



HAL
open science

The parameter set in an adaptive control Monte Carlo experiment: Some considerations

Marco P. Tucci, David A. Kendrick, Hans M. Amman

► To cite this version:

Marco P. Tucci, David A. Kendrick, Hans M. Amman. The parameter set in an adaptive control Monte Carlo experiment: Some considerations. *Journal of Economic Dynamics and Control*, 2010, 34 (9), pp.1531. 10.1016/j.jedc.2010.06.014 . hal-00732676

HAL Id: hal-00732676

<https://hal.science/hal-00732676>

Submitted on 16 Sep 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Author's Accepted Manuscript

The parameter set in an adaptive control Monte Carlo experiment: Some considerations

Marco P. Tucci, David A. Kendrick, Hans M. Amman

PII: S0165-1889(10)00144-2
DOI: doi:10.1016/j.jedc.2010.06.014
Reference: DYNCON 2444

To appear in: *Journal of Economic Dynamics & Control*



www.elsevier.com/locate/jedc

Cite this article as: Marco P. Tucci, David A. Kendrick and Hans M. Amman, The parameter set in an adaptive control Monte Carlo experiment: Some considerations, *Journal of Economic Dynamics & Control*, doi:10.1016/j.jedc.2010.06.014

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting galley proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

The Parameter Set in an Adaptive Control Monte Carlo Experiment: Some Considerations¹

Marco P. Tucci , David A. Kendrick , Hans M. Amman

*Dipartimento di Economia Politica, Università di Siena, Piazza S. Francesco, 7, 53100
Siena, Italy*

Department of Economics, University of Texas, Austin, Texas 78712, USA.

*Utrecht School of Economics, Utrecht University, Heidelberglaan 8, 3584 CS Utrecht, the
Netherlands.*

Abstract

Comparisons of various methods for solving stochastic control economic models can be done with Monte Carlo methods. These methods have been applied to simple one-state, one-control quadratic-linear tracking models; however, large outliers may occur in a substantial number of the Monte Carlo runs when certain parameter sets are used in these models. Building on the work of Mizrach (1991) and Amman and Kendrick (1994, 1995), this paper tracks the source of these outliers to two sources: (1) the use of a zero for the penalty weights on the control variables and (2) the generation of near-zero initial estimate of the control parameter in the systems equations by the Monte Carlo routine. This result leads to an understanding of why both the unsophisticated Optimal Feedback (Certainty Equivalence) and the sophisticated Dual methods do poorly in some Monte Carlo comparisons relative to the moderately sophisticated Expected Optimal Feedback method.

Key words: Active learning, dual control, optimal experimentation, stochastic optimization, time-varying parameters, numerical experiments.

JEL Classification: C63, E61.

Email addresses: `tucci@unisi.it` (Marco P. Tucci),
`kendrick@eco.utexas.edu` (David A. Kendrick), `amman@uu.nl` (Hans M. Amman).

¹ We would like to thank an anonymous referee for valuable comments.
Corresponding author Marco P. Tucci, document name *main.tex*.

1 Introduction

One of the outstanding problems in the application of stochastic control methods to economic models is the relative performance of different ways of treating the uncertainties in these models and thus of different methods for solving the models. For example, there has been considerable comparison of Optimal Feedback (*OF*), Expected Optimal Feedback (*EOF*) and Dual Control (*DC*) methods using Monte Carlo runs.² The first of these methods is a *certainty equivalence* approach that ignores all the uncertainties in the model except the additive noise terms in the systems equations. The second method considers additive noise terms as well as parameter uncertainty in the system equations but uses passive learning. The third method considers both additive noise and parameter uncertainty and adds consideration of active learning methods that include perturbation early in time to gain improved estimates of the model parameters later in time.

The comparison of these methods has mostly been done applying Monte Carlo methods to simple one-state, one-control quadratic-linear tracking models such as the MacRae (1975) model and the Beck and Wieland (2002) model.³ However, in some cases these comparisons have encountered a significant number of outliers, viz Amman, Kendrick and Tucci (2007). The source of these outliers was a puzzle – particularly so because they occurred prominently in the least sophisticated (Optimal Feedback) and the most sophisticated (Dual Control) methods but not in the intermediate sophistication method (Expected Optimal Feedback).

In this paper we track the source of this puzzle to the choice of parameters (i.e. the parameter set) in these models. In particular, we find that the use of a zero for the penalty weight on the control variables in the criterion function (the λ parameter) can cause problems. This is true when this weight is used in conjunction with a parameter value near zero for the coefficient in the system equation that is multiplied by the control variable (the b coefficient). If the variance of the b parameter is sufficiently large, then the Monte Carlo routine will generate values of this estimate around zero in a meaningful number of runs. The combination of a zero value for λ and a near zero value for b results in division by a number near zero and thus causes large values for some components of the calculations and therefore outliers. However, this occurs in the Optimal Feedback and the Dual Control methods but not in the Expected Optimal Feedback methods as is discussed below.

The MacRae (1975) and the Beck and Wieland (2002) models are identical in struc-

² For a classification of various methods of solving stochastic control models see Kendrick and Amman (2006).

³ This work is a part of the Methods Comparison Project that seeks to compare various ways of solving economic stochastic control models. For examples of the various approaches under consideration see Cosimano (2008), Cosimano and Gapen (2006), Amman, Kendrick and Tucci (2007) and Beck and Wieland (2002).

ture; however they use different sets of parameter values. In particular the Beck and Wieland model uses a zero for the weight, λ , on the control variables in the criterion function. Consequently it is this model where the outlier problem is pronounced, particularly when the variance of the b parameter is relatively large.

In tracking down the source of the outliers we have made considerable use of the research results on nonconvexities in the cost-to-go function of adaptive control problems. This work began when the Tse and Bar-Shalom (1973) algorithm for solving adaptive control problems as used by Kendrick (1978) and, Norman, Norman and Palash (1979) found nonconvexities in the cost-to-go function of a ten-period three-parameter version of the classic MacRae (1972) problem. However, the computer codes were sufficiently complex that it was difficult to be sure that the phenomena were not caused by programming errors. Therefore, at the time it was uncertain whether nonconvexities were

- (1) caused by a programming error;
- (2) caused by the particular choice of parameters;
- (3) fundamental to this class of problems.

Over a decade later Mizrach (1991) and Amman and Kendrick (1994, 1995) returned to this problem showing, by analytical methods, that nonconvexities are fundamental in this type of models.

In the Tse and Bar-Shalom (1973) approach for solving active learning stochastic control problems, the total cost-to-go is approximated by a function which can be decomposed into three terms: deterministic, cautionary and probing. Nonconvexities are caused by the probing and/or cautionary components of the cost-to-go. When only the parameter associated with the control variable is unknown, the nonconvexities are determined by the changes in the path of future state variables induced by modifying the value of the control. If more than one parameter is unknown the covariance between the unknown parameters, e.g. the covariance between the unknown parameter associated with the control and that associated with the state, and the relative magnitude of their variances is another possible source of nonconvexities (Mizrach (1991)).⁴

The situation with only one unknown parameter is further investigated in Amman

⁴ According to Mizrach (1991):

increases in the control variables ultimately do reduce the parametric uncertainty to zero in the limit ... there will often be at least one region in which increases in the control value raise parametric uncertainty before it begins to decline.

This phenomenon can generate nonconvexities, see Mizrach (1991, pages 516, 526-534). Furthermore when the planning time horizon is short initial conditions are important. For instance in the 2 periods case, different values of x_0 affect the variance term in an interesting way, Mizrach (1991, page 534).

and Kendrick (1994, 1995). They find that nonconvexities depend upon the magnitude of the initial variance of the unknown parameter, the level of the parameter itself and the variance of the additive noise in the system equation. Moreover, they show that particular combinations of the parameters and of the penalty weights in the objective function can turn the usually convex cautionary term of the cost-to-go into a concave term which, under certain circumstances, may generate a concave cost-to-go. These results are important because they clearly indicate that nonconvexities are fundamental in the mathematics of the problem (Amman and Kendrick (1995, page 456)) and can arise even in very simple adaptive control problems.

