optimal Learning by Experimentation," *Review of Economic Studies*, **58**, 621-654.
- Amman, Hans M. and David A. Kendrick (1994), "Active Learning: Monte Carlo Results", *Journal of Economic Dynamics and Control*, **18**, 119-124.
- Amman, Hans M. and David A. Kendrick (1995), "Nonconvexities in Stochastic Control Models," *International Economic Review*, **36**, 455-475.
- Amman, Hans M. and David A. Kendrick (1997), "Erratum to 'Active Learning: Monte Carlo Results'", *Journal of Economic Dynamics and Control*, **21**, 1613-4.
- Amman, Hans M. and David A. Kendrick (1999a), "Linear Quadratic Optimization for Models with Rational Expectations", with Hans M. Amman, *Macroeconomic Dynamics*, **3**, p. 534-543.
- Amman, Hans M. and David A. Kendrick (1999b), "Should Macroeconomic Policy Makers Consider Parameter Covariances", *Computational Economics*, **14**, December, 263-267.

Amman, Hans M. and David A. Kendrick (1999c), "The Duali/Dualpc Software for Optimal Control Models: User's Guide", Center for Applied Research in Economics, The University of Texas, Austin, Texas.

Amman, Hans M. and David Kendrick (1999d), "Programming Languages in Economics", *Computational Economics*, **14**, pp. 151-181.

Amman, Hans M. and David A. Kendrick (2000), "Stochastic Policy Design in a Learning Environment with Rational Expectations", *Journal of Optimization Theory and Applications*, **105**, 509-520.

Amman, Hans M. and David A. Kendrick (2002), "Mitigation of the Lucas Critique with Stochastic Control Methods" forthcoming in the *Journal of Economic Dynamics and Control*.

Anderson, Gary and George Moore (1985), "A Linear Algebraic Procedure for Solving Linear Perfect Foresight Models", *Economics Letters*, **17**, 247-252.

Aoki, Masanao (1967), *Optimization of Stochastic Systems*, Academic Press, New York.

Athans, Michael and Gregory C. Chow (1972), "Introduction to Stochastic Control and Economic Systems", *Annals of Economic and Social Measurement*, **1**, 375-384.

Basar, Tamer and T. G. Olsder (1999), *Dynamic Noncooperative Game Theory*, SIAM Series in Classics in Applied Mathematics, Philadelphia. (Revised version of the 1995 Academic Press book with the same title.)

Beck, Gunter W. and Volker Wieland (2002), "Learning and Control in a Changing Economic Environment," *Journal of Economic Dynamics and Control*, **26**, 1359-1378.

Belsley, David (ed) (1994), *Computational Techniques for Econometrics and Economic Analysis*, Kluwer Academic Publishers, Dordrecht, The Netherlands.

Binder, Michael and M. Hashem Pesaran (2000), "Solution of Finite Horizon Multivariate Linear Rational Expectations Models with Sparse Linear Systems", *Journal of Economic Dynamics and Control*, **24**, 325-346.

Blanchard, O.J. and C.M. Kahn (1980), "The Solution of Linear Difference Models under Rational Expectations", *Econometrica*, **48**, 1305-1311.

Brainard, W. (1967) "Uncertainty and the Effectiveness of Policy", *American Economic Review*, **57**, 411-425.

Breslaw, Jon (1994), "GAUSSX, Econotron Software, Ver. 3.3", Econotron Software, 4164 Henri Julien, Montreal, P.Q., Canada H2W-2K3.

Brock, William A. and D. Dechert (1991), "Nonlinear Dynamical Systems: Instability and Chaos in Economics", in *Handbook of Mathematical Economics IV*, W. Hildenbrand and H. Sonnenschein (eds), North Holland, 2210-2235.

Brooke, Anthony, David A. Kendrick and Alexander Meeraus (1982), *GAMS: A User's Guide*, The Scientific Press, South San Francisco, currently available at the GAMS web cite <http://www.gams.com>.

Bryson, Arthur and Yu-Chi Ho (1969), *Applied Optimal Control*, Ginn and Company, Waltham, Mass.

Chiarella, Carl (1988), "The Cobweb Model: Its Instability and the Onset of Chaos", *Economic Modeling*, **5**, 377-384.

Chow, Gregory C. (1973), "Effects of Uncertainty on Optimal Control Policies", *International Economic Review*, **14**, 632-645.

Chow, Gregory C. (1975), *Analysis and Control of Dynamic Systems*, Wiley, New York.

Coenen, Guenter, Andrew Levin and Volker Wieland (2001), "Data Uncertainty and the Role of Money as an Information Variable for Monetary Policy", European Central Bank Working Paper No. 84, Frankfurt, Germany.