These conclusions have substantial computational consequences. When nonconvexities are absent, the stochastic control problem can be solved with gradient methods. However the *gradient solution* may be a local optimum, rather than a global optimum, when nonconvexities are present. For this reason a globally optimizing algorithm should be generally preferred except in the one control case where a grid search may be equally effective. Indeed as shown in Tucci (1998, 2004), nonconvexities may be much more common and subtle than the theoretical results suggest. For example, Amman and Kendrick (1995, page 465) find that when the MacRae (1972) parameter set is used, the nonconvexity appears when the variance of the estimate of the unknown parameter is set to 2. On the other hand Tucci (1998, 2004) reports nonconvexities in 28% of the cases of a Monte Carlo experiment with the same parameter set and a variance equal to .5 for the unknown parameter. Therefore the need for a globally optimizing algorithm may be more stringent than generally believed.

One side effect of the line of research discussed above on the origin of nonconvexities is that the analytical results contained in these papers allow one to fully characterize these three components of the cost-to-go for the simplest one-state, one-control, one unknown parameter, quadratic linear adaptive control problem with a time horizon of two periods. It is therefore possible to compare the *average* or *representative* cost-to-go associated with different parameter sets and to study the impact of a certain parameter set on the individual runs of a Monte Carlo experiment.⁵ The former may help to understand the basic characteristics of a certain parameter set. The latter may be useful to reconcile the theoretical results in Mizrach (1991) and Amman and Kendrick (1995) with the computational findings in Tucci (1998, 2004) and to shed some light on the *outlier problem* mentioned in Amman et al. (2007).

In the following a simple, but fairly general, one-state, one-control, one unknown parameter, quadratic linear adaptive control problem is presented first (Section 2). Then the probing, cautionary and deterministic components of its dynamic pro-

⁵ By parameter set is here intended both the value of the parameters and their covariances and the values used for the penalty matrices, desired paths for the states and controls and the initial state.

gramming cost-to-go are characterized (Sections 3 through 5). This is done for the case in which the planning time horizon is 2, following closely Amman and Kendrick (1994, 1995). In Section 6 it is shown that this model encompasses as special cases both the MacRae (1972) and Beck and Wieland (2002) models and parameter sets and the characteristics of the *representative* cost-to-go associated with these two problems are discussed. At this point the impact of a certain parameter set on the individual runs of a Monte Carlo experiment is investigated (Section 7). Its impact on the *OF* and *EOF* solutions is discussed in Section 8. Conclusions are summarized in Section 9.

2 The Problem

Consider a simple control problem with one state, one control and a time horizon of N periods in which the policy maker wants to find u_0, u_1, \dots, u_{N-1} to minimize

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

where E is the expectation operator, subject to

$$x_{k+1} = \alpha x_k + \beta_k u_k + \gamma + \varepsilon_{k+1} \quad \text{for } k = 0, 1, \dots, N-1 \quad (2)$$

with x_k and u_k the state and control variables, respectively. Also α , β_k and γ are the parameters of the system equation and ε_{k+1} is an error term identically and independently distributed (i.i.d.) normal with mean zero and variance q . Finally, the initial state x_0 and the penalty weights w 's and λ 's are given constants. The parameter associated with the control is assumed time-varying with the following law of motion⁶

$$\beta_{k+1} = d\beta_k + (1-d)\beta + \eta_{k+1} \quad \text{for } k = 0, 1, \dots, N-1 \quad (3)$$

where d is the transition parameter, β is the unconditional mean of the stochastic parameter and η_{k+1} is the additive error term i.i.d. normal with mean zero and variance σ_η^2 . Also, the states are measured without error.⁷

⁶ Tucci (1997, 1998) uses a similar model for the parameter associated with the control variable. Tucci (2004) and Tucci et al (2007) consider the case where all the parameters may be time-varying.

⁷ This is equivalent to setting $H = I$ and $R = 0$ in Kendrick (1981, Chapters 10-11) or

It is worth while to point out that equation (3) includes as special cases several relevant situations. For instance, when d and σ_η^2 are zero, β_k reduces to the usual time-invariant case. On the other hand, if $d = 0$ but $\sigma_\eta^2 \neq 0$ equation (3) describes a random parameter, i.e. a parameter varying randomly about the fixed mean β . When $\beta = 0$, equation (3) describes the usual vector-autoregressive process of order one with mean zero.⁸ Finally, parameters following a random walk are modeled setting $d = 1$. Alternatively, equation (3) may be used to model the lack of knowledge about the parameters. For instance when the true parameter associated with the control is constant but unknown, setting $d = 0$ and $\sigma_\eta^2 \neq 0$ in the parameter transition equation allows one to interpret β_k in equation (2) as the time-varying estimate, based on observations up to time k , of the unknown β used to determine optimal control. In this case σ_η^2 should be interpreted as the variance of the estimate based on all information available at time k .⁹

Following the Tse and Bar-Shalom (1973) method for solving active learning stochastic control problems, Kendrick (1981, 2002) and Tucci (1997, 2004) compute the approximate cost-to-go at different values of the control and then choose that value which yields the minimum approximate cost.¹⁰ This approximate cost-to-go can be decomposed into three terms and, for the present problem, written as

$$J_N = J_{D,N} + J_{C,N} + J_{P,N} \quad (4)$$

where J_N is the total cost-to-go with N periods remaining and $J_{D,N}$, $J_{C,N}$ and $J_{P,N}$ are the deterministic, cautionary and probing components, respectively. In equation (4) the deterministic component includes only terms which are not stochastic. The cautionary term includes uncertainty only in the next time period and the probing term contains uncertainty in all future time periods. Thus the probing term includes the motivation to perturb the controls in the present time period in order to reduce future uncertainty about parameter values.¹¹

3 The Probing Component

Amman and Kendrick (1994, 1995) consider the case where the planning horizon is $N = 2$ and all the parameters in equation (2) are constant, but β is unknown.

⁷ Tucci (2004, Chapters 2-5).

⁸ See, e.g., Harvey (1981, Chapter 2).

⁹ Sometimes the more cumbersome notation $\sigma_{b,k}^2$, with $\sigma_{b,k}^2 \equiv \sigma_\eta^2$, may be preferred to stress the fact that it is the variance of the estimate of the unknown parameter β based on the information available at time k , when this is the appropriate interpretation.

¹⁰ See Kendrick (1981, 2002, Chapters 9-10) or Tucci (2004, Chapter 2) for details.

¹¹ See Kendrick (1981, pages 97-98)) for an introduction to this decomposition.

Therefore the optimal control is determined using the estimate of β based on all the information available at time 0, i.e. $\beta_0 \equiv b$, with a variance equal to $\sigma_\eta^2 \equiv \sigma_b^2$. They show that the probing component takes the form ¹²

$$J_{P,2} = \frac{1}{2} \left[\frac{g(u_0)}{h(u_0)} \right] \quad (5)$$

where $g(u_0)$ is a quadratic function reflecting the effect of the control u_0 on future states x_1 and x_2 . ¹³ Also, $h(u_0)$ is a quadratic function whose inverse defines the relationship between the control and the updated variance of the unknown parameter b . ¹⁴ Consequently the probing component is the ratio of two quadratic functions that can be written as ¹⁵

$$h(u_0) = \left(\frac{1}{\sigma_b^2 q} \right) (\sigma_b^2 u_0^2 + q) \quad (6)$$

and

$$g(u_0) = \left(\frac{w_2^2}{\lambda_1 + b^2 w_2} \right) (b u_{01} + x_{02} - \tilde{x}_2)^2 \quad (7)$$

with u_{01} and x_{02} the nominal, or CE, values of u_1 and x_2 defined as ¹⁶

$$u_{01} = \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) \left[\alpha b^2 w_2 u_0 + \alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 (\gamma - \tilde{x}_2) - \lambda_1 \tilde{u}_1 \right] \quad (8)$$

$$\begin{aligned} x_{02} = & b \left(\alpha - \frac{\alpha b^2 w_2}{\lambda_1 + b^2 w_2} \right) u_0 + \alpha^2 x_0 + \alpha \gamma + \gamma \\ & + b \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) \left[\alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 (\gamma - \tilde{x}_2) - \lambda_1 \tilde{u}_1 \right] \end{aligned} \quad (9)$$

¹² See Amman and Kendrick (1994) for a detailed derivation.

¹³ The term $g(u_0)$ is the same as the *script R* term in Mizrach (1991) and Amman and Kendrick (1995) and the Θ term in Tucci (2004, Chapters 2-5).

¹⁴ The updated variance of the parameter can be obtained using the Kalman filter and is independent of the actual observations. See, e.g., Tucci (2004, Chapter 2) for a brief discussion or Harvey (1981, Chapter 4) for a more technical presentation.

¹⁵ See Amman and Kendrick (1994) for details.

¹⁶ Equations (8)-(9) correspond to (2.17) and (2.19) in Amman and Kendrick (1994), respectively.

Substituting (8)-(9) into (7) and simplifying yields

$$g(u_0) = \phi_1 (\phi_2 u_0 + \phi_3)^2 \quad (10)$$

with

$$\begin{aligned} \phi_1 &= \left(\frac{w_2^2}{\lambda_1 + b^2 w_2} \right) \\ \phi_2 &= \alpha b \left(1 - \frac{2b^2 w_2}{\lambda_1 + b^2 w_2} \right) \\ \phi_3 &= 2b \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) [\alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 (\gamma - \tilde{x}_2) - \lambda_1 \tilde{u}_1] \\ &\quad + \alpha^2 x_0 + \alpha \gamma + \gamma - \tilde{x}_2 \end{aligned} \quad (11)$$

So the probing component of the cost-to-go is the ratio of two quadratic functions in u_0 .¹⁷

From (5) with h and g defined as in (6) and (10)-(11), respectively, it follows that the probing function in this case is

$$J_{P,2} = \left(\frac{\sigma_b^2 q \phi_1}{2} \right) \frac{(\phi_2 u_0 + \phi_3)^2}{\sigma_b^2 u_0^2 + q} \quad (12)$$

From equation (11) ϕ_1 will be positive regardless of b and u_0 when the λ 's and w 's are positive. Under this condition the first term in parenthesis on the right-hand side of (12) will be positive provided that the variances σ_b^2 and q are nonzero. If these conditions are met, $J_{P,2}$ is non-negative for all values of u_0 .

Next Amman and Kendrick (1994, 1995) consider the roots of the probing term. The first derivative of equation (12) with respect to the control is

¹⁷ To investigate the behavior of this component of the probing cost, Mizrach (1991) studies its partial derivative with respect to the control at time 0, i.e. $[\partial g(u_0)/\partial u_0] = 2\phi_1 (\phi_2 u_0 + \phi_3) [\partial (\phi_2 u_0 + \phi_3) / \partial u_0]$. Given that ϕ_1 is positive regardless of b and u_0 when the λ 's and w 's are positive, the signs of the term in parenthesis and of the partial derivative appearing on the right-hand side are critical. If both of them are positive an increase in the control increases the volatility of the future states path and, depending upon how fast the variance of the uncertain parameter decreases when the control is increased, this can cause the appearance of nonconvexities in the probing component of the cost-to-go. This derivative is indeed positive when the set of parameters is $\alpha = .7$, $b = -3.5$, $c = 3.5$, $q = -3.5$, $\sigma_b^2 = 2.5$, $x_0 = 2$, $\lambda_1 = \lambda_2 = w_1 = w_2 = 1$, $\tilde{x}_0 = \tilde{x}_1 = \tilde{x}_2 = 0$ and $\tilde{u}_0 = \tilde{u}_1 = \tilde{u}_2 = 0$. In this case movements in the control variable produce changes in $g(u_0)$ large enough to overcome the drop in the updated variance of b caused by the same movement in the control and give rise to nonconvexities in the probing component Mizrach (1991).

$$\frac{\partial J_{P,2}}{\partial u_0} = \frac{\sigma_b^2 q \phi_1 \left[(\phi_2^2 u_0 + \phi_2 \phi_3) (\sigma_b^2 u_0^2 + q) - \sigma_b^2 u_0 (\phi_2 u_0 + \phi_3)^2 \right]}{(\sigma_b^2 u_0^2 + q)^2} \quad (13)$$

Since both the term outside the square brackets in the numerator of equation (13) and the denominator are strictly positive as long as the λ 's, w 's, σ_b^2 and q are positive, they focus on the condition under which the bracketed term would be zero. Expansion of the term in brackets yields

$$-\phi_2 \phi_3 \sigma_b^2 u_0^2 + (\phi_2^2 q - \sigma_b^2 \phi_3^2) u_0 + \phi_2 \phi_3 q \quad (14)$$

This expression is in general a quadratic function in u_0 with roots

$$r_1 = -\frac{\phi_3}{\phi_2} \quad \text{and} \quad r_2 = \frac{\phi_2 q}{\phi_3 \sigma_b^2} \quad (15)$$

So the first derivative of the probing function with respect to the control in period zero has two roots. In other words the probing function has two extreme points.

At this point Amman and Kendrick (1994, 1995) evaluate the function at these two extremes. Substitution of the first root r_1 into equation (12) yields $J_{P,2} = 0$. To evaluate the function at the second root they substitute r_2 into (12) to obtain

$$J_{P,2} = \underbrace{\left(\frac{\sigma_b^2 q \phi_1}{2} \right)}_A \underbrace{\left(\frac{\phi_2^2 q + \phi_3^2 \sigma_b^2}{\phi_3 \sigma_b^2} \right)^2}_B \underbrace{\left\{ \frac{\phi_2^2 q^2 + \phi_3^2 \sigma_b^2 q}{\phi_3^2 \sigma_b^2} \right\}^{-1}}_C \quad (16)$$

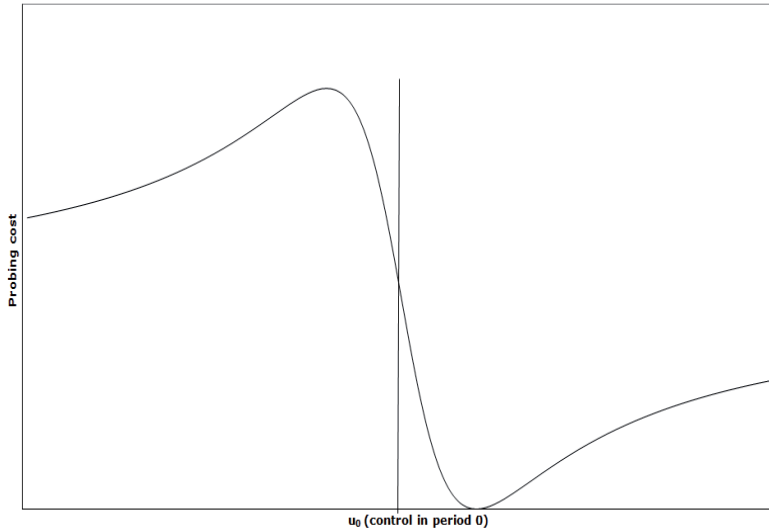
If the variance of the parameter, σ_b^2 , and the variance of the additive error term, q , are nonzero then the term A will be strictly positive given the fact from equation (11) that ϕ_1 is positive. Term B is a square so it will be strictly positive provided that either of the variances are non-zero. Term C will also be positive under the same restriction. Therefore the probing function $J_{P,2}$ in general will be strictly positive at the second root provided that σ_b^2 and q are non-zero.

To check the limit of the probing term as u_0 goes to either plus or minus infinity use (5) with h and g defined as in (6) and (10)-(11), respectively, to obtain

$$\lim_{u_0 \rightarrow \pm\infty} J_{P,2} = \frac{\sigma_b^2 q \phi_1 \phi_2^2}{2\sigma_b^2} = \frac{q \phi_1 \phi_2^2}{2} \quad (17)$$

The right hand side of (17) is positive so long as q is nonzero so the probing term approaches a positive limit when u_0 approaches either plus or minus infinity.

Figure 1.
The probing cost function $J_{P,2}$



Summarizing the probing function (5) with h and g defined as in (6) and (10)-(11), respectively, has the general form shown in Figure 1, page 10.¹⁸ The function is non-negative at all values of u_0 . The first derivative of the function has two roots which correspond to the maximum and minimum shown. At the minimum the value of the function is zero. The limits of the function as u_0 approaches plus or minus infinity are strictly positive and identical. The intuition associated with this function is that the probing component is decreased by actions which make u_0 large in either the positive or negative directions. Thus the perturbations to the system which come from making u_0 large in absolute value decrease the variance of the parameters in future time periods and therefore decrease the total cost-to-go.