Craine, Roger (1979), "Optimal Monetary Policy with Uncertainty", *Journal of Economic Dynamics and Control*, **1**, 59-83.

Craine, Roger, Arthur Haverner, and Peter Tinsley (1976), "Optimal Macroeconomic Control Policies", *Annals of Economic and Social Measurement*, **5**, 191-204.

Currie, D. and P. Levine (1993a), "The Design of Feedback Rules in Linear Stochastic Rational Expectations Models," Ch. 4, pp. 95-121 in Currie and Levine (1993b).

Currie, D. and P. Levine (1993b), *Rules, Reputation and Macroeconomic Policy Coordination*, Cambridge University Press, Cambridge, United Kingdom.

Dechert, W. D. and Cars H. Hommes (2000), "Complex Nonlinear Dynamics and Computational Methods", *Journal of Economic Dynamics and Control*, **24**, 651-662.

Easley, D. and N. M. Kiefer (1988), "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, **56**, 1045-1064.

Fair, Ray C. (1978), "The Use of Optimal Control Techniques to Measure Economic Performance", *International Economic Review*, Vol. **19**, 289-309.

Fair, Ray C. (2002a), FAIRMODEL, <http://fairmodel.econ.yale.edu>.

- Fair, Ray C. (2002b), "Optimal Control and Stochastic Simulation of Large Nonlinear Models with Rational Expectations," forthcoming in *Computational Economics*.
- Fair, Ray C. and John B. Taylor (1983), "Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models", *Econometrica*, **51**, 1169-1185.
- Farison, J. B., R. E. Graham, and R. C. Shelton (1967), "Identification and Control of Linear Discrete Systems", *IEEE Trans. Autom. Control*, Vol.**AC-12**, No. 4, 438-442.
- Fisher, P. G., S. Holly and A. J. Hughes Hallett (1986), "Efficient Solution Techniques for Dynamic Nonlinear Rational Expectations Models," *Journal of Economic Dynamics and Control*, **10**, 139-145.
- Giannoni, Marc (2002), "Does Model Uncertainty Justify Caution?: Robust Optimal Monetary Policy in a Forward Looking Model," *Macroeconomics Dynamics*, **6**, 111-144.
- Hall, Robert E. and John B. Taylor (1993), *Macroeconomics*, 4th edition, W. W. Norton & Company, New York.
- Hansen, Lars Peter and Thomas J. Sargent (2001), *Elements of Robust Control and Filtering for Macroeconomics*, draft downloaded from Sargents web site at www.stanford.edu/~sargent, version of 23 March.
- Hansen, Lars Peter, Thomas J. Sargent and Thomas D. Tallarini, Jr. (1999), "Robust Permanent Income and Pricing", *Review of Economic Studies*, **66**, 873-907.
- Healey, A. J. and S. Summers (1974), "A Suboptimal Method for Feedback Control of the St. Louis Econometric Model", *Trans. ASME, J. Dynam. Syst., Meas. Control*, **96**, pp.446-454.
- Henderson, D. and S. J. Turnovsky (1972), "Optimal Macroeconomic Policy Adjustment under Conditions of Risk", *Journal of Economic Theory*, **4**, 58-72.
- Henderson, Dale W. and Warwick J. McKibbin (1993), "A Comparison of Some Basic Monetary Policy Regimes for Open Economies," International Finance Discussion Papers 458, Board of Governors of the Federal Reserve System, Washington, D.C.
- Herbert, Ric D. (1998), *Observers and Macroeconomic Systems*, Kluwer Academic Publishers, Boston/Dordrecht/London.
- Herbert, Ric D. and R. D. Bell (1997), "Visualisation in the Simulation and Control of Economic Models", *Computational Economics*, **10**, 107-118.
- Holt, Charles C. (1962), "Linear Decision Rules for Economic Stabilization and Growth", *Quarterly Journal of Economics*, **76**, 20-45.

Holt, Charles, Franco Modigliani, John F. Muth and Herbert A. Simon (1960), *Planning Production, Inventories, and Work Force*, Prentice-Hall, Inc., Englewood Cliffs, New Jersey.

Hughes-Hallett, Andrew and Peter McAdam (eds) (1999), *Analyses in Macroeconomic Modeling*, Kluwer Academic Publishers, Boston/Dordrecht/London.

Judd, Kenneth L. (1998), *Numerical Methods in Economics*, M.I.T. Press, Cambridge, Mass.

Juillard, Michel (1996), "DYNARE: A Program for the Resolution and Simulation of Dynamic Models with Forward Variables through the Use of a Relaxation Algorithm", CEPREMAP Working Paper No. 9602, Paris.