4 The Cautionary Component

Following Amman and Kendrick (1994, equation 5.17), the cautionary component of the approximate cost-to-go can be written as

¹⁸ As pointed out by an anonymous referee, it is also possible that the curve cuts the axis with a positive slope.

$$J_{C,2} = \frac{\sigma_b^2 w_2}{2} (\alpha u_0 + u_{01})^2 + \frac{\sigma_b^2}{2} \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2 u_0 + b w_2 u_{01} + w_2 x_{02} - w_2 \tilde{x}_2)^2 + \frac{q}{2} \left[\alpha^2 w_2 + w_2 + w_1 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2)^2 \right] + \frac{\sigma_b^2 w_1}{2} u_0^2 \quad (18)$$

with u_{01} and x_{02} defined as in (8) and (9), respectively. As is shown in Appendix A, when these quantities are replaced by their definitions, equation (18) takes the form¹⁹

$$J_{C,2} = \delta_1 u_0^2 + \delta_2 u_0 + \delta_3 \quad (19)$$

with

$$\begin{aligned} \delta_1 &= \frac{\sigma_b^2}{2} \left[\nu_1^2 \left(w_2 - \frac{4b^2 w_2^2}{\lambda_1 + b^2 w_2} \right) + w_1 \right] \\ \delta_2 &= \sigma_b^2 w_2 \nu_1 \left\{ \nu_2 - \frac{2b w_2 (2b \nu_2 + \nu_3)}{\lambda_1 + b^2 w_2} \right\} \\ \delta_3 &= \frac{\sigma_b^2}{2} w_2 \left\{ \nu_2^2 - \frac{w_2 (2b \nu_2 + \nu_3)^2}{\lambda_1 + b^2 w_2} \right\} + \frac{q}{2} \left\{ \alpha^2 w_2 + w_2 + w_1 - \frac{(\alpha b w_2)^2}{\lambda_1 + b^2 w_2} \right\} \end{aligned} \quad (20)$$

where

$$\begin{aligned} \nu_1 &= \alpha \left(1 - \frac{b^2 w_2}{\lambda_1 + b^2 w_2} \right) \\ \nu_2 &= \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) [\alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 (\gamma - \tilde{x}_2) - \lambda_1 \tilde{u}_1] \\ \nu_3 &= \alpha^2 x_0 + \alpha \gamma + \gamma - \tilde{x}_2 \end{aligned} \quad (21)$$

This component of the cost-to-go is convex if $\delta_1 > 0$. When the penalty weights w_1 and w_2 are positive

$$1 - \frac{4b^2 w_2}{\lambda_1 + b^2 w_2} = \frac{\lambda_1 - 3b^2 w_2}{\lambda_1 + b^2 w_2} > 0 \quad (22)$$

or, alternatively, $\lambda_1 > 3b^2 w_2$ is a sufficient condition for δ_1 to be positive.²⁰ In summary, even if the cautionary term is concave, the cost-to-go will usually be convex because the deterministic term is always convex and is frequently larger

¹⁹ See also Amman and Kendrick (1995).

²⁰ This condition however is not necessary. It may happen that it is not satisfied and δ_1 is still positive. The bracket term in (20) may be positive thanks to the squared term in (21), or to w_2 , being sufficiently small to make the effect of w_1 overwhelming.

in magnitude than the cautionary term. Only when the parameter variance gets large the concave cautionary term can overcome the deterministic oneto produce a concave cost-to-go (see Amman and Kendrick, 1995, page 466).

5 The Deterministic Component and the Total Cost-to-Go

Finally Amman and Kendrick (1994, equation 6.8) write the deterministic component as

$$J_{D,2} = \frac{\lambda_0}{2} (u_0 - \tilde{u}_0)^2 + \frac{w_2}{2} (x_{02} - \tilde{x}_2)^2 + \frac{w_1}{2} (x_{01} - \tilde{x}_1)^2 + \frac{\lambda_1}{2} (u_{01} - \tilde{u}_1)^2 \quad (23)$$

with u_{01} and x_{02} defined as in (8) and (9), respectively and

$$x_{01} = \alpha x_0 + b u_0 + \gamma$$

Replacing these quantities in equation (23), squaring and rearranging the terms yields a convex quadratic function of the initial control value u_0 which can be written

$$J_{D,2} = \psi_1 u_0^2 + \psi_2 u_0 + \psi_3 \quad (24)$$

with

$$\begin{aligned} \psi_1 &= \frac{\lambda_0}{2} + \frac{w_2}{2} (b\nu_1)^2 + \frac{w_1 b^2}{2} + \frac{\lambda_1}{2} \left(-\frac{1}{\lambda_1 + b^2 w_2} \right)^2 (\alpha b^2 w_2)^2 \\ \psi_2 &= -\lambda_0 \tilde{u}_0 + w_2 b \nu_1 (b\nu_2 + \nu_3) + w_1 (\alpha x_0 + \gamma - \tilde{x}_1) b + \left(-\frac{\lambda_1}{\lambda_1 + b^2 w_2} \right) \alpha b^2 w_2 (\nu_2 - \tilde{u}_1) \\ \psi_3 &= \frac{\lambda_0}{2} \tilde{u}_0^2 + \frac{w_2}{2} (b\nu_2 + \nu_3)^2 + \frac{w_1}{2} (\alpha x_0 + \gamma - \tilde{x}_1)^2 + \frac{\lambda_1}{2} (\nu_2 - \tilde{u}_1)^2. \end{aligned} \quad (25)$$

Then, as in equation (4), the total cost-to-go for period $k = 0$, i.e. with two periods remaining, can be written as

$$J_2 = J_{D,2} + J_{C,2} + J_{P,2} \quad (26)$$

Substitution of (12), (19) and (24) into (26) yields²¹

$$J_2 = (\psi_1 + \delta_1) u_0^2 + (\psi_2 + \delta_2) u_0 + (\psi_3 + \delta_3) + \left(\frac{\sigma_b^2 q}{2} \right) \frac{\phi_1 (\phi_2 u_0 + \phi_3)^2}{(\sigma_b^2 u_0^2 + q)} \quad (27)$$

6 The Representative Cost-to-Go

The model analyzed in the previous sections encompasses as special cases both the MacRae (1972) and Beck and Wieland (2002) models and parameter sets.²² It is therefore possible to characterize the *representative* cost-to-go of these two popular problems. The parameter set used in MacRae (1972) is

$$\begin{aligned} \alpha &= .7 & \beta_0 \equiv b &= -.5 & \gamma &= 3.5 & \sigma_\varepsilon^2 = q &= .2 \\ w_k &= 1 & \forall k & & \lambda_k &= 1 & \forall k & & \sigma_\eta^2 \equiv \sigma_b^2 &= .5 & x_0 = \sigma_\alpha^2 = \sigma_\gamma^2 &= 0 \end{aligned} \quad (28)$$

and the desired path for the state and control is set equal to zero. Then, as already noticed in Amman and Kendrick (1994, 1995), the coefficients of the probing cost are

$$\begin{aligned} \phi_1 &= \left(\frac{1}{1+b^2} \right) = .8 \\ \phi_2 &= \alpha b \left(\frac{1-b^2}{1+b^2} \right) = -.21 \\ \phi_3 &= (\alpha\gamma + \gamma) \left(\frac{1-b^2}{1+b^2} \right) = 3.57 \end{aligned} \quad (29)$$

and this function reaches its minimum 0 at $r_1 = 17$ and its maximum $J_{P,2} = 2.553$ at $r_2 = -.023529$. The limits of the probing cost function as u_0 approaches plus

²¹ As noticed in Amman and Kendrick (1994, 1995), the derivative of this function is a complex function of u_0 and the function has several extreme points. However by taking the second derivative of the cost-to-go at $u_0 = 0$ it yields $(\partial^2 J_2 / \partial u_0^2)|_{u_0=0} = 2\psi_1 + 2\delta_1 + \sigma_b^2 \psi_1 \psi_2^2 - (\sigma_b^4 \phi_1 \phi_3^2 / q)$. Then they conclude that a higher variance, as well as a lower value of q , usually makes nonconvexities more likely.

²² The Excel spreadsheet that we have used for many of the calculations and the creation of the graphs in this and the following section is available from Marco Tucci (tucci@unisi.it).

or minus infinity are equal to $J_{P,2} = .0035$ and it has the general form shown in Figure 1, page 10.

The coefficients for the cautionary component of the approximate cost-to-go are

$$\begin{aligned}\nu_1 &= \alpha \left(\frac{1}{1+b^2} \right) = .56 \\ \nu_2 &= \left(-\frac{1}{1+b^2} \right) b(\alpha\gamma + \gamma) = 2.38 \\ \nu_3 &= \alpha\gamma + \gamma = 5.95\end{aligned}\tag{30}$$

therefore

$$\begin{aligned}\delta_1 &= \frac{\sigma_b^2}{2} \left[\nu_1^2 \left(\frac{1-3b^2}{1+b^2} \right) + 1 \right] = .2657 \\ \delta_2 &= \sigma_b^2 \nu_1 \left\{ \nu_2 - \frac{2b(2b\nu_2 + \nu_3)}{1+b^2} \right\} = 1.4661 \\ \delta_3 &= \frac{\sigma_b^2}{2} \left[\nu_2^2 - \frac{(2b\nu_2 + \nu_3)^2}{1+b^2} \right] + \frac{q}{2} \left[\alpha^2 + 2 - \frac{(\alpha b)^2}{1+b^2} \right] = -.8937\end{aligned}\tag{31}$$

and the cautionary component is convex. Finally the deterministic component

$$J_{D,2} = \psi_1 u_0^2 + \psi_2 u_0 + \psi_3\tag{32}$$

has coefficients

$$\begin{aligned}\psi_1 &= \frac{1}{2} \left[1 + (b\nu_1)^2 + b^2 + \left(-\frac{1}{1+b^2} \right)^2 (\alpha b^2)^2 \right] = .674 \\ \psi_2 &= b\nu_1 (b\nu_2 + \nu_3) + \gamma b + \left(-\frac{1}{1+b^2} \right) \alpha b^2 \nu_2 = -3.416 \\ \psi_3 &= \frac{1}{2} (b\nu_2 + \nu_3)^2 + \frac{1}{2} \gamma^2 + \frac{1}{2} \nu_2^2 = 20.286.\end{aligned}\tag{33}$$

Using these values the *representative* cost-to-go for the MacRae (1972) problem can be plotted as in Figure 2, page 15, with the *flat looking* portion around control 0 due to the fact that in that portion the cost-to-go decreases at a slower pace. The optimal cost in Figure 2, page 15, is 18.8605 at $u_0 = 1.32691$.²³

Next we shift from the MacRae (1972) model to the Beck and Wieland (2002) model. The parameter set for that model is

²³ We thank an anonymous referee for providing a better approximation of the optimal cost than the one originally found by the authors.

Figure 2. Approximate cost-to-go in the first time period for the MacRae (1972) model in the Monte Carlo *representative* run, i.e. when b is -0.5 .

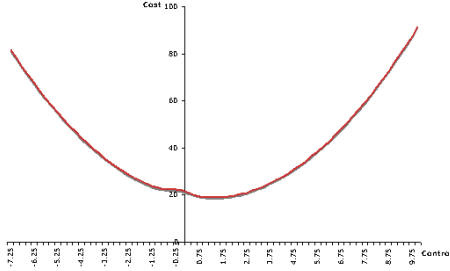
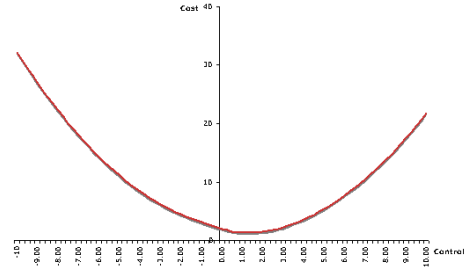


Figure 3. Approximate cost-to-go in the first time period for the Beck and Wieland (2002) model in the Monte Carlo *representative* run, i.e. when b is -0.5 .



$$\begin{aligned}
 \alpha &= 1. & \beta_0 \equiv b &= -0.5 & \gamma &= 0.0 & \sigma_\varepsilon^2 = q &= 1.0 \\
 w_k &= 1 \quad \forall k & \lambda_k &= 0 \quad \forall k & \sigma_\eta^2 \equiv \sigma_b^2 &= .25 & \sigma_\alpha^2 = \sigma_\gamma^2 &= 0
 \end{aligned} \tag{34}$$

and the desired paths for the state and control variables are set equal to zero. Also the initial state is $x_0 = 1$ as set in Amman et al. (2007). In this case the coefficients of the probing component are

$$\phi_1 = \left(\frac{1}{b^2}\right) = 4, \quad \phi_2 = \alpha b(-1) = .5, \quad \phi_3 = -x_0 = -1 \tag{35}$$

and this function reaches its minimum 0 at $r_1 = -2$ and its maximum $J_{P,2} = 1$ at $r_2 = 2$. The limits of the probing cost function as u_0 approaches plus or minus infinity are equal to $J_{P,2} = .5$ and, again, it has the general form shown in Figure 1, page 10. From equation (35) it follows that setting no penalty on the use of the control variable, i.e. $\lambda_k = 0$, makes the probing component of the cost-to-go extremely sensitive to the value of b . Thus in equation (11) with $w_2 = 1$, the parameter ϕ_1 reduces to the inverse of b^2 . Furthermore assuming a null intercept in equation (11) with $w_2 = 1$, i.e. $\gamma = 0$, makes ϕ_3 solely dependent on the initial state.

The coefficients for the cautionary component of the approximate cost-to-go are

$$\nu_1 = \alpha \left(1 - \frac{b^2}{b^2}\right) = 0, \quad \nu_2 = \left(-\frac{1}{b^2}\right) \alpha^2 b x_0 = -\frac{1}{b} = 2, \quad \nu_3 = \alpha^2 x_0 = 1 \quad (36)$$

therefore

$$\begin{aligned} \delta_1 &= \frac{\sigma_b^2}{2} \left[\nu_1^2 \left(1 - \frac{4b^2}{b^2}\right) + 1 \right] = .125 \\ \delta_2 &= \sigma_b^2 \nu_1 \left\{ \nu_2 - \frac{2(2b\nu_2 + \nu_3)}{b} \right\} = 0 \\ \delta_3 &= \frac{\sigma_b^2}{2} \left[\nu_2^2 - \frac{(2b\nu_2 + \nu_3)^2}{b^2} \right] + q = q = 1 \end{aligned} \quad (37)$$

and the cautionary component is convex. From equation (37) it is apparent that δ_3 is always equal to q in this model. Moreover these coefficients appear to be independent of the value of b with δ_2 always equal to zero. Finally the deterministic component

$$J_{D,2} = \psi_1 u_0^2 + \psi_2 u_0 + \psi_3 \quad (38)$$

has coefficients

$$\begin{aligned} \psi_1 &= \frac{1}{2} (b\nu_1)^2 + \frac{1}{2} b^2 = .125 \\ \psi_2 &= b\nu_1 (b\nu_2 + \nu_3) + \alpha x_0 b = -.5 \\ \psi_3 &= \frac{1}{2} (b\nu_2 + \nu_3)^2 + \frac{1}{2} (\alpha x_0)^2 = .5 \end{aligned} \quad (39)$$

and the *representative* cost-to-go function for the Beck and Wieland (2002) model can be plotted as in Figure 3, page 15. The optimum cost in Figure 3, page 15, is 1.32 at $u_0 = 1.25$.

7 The Parameter Set and Monte Carlo Experiments

The analysis carried out in the last section focuses on the shape of the *representative* cost-to-go. Therefore it sheds little light on the results obtained in each run of a Monte Carlo experiments carried out using the MacRae (1972), or Beck and Wieland (2002), models and parameter sets. In a typical Monte Carlo experiment the initial estimate of the unknown parameter β , i.e. b , in each run is drawn from a distribution having the true value of the unknown parameter as mean and the es-

estimated variance σ_b^2 as variance (Figure 4, page 18).²⁴ This means that, when the MacRae problem is used, the actual value for b is between -1.2 and 0.2 in approximately 68% of the Monte Carlo runs, in the interval (-1.9, 0.9) approximately 95% of times and over 99% of time between -2.6 and 1.6.²⁵ The same intervals for the Beck and Wieland problem are, respectively, (-1.0, 0.0), (-1.5, 0.5) and (-2.0, 1.0).