Kalchbrenner, J. and Peter Tinsley (1976), "On the Use of Feedback Control in the Design of Aggregate Monetary Policy," *American Economic Review*, 66, 349-355.

Kalchbrenner, J. and Peter Tinsley (1977), "On Filtering Auxiliary Information in Short-Run Monetary Policy," *Journal of Monetary Economics*, Carnegie-Rochester Supplement (ed. K. Brunner and A. Meltzer), 39-84.

Karakitsos, Elias and Berc Rustem (1984), "Optimally Derived Fixed Rules and Indicators", *Journal of Economic Dynamics and Control*, **8**, 33-64.

Karakitsos Elias and Berc Rustem (1985), "Optimal Fixed Rules and Simple Feedback Laws for Nonlinear Econometric Models", *Automatica*, **21**, 169-180.

Kendrick, David A. and Lance Taylor (1970), "Numerical Methods and Nonlinear Planning Models," *Econometrica*, **38**, pp. 453-467.

Kendrick, David A., and Joseph Majors (1974), "Stochastic Control with Uncertain Macroeconomic Parameters", *Automatica*, Vol. **10**, 587-594.

Kendrick, David A. (1978), "Non-convexities from Probing an Adaptive Control Problem", *Journal of Economic Letters*, **1**, 347-351.

Kendrick, David A. (1981), *Stochastic Control for Economic Models*, McGraw-Hill Book Company, New York.

Kendrick, David A. (1982), "Caution and Probing in a Macroeconomic Model", *Journal of Economic Dynamics and Control*, **4**, 149-170.

Kendrick, David A. (2002), "A Classification System for Economic Stochastic Control Models", Working Paper, Department of Economics, The Univ. of Texas, Austin, TX 78712.

Kiefer, N. and Y. Nyarko (1989), "Optimal Control of an Unknown Linear Process with Learning," *International Economic Review*, **30**, 571-586.

Kozicki, Sharon and Peter A. Tinsley (2002), "Dynamic Specification in Optimizing Trend-Deviation Macro Models," *Journal of Economic Dynamics and Control*, **26**, 1585-1612.

Lee, Myong Hwal (1998), "Analysis of Optimal Macroeconomic Policy Design", Ph.D. Dissertation, Department of Economics, The University of Texas, Austin, Texas 78712.

Leeson, Robert (2000), *A. W. H. Phillips: Collected Works in Contemporary Perspective*, Cambridge University Press, Cambridge, United Kingdom.

Levin, Andrew, Volker Wieland, and John C. Williams, (1999), "The Performance of Forward-Looking Monetary Policy Rules under Model Uncertainty", in Taylor (1999).

Lin, Kuan-Pin (2001), *Computational Econometrics: GAUSS Programming for Econometricians and Financial Analysts*, ETEXT Publishing.

Livesey, David A. (1971), "Optimizing Short-Term Economic Policy", *Economic Journal*, **81**, 525-546.

Lucas, Robert (1976), "Econometric Policy Evaluation: A critique", in K. Brunner and A. H. Meltzer (editors), *The Phillips Curve and the Labor Markets, 19-46*, Supplemental series to the *Journal of Monetary Economics*.

MacRae, Elizabeth Chase (1972), "Linear Decision with Experimentation", *Annals of Economic and Social Measurement*, **1**, 437-447.

MacRae, Elizabeth Chase (1975), "An Adaptive Learning Role for Multiperiod Decision Problems", *Econometrica*, **43**, 893-906.

Math Works (1999), *Using MATLAB Version 5*, The Math Works Inc., 24 Prime Park Way, Natick, MA 01760,

Mercado, P. Ruben, David A. Kendrick and Hans M. Amman (1998), "Teaching Macroeconomics with GAMS", *Computational Economics*, **12**, 125-149.

Mercado, P. Ruben and David A. Kendrick (1999), "Computational Methods for Macro Policy Analysis: Hall and Taylor's Model in Dual", Chapter 8 in Hughes-Hallett and McAdam (1999), 179-206.

Mercado, P. Ruben and David A. Kendrick (2000), "Caution in Macroeconomic Policy: Uncertainty and the Relative Intensity of Policy", *Economics Letters*, **68**, 37-41.

Mercado, P. Ruben (2001), "The Timing of Uncertainty and the Intensity of Policy", Paper presented at the Seventh International Conference of the Society of Computational Economics, Yale University, New Haven, June 28-29.

Mizrach, B. (1991), "Non-convexities in a Stochastic Control Problem with Learning", *Journal of Economic Dynamics and Control*, **15**, 515-538.