It is therefore useful to study the shape of the approximate cost-to-go when the parameter b takes on values in these different intervals. Even though the probability that b takes on a specific value is zero, as is well known, the probability that it falls in a certain interval, no matter how small, is not. Amman and Kendrick (1995, page 470) consider the effects of changes in the b parameter and they notice that for the MacRae problem *the cost-to-go function will be convex for values of b either substantially above or below zero but will be nonconvex for values of b close to zero*. However they do not explicate the relationship between this result and the outcome of Monte Carlo experiments. This is the main task of the present section.

In the MacRae problem, as originally suggested in Amman and Kendrick (1995), the approximate cost-to-go shows a nonconvexity when b is close to 0. Indeed a nonconvexity appears in all those runs where the value of b is between -0.45 and 0.45 and it becomes more pronounced as b gets closer to 0 (Figure 5, page 19, and Figure 6, page 19).²⁶ The probability for b to fall in the interval (-0.45, 0.45) in this problem is equal to the probability for the standard normal to be in the interval (0.07143, 1.3571). Therefore in a Monte Carlo experiment based on the MacRae (1972) problem nonconvexities will be encountered in about 38% of the runs which is even higher than the result, namely 28%, reported in Tucci (1998, page 216) and Tucci (2004, page 105). When the value for b is outside the interval (-0.45, 0.45) the cost-to-go appears as in Figure 7, page 19.

Some relevant information about a certain set of parameters may be gained studying the relationship between the stochastic parameter b and the optimal control, for the first period, and its associated cost-to-go. This can be done, for each value of b , by computing the coefficients of the approximated cost-to-go, see Appendix B, and

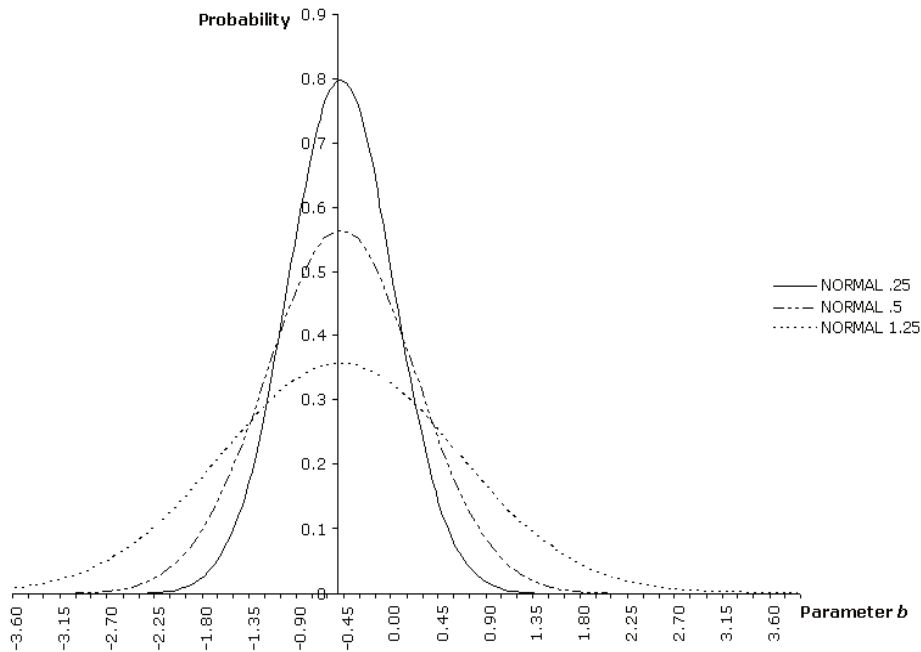
²⁴ This is equivalent to assuming that the estimator of the unknown parameter is unbiased.

²⁵ In the first time period (0 time period) when computing u_0 the only uncertain element is the b , i.e. θ_0 . Therefore u_0 depends only on the uncertain b but u_1 will depend also on the additive noise terms as well as the uncertain parameter.

²⁶ In the MacRae (1972) problem the optimal control for period 0 is -1.25, with an associated cost of 19.163, when $b = 0.4$. A local optimum is at $u_0 = 0.75$ with a cost of 21.96. The maximum cost between these two optima is 22.68 and is associated with $u_0 = 0$. When $b = -0.4$ the cost function looks like in a mirror with the global optimum at $u_0 = 1.25$, the local optimum at $u_0 = -0.75$ with a cost of 19.163 and 21.96, respectively, and the maximum cost between the two optima is again 22.68 at $u_0 = 0$. The cost function is perfectly symmetric around $u_0 = 0$ when $b = 0$. In this case there are *two global optima* at $u_0 = -1.25$ and $u_0 = 1.25$ with a cost equal to 18.392, and the maximum cost between the two optima is 24.075 at $u_0 = 0$.

Figure 4.

Probability density function of the unknown parameter b with the Beck and Wieland (2002) parameter set (NORMAL.25), with the MacRae (1972) parameter set (NORMAL.5) and with the Amman et al. (2007) parameter set (NORMAL1.25)



selecting the optimal control. When b is in the interval $(-3.0, -0.02)$, the optimal control is always between 1.00 and 2.00 with a cost ranging from 2.08 to 19.16 (Figure 8, page 20).²⁷

In the Beck and Wieland problem, it follows from equation (35) that when $b = 0$ the parameter ϕ_1 is infinity and as reported in Appendix B most of the relevant coefficients of the approximate total cost-to-go are not available because they include a *division by zero* term. When the parameter b happens to be around 0, say 0.02 as in Figure 9, page 21, ϕ_1 is extremely high and the approximate total cost-to-go takes on almost completely the shape of its probing component, see Appendix B.

²⁷ The optimal control associated with $b = 3.0$ is the opposite of the optimal control for $b = -3.0$ and the approximate cost-to-go is the same for the two controls. The same is true for all the positive values of b .

Figure 5. Approximate cost-to-go in the first time period for the MacRae (1972) problem in the Monte Carlo run where b is 0.4.

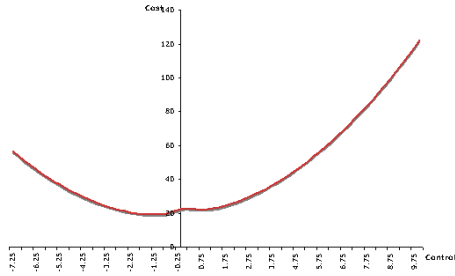


Figure 6. Approximate cost-to-go in the first time period for the MacRae (1972) problem in the Monte Carlo run where b is 0.0.

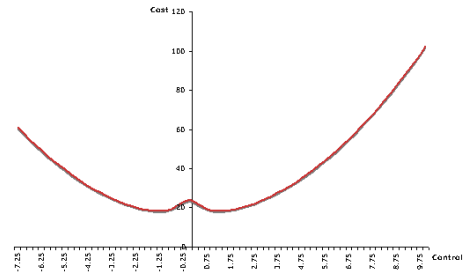
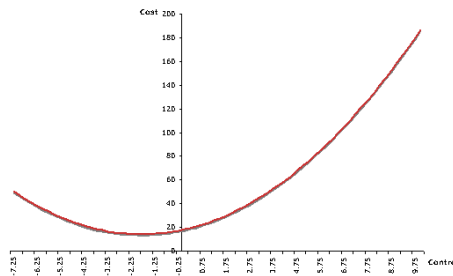


Figure 7. Approximate cost-to-go in the first time period for the MacRae (1972) problem in the Monte Carlo run where b is 0.9. The optimal control is -1.75 with a cost of 14.127.