Norman, A. L. (1976), "First Order Dual Control," *Annals of Economic and Social Measurement*, **5**, 311-21.

Norman, A. (1981), "On the Control of Structural Models," and "A Reply", *Journal of Econometrics*, **15**, 13-24, 29.

Norman, A. (1994), "On the Complexity of Linear Quadratic Control," *European Journal of Operations Research*, **73**, 319-330.

Norman, A. L. and W. Jung (1977), "Linear Quadratic Theory for Models with Long Lags," *Econometrica*, **45**, 905-17.

Norman, A., M. Norman and C. Palash (1979), "Multiple Relative Maxima in Optimal Macroeconomic Policy: An Illustration", *Southern Economic Journal*, **46**, 274-279.

Onatski, Alexei and James H. Stock (2000), "Robust Monetary Policy Under Model Uncertainty in a Small Model of the U.S. Economy", Working Paper 7490, National Bureau of Economic Research, Cambridge, MA.

Oudiz, G. and J. Sachs (1985), "International Policy Coordination in Dynamic Macroeconomic Models", in Willem H. Buiter and R. C. Marston, *International Economic Policy Coordination*, Cambridge University Press, Cambridge, England.

Phillips, A. W. (1954), "Stabilization Policy in a Closed Economy", *Economic Journal*, **64**, 290-323.

Pindyck, Robert S. (1972), "An Application of the Linear Quadratic Tracking Problem to Economic Stabilization Policy", *IEEE Trans. Autom. Control*, **AC-17**, 287-300.

Pindyck, Robert S. (1973a), *Optimal Planning for Economic Stabilization*, North-Holland, Amsterdam.

Pindyck, Robert S. (1973b), "Optimal Policies for Economic Stabilization", *Econometrica*, **41**, 529-560.

Prescott, Edward C. (1972), "The Multi-Period Control Problem under Uncertainty", *Econometrica*, **40**, 1043-1058.

- Rust, John (1996), "Using Randomization to Break the Curse of Dimensionality" *Econometrica*.
- Rustem, Berc (1992), "A Constrained Min-Max Algorithm for Rival Models of the Same Economic System", *Math Programming*, **53**, 279-295.
- Rustem, Berc (1994), "Robust Min-Max Decisions with Rival Models", in Belsley (1994), pp. 109-136.
- Rustem, Berc (1998), *Algorithms for Nonlinear Programming and Multiple Objective Decisions*, John Wiley & Sons, New York.
- Rustem, Berc, Volker Wieland and Stanislav Zakovic, (2001), "A Continuous Min-Max Problem and Its Application to Inflation Targeting", Ch. 11, pp. 201-219 in Zaccour (2002).
- Rustem, Berc and M. Howe (2002), *Algorithms for Worst-Case Design with Applications to Risk Management*, Princeton University Press, Princeton, New Jersey.
- Sarris, A. H. (1973), "A Bayesian Approach to the Estimation of Time-Varying Regression Coefficients", *Annals of Economic and Social Measurement*, **2**, 501-523.
- Shupp, Franklin R. (1972), "Uncertainty and Stabilization Policies for a Nonlinear Macroeconomic Model", *Quarterly Journal of Economics*, **80**, 94-110.
- Shupp, Franklin R. (1976), "Uncertainty and Optimal Stabilization Policies", *J. Public Finances*, **6**, 243-253.
- Simon, H. A. (1956), "Dynamic Programming under Uncertainty with a Quadratic Criterion Function", *Econometrica*, **24**, 74-81.
- Sims, Christopher A. (1996), "Solving Linear Rational Expectations Models", Research Paper, Department of Economics, Yale University.
- Swamy, P.A.V.B. and Peter A. Tinsley (1980), "Linear Prediction and Estimation Methods for Regression Models with Stationary Coefficients," *Journal of Econometrics*, 108-42.
- Taylor, John B. (1974), "Asymptotic Properties of Multiperiod Control Rules in the Linear Regression Model", *International Economic Review*, Vol.**15**, 472-482.
- Taylor, John B. (1993a), *Macroeconomic Policy in a World Economy*, W. W. Norton & Company, New York.
- Taylor, John B. (1993b), "Discretion Versus Policy Rules in Practice", *Carnegie-Rochester Conference Series on Public Policy*, **39**, 195-214.

Taylor, John B. (ed) (1999), *Monetary Policy Rules*, Univ. of Chicago Press, Chicago.

Tetlow, Robert (2000), "The Fed Is Not as Ignorant as You Think", presented at the Society of Computational Economics Conference in Barcelona, June 2000.