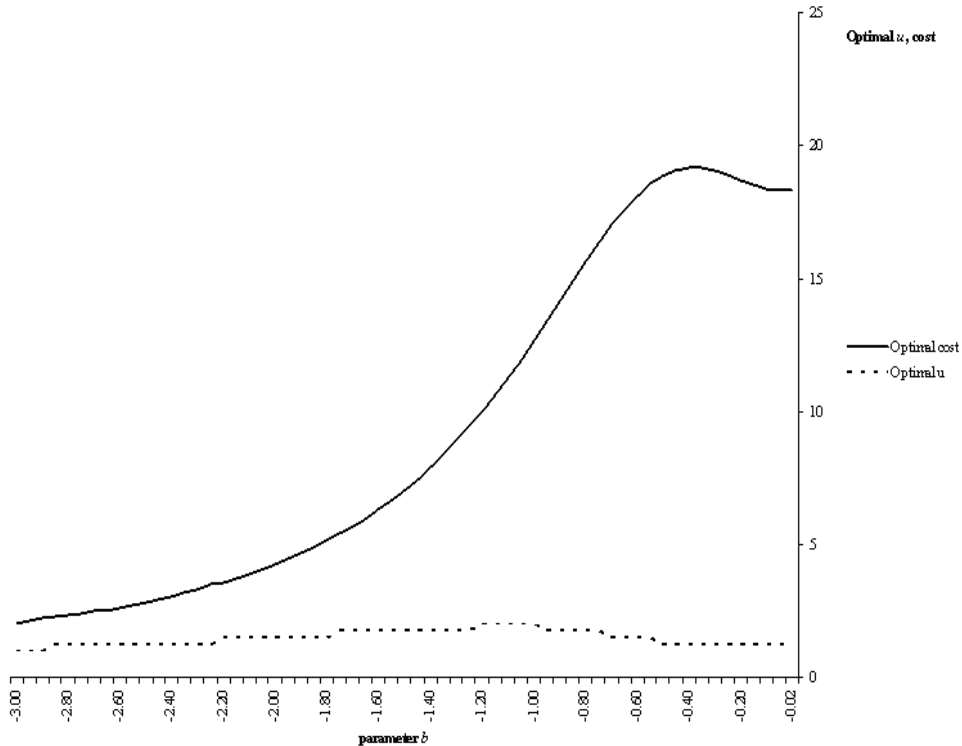
Also recall that in the Beck and Wieland model the weight λ_1 on the control variable deviation in the criterion function is set to zero. When this is combined with a Monte Carlo run in which the coefficient b on the control variable term in the systems equations is near zero as was discussed above, then from equation (11) the coefficient

$$\phi_1 = \left(\frac{w_2^2}{\lambda_1 + b^2 w_2} \right)$$

in the probing term of the cost-to-go becomes very large. As Figure 9, page 21, shows the probing term is large enough in this case to make the entire cost-to-go function nonconvex.

The intuition here is that the probing component in the cost-to-go provides a mea-

Figure 8.
Relationship between b and the optimal control in the first time period, and its approximate cost-to-go, for the MacRae (1972) problem.



sure of future cost depending on the amount of probing. If there is not enough probing to move the state variable x_1 around, the cost is high because the unknown parameter value is not learned effectively and its variance remains high. Thus there is an incentive, as Figure 9, page 21, shows, for the control to be made large in either the positive or negative direction in order to learn the parameter and reduce the cost. However, when b is very small it takes large movements in the control variable u_0 to bring about shifts in the state variable x_1 in order to learn the coefficient b and reduce its variance. Also, in the Beck and Wieland model, with λ_1 equal to zero there is no cost in the criterion function associated with using large positive or negative values of the control. Thus there is a need, when b is near zero, for large values of the control in order to learn and there is no constraint from the criterion cost. These two factors together bring about the shape of the cost-to-go function shown in Figure 9, page 21.

Moreover it should be noticed that in the Beck and Wieland (2002) problem the cautionary cost is constant in the various runs, i.e. is independent of b , as indicated by the δ coefficients reported in Table B.2, page 33, of Appendix B. When the

Figure 9. Approximate cost-to-go in the first time period for the Beck and Wieland (2002) problem in the Monte Carlo run where b is 0.02.

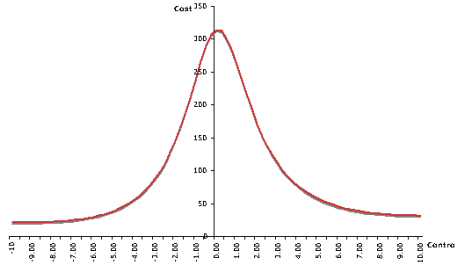


Figure 10. Approximate cost-to-go in the first time period for the Beck and Wieland (2002) problem in the Monte Carlo run where b is 0.1.

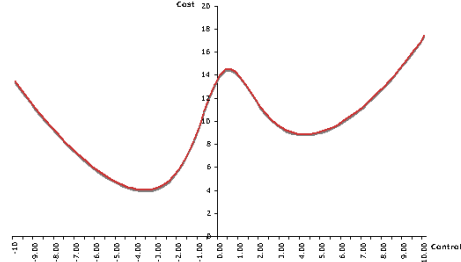


Figure 11. Approximate cost-to-go in the first time period for the Beck and Wieland (2002) problem in the Monte Carlo run where b is 0.2.

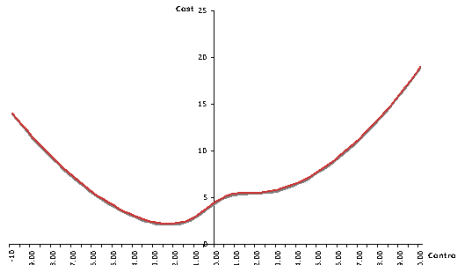
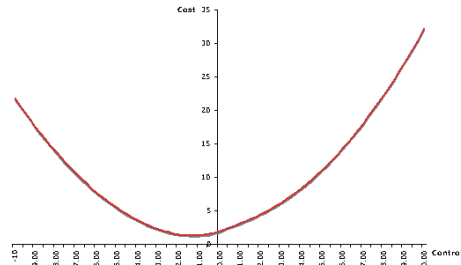


Figure 12. Approximate cost-to-go in the first time period for the Beck and Wieland (2002) problem in the Monte Carlo run where b is 0.5.



absolute value of b is in the interval $(0.03, 0.19)$ the cost-to-go function has a *sombrero like* shape as in Figure 10, page 21.²⁸ The nonconvexity is very pronounced for values of b close to the lower end of the interval and it becomes less and less prominent as b moves toward the higher end of the interval. For $b = 0.2$ the cost-to-go appears as in Figure 11, page 21. The *flat looking* portion around a control equal to 1 is due to the fact that in that portion the cost-to-go increases at a slower pace. Namely, when the control at time zero is increased from 0.5 to 0.75, 1, 1.25,

²⁸ In the Beck and Wieland (2002) problem the optimal control for period 0 is 3.75, with an associated cost of 4.034, when $b = -0.1$. A local optimum is at $u_0 = -4.25$ with a cost of 8.875. The maximum cost between these two optima is 14.55 and is associated with $u_0 = -0.50$. From Figure 10, page 21, when $b = 0.1$ the cost function looks like in a mirror with the global optimum at $u_0 = -3.75$, the local optimum at $u_0 = 4.25$ with a global optimum cost of 4.034 and a local optimum cost of 8.875, respectively, and the maximum cost between the two optima is again 14.55 at $u_0 = 0.50$.

1.5, 1.75 and 2 the approximate cost-to-go increases from 5.20 to 5.35, 5.45, 5.49, 5.51, 5.52 and 5.54 respectively.

The probability for b to fall in the interval $(-0.2, 0.2)$ in the Beck and Wieland problem is equal to the probability for the standard normal to be in the interval $(0.6, 1.4)$. Therefore, in a Monte Carlo experiment based on this model and parameter set, from mild to severe nonconvexities will be encountered in about 20% of the runs. Half of them will be severe and around 2% will be associated with a cost-to-go as in Figure 9, page 21. When the value for b is outside the interval $(-0.2, 0.2)$ the cost-to-go appears as in Figure 12, page 21. When the value for b is outside the interval $(-0.2, 0.2)$ the cost-to-go appears as in Figure 12, page 21.

The probability that the value of b used in the Monte Carlo run falls in the intervals $(-0.1, 0.1)$ and $(-0.02, 0.02)$ in the Beck and Wieland (2002) problem is equal to the probability for the standard normal to be in the intervals $(0.8, 1.2)$ and $(.96, 1.04)$, respectively.

Summarizing (as shown in Figure 13, page 23) when b is in the interval $(-3.0, -0.1)$, the optimal control goes from .25 to 3.75 with a cost ranging from 1.04 to 4.03.²⁹ As b gets closer to 0 both the optimal control and its associated cost skyrocket. At $b = -.02$ the former is 9.25 and the latter is 21.30.