Tetlow, Robert and Peter von zur Muehlen (2001a), "Simplicity versus Optimality: The Choice of Monetary Policy Rules When Agents Must Learn," *Journal of Economic Dynamics and Control*, **25**, 245-279.

Tetlow, Robert and Peter von zur Muehlen (2001b), "Robust Monetary Policy with Misspecified Models: Does Model Uncertainty Always Call for Attenuated Policy?," *Journal of Economic Dynamics and Control*, **25**, 911-949.

Tetlow, Robert and Peter von zur Muehlen (2001c), "Avoiding Nash Inflation: Bayesian and Robust Responses to Model Uncertainty", working paper, Board of Governors of the Federal Reserve System, Washington, D.C.

Theil, H. (1957), "A Note on Certainty Equivalence in Dynamic Planning", *Econometrica*, **25**, pp.346-349.

Tinsley, P., R. Craine, and A. Havenner (1974), "On NEREF Solutions of Macroeconomic Tracking Problems", 3d NBER Stochastic Control Conf., Washington..

Tinsley, Peter, Peter von zur Muehlen and G. Fries (1982), "The Short-Run Volatility of Money Stock Targeting," *Journal of Monetary Economics*, 215-237.

Tse, Edison and Yaakov Bar-Shalom (1973), "An Actively Adaptive Control of Linear Systems with Random Parameters", *IEEE Transactions on Automatic Control*, **AC-18**, 109-117.

Tucci, Marco (1989), *Time Varying Parameters in Adaptive Control*, Ph.D. Dissertation, Dept. of Economics, University of Texas, Austin, Texas, 78712.

Tucci, Marco (1997), "Adaptive Control in the Presence of Time-Varying Parameters", *Journal of Economic Dynamics and Control*, **22**, 39-47.

Tucci, Marco (1998), "The Nonconvexities Problem in Adaptive Control Models: A Simple Computational Solution", *Computational Economics*, **12**, 203-222.

Tucci, Marco (2002), "A Note on Global Optimization in Adaptive Control, Econometrics and Macroeconomics", *Journal of Economic Dynamics and Control*, **26**, 1739-1764

Turnovsky, Stephen J. (1973), "Optimal Stabilization Policies for Deterministic and Stochastic Linear Systems," *Review of Economic Studies*, **40**, 79-96.

Turnovsky, Stephen J. (1975), "Optimal Choice of Monetary Instruments in a Linear Economic Model with Stochastic Coefficients", *J. Money Credit Banking*, 7, 51-80.

Turnovsky, Stephen J. (1977), "Optimal Control of Linear Systems with Stochastic Coefficients and Additive Disturbances", Ch. 11, in Pitchford and Turnovsky (1977).

Turnovsky, Stephen J. (2000), *Methods of Macroeconomic Dynamics*, M.I.T. Press, Cambridge, Mass.

Tustin, Arnold (1953), *The Mechanism of Economic Systems*, Heinemann.

Vallee, Th., Christophe Deissenberg and Tamer Basar (1999), "Optimal Open Loop Cheating in Dynamic Reversed Linear-Quadratic Stackelberg Games," *Annals of Operations Research*, **88**, 247-266.

von zur Muehlen, Peter (1982), "Activist vs. Non-activist Monetary Policy: Optimal Rules Under Extreme Uncertainty", Mimeo, Board of Governors of the Federal Reserve System, Washington, D.C.

Wieland, Volker (2000a) "Learning by Doing and the Value of Optimal Experimentation", *Journal of Economics Dynamics and Control*, **24**, 501-543.

Wieland, Volker (2000b), "Monetary Policy, Parameter Uncertainty and Optimal Learning", *Journal of Monetary Economics*, **46**, 199-228.

Wieland, Volker (2001), "Learning, Stabilization and Credibility: Optimal Monetary Policy in a Changing Economy", paper presented at the 7th International Conference of the Society of Computational Economics, Yale University, New Haven, CT.

Williams, John C. (1999), "Simple Rules for Monetary Policy," Finance and Economics Discussion Series 1999-12, Board of Governors of the Federal Reserve System, Washington, D.C.

Woodford, Michael (1999), "Optimal Monetary Policy Inertia", Working Paper 7261, National Bureau of Economic Research, Cambridge, MA.

Zaccour, Georges (2002), *Decision and Control in Management Science, Essays in Honor of Alan Haurie*, Kluwer Academic Publishers, Boston/Dordrecht/London.

Zadrozny, Peter and Baoline Chen (1999), "Perturbation Solutions of Nonlinear Rational Expectations Models," presented at the Fifth International Conference of the Society for Computational Economics, Boston College, June 1999.