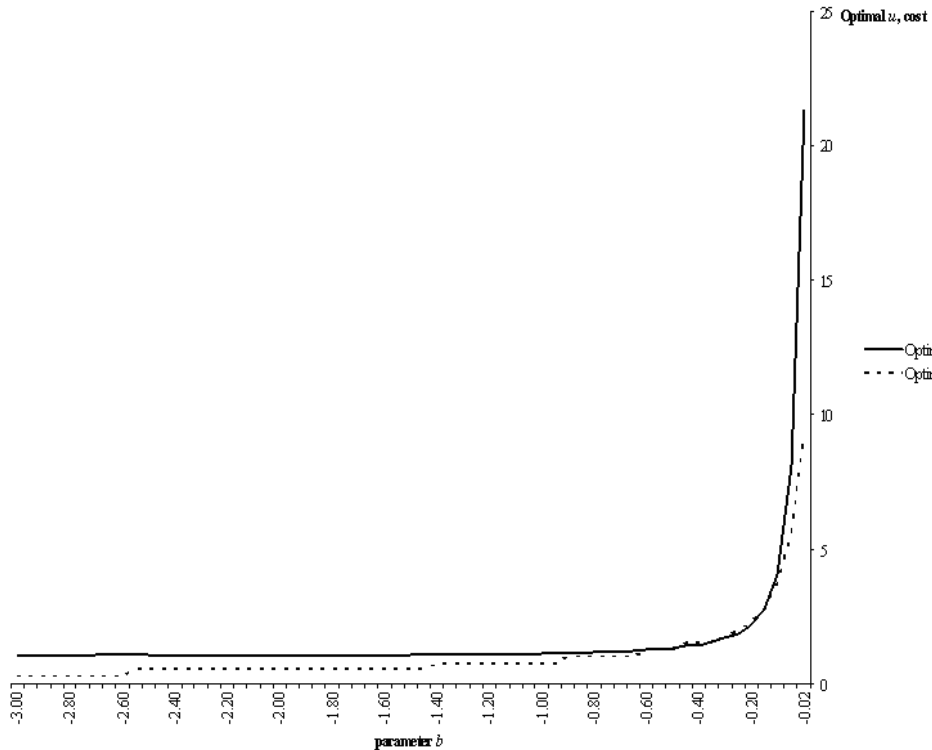
The discussion of this section reveals that the variance of the unknown parameter estimate plays a *dual role* in making nonconvexities more or less likely to occur in a certain Monte Carlo experiment. On one hand it enters the coefficients of the approximate cost-to-go and its second order derivative with respect to the control, as pointed out in Amman and Kendrick (1994, 1995), effecting directly the shape of the cost function. On the other hand it affects the probability with which values of the parameter drawn from certain intervals of the parameter space are used. Therefore the role of σ_b^2 is more subtle than so far expected. Nonconvexities may appear frequently even in cases where the *representative* run shows no indication of it and a smaller variance does not necessarily imply fewer nonconvexities. For instance, in the Beck and Wieland (2002) problem when $\sigma_b^2 = 1.25$ the probability that b is drawn from the interval $(-0.02, 0.02)$ is approximately 1%, instead of 2% as with the original data set where $\sigma_b^2 = 0.25$.³⁰ This means that cost functions similar to that reported in Figure 9, page 21, will be encountered less frequently with this higher variance.

This completes the discussion of the main line of argument of this paper about

²⁹ Again the optimal control associated with a positive b is the opposite of the optimal control for a negative b and the approximate cost-to-go is the same in both cases.

³⁰ This probability is approximately equal to the probability that the standard normal is in the interval $(0.43, 0.47)$.

Figure 13.
Relationship between b and the optimal control in the first time period, and its approximate cost-to-go, for the Beck and Wieland (2002) problem.



the effect of parameter sets on the adaptive or dual control solutions. However, before closing it is appropriate to add a short additional section on the effect of the parameter set selected for a certain numerical experiment on the optimal feedback (OF) and expected optimal feedback (EOF) solutions since it is anticipated that comparisons of these two methods may occur with some frequency in the future.

8 The Importance of Penalty Weights on the OF and EOF Solutions

Suppose that the optimal feedback (OF) control, the expected optimal feedback (EOF) control and Dual control (DC) are compared.³¹ Dual control minimizes the approximate cost-to-go defined in equation (27). The EOF control is determined using the feedback rule

³¹ See Kendrick and Amman (2006) for a description of the three different methods.

$$u_0 = G_0 x_0 + g_0 \quad (40)$$

with

$$\begin{aligned} G_0 &= -[\lambda_0 + K_1 (b^2 + \sigma_b^2)]^{-1} (\alpha K_1 b) \\ g_0 &= -[\lambda_0 + K_1 (b^2 + \sigma_b^2)]^{-1} (b K_1 \gamma + b p_1) \end{aligned} \quad (41)$$

and the Riccati quantities K_1 and p_1 defined as

$$\begin{aligned} K_1 &= w_1 + \alpha^2 K_2 - (\alpha b K_2) [\lambda_1 + K_2 (b^2 + \sigma_b^2)]^{-1} (\alpha b K_2) \\ p_1 &= (\alpha \gamma K_2) + \alpha p_2 - (\alpha b K_2) [\lambda_1 + K_2 (b^2 + \sigma_b^2)]^{-1} [b \gamma K_2 + b p_2] \end{aligned} \quad (42)$$

where the terminal conditions are given by $K_2 = 1$ and $p_2 = 0$.³² The OF control is obtained using equations (40)-(42) with $\sigma_b^2 = 0$. When the parameter b is very close to zero, say $b = 0.02$, the EOF control is equal to 0, OF control is infinity and the DUAL control diverges toward very high or low values trying to minimize a function like that depicted in Figure 9 on page 21. Consequently, in a Monte Carlo experiment based on the Beck and Wieland (2002) problem both OF and Dual control will exhibit around 10% of very unfavorable runs and around 2% of severe outliers compared with EOF. As indicated in the sixth column of Table 1, page 32, the EOF control is not too affected by a parameter value for b close or equal to 0 in the Beck and Wieland (2002) problem. It goes from a maximum of 1 to a minimum of minus 1 for the selected b 's.

Completely different is the situation for the OF control which, for the set of selected values of b , ranges from infinity to -100 as reported in the last column of Table 1, page 32. The situation looks more balanced when the MacRae problem is used. As reported in Table 2, page 33, the difference between EOF and OF control is not too big and as shown in Figure 6, page 19, the shape of the cost-to-go is not such as to cause the Dual algorithm to diverge.

9 Conclusion

As work on the application of various stochastic control methods to economic models proceeds Monte Carlo and similar stochastic methods will probably be used to compare the performance of the various methods on simple models. When this is

³² See, e.g., Kendrick (1981, 2002, Chapter 6).

done outliers of the type we have encountered in doing such studies on the MacRae (MacRae (1972)) and the Beck and Wieland (Beck and Wieland (2002)) models are likely to occur. We have traced a major cause of these outliers to the choice of parameter sets and in particular to the weight on the control term in the criterion function and the estimate of the unknown parameter that is multiplied by the control variable. When $\lambda = 0$ and b is drawn near zero in the Monte Carlo runs the cost-to-go function is likely to be nonconvex and the criterion value may be large.

Therefore, in conducting Monte Carlo experiments to compare various methods of solving stochastic control methods in quadratic-linear tracking models it is important to be careful in the choice of parameter sets. We would recommend that whenever possible that calculations similar to those we have performed here be done before large Monte Carlo runs are done. As indicated above, we are happy to provide to interested users the spreadsheet we have used to make the calculations in Sections 7 and 8 in this paper and which can easily be modified for other parameter sets in one-state, one-control, one uncertain parameter cases.

Table 1

Expected optimal feedback (EOF) control and optimal feedback ((OF) control for the Beck and Wieland (2002) parameter set with $x_0 = 1$ at various values of b (Note: n.a.= not available because *dividing by zero*)

EOF b	CASE K_1	p_1	G_0	g_0	u_0	OF K_1	CASE p_1	G_0	g_0	u_0
-2.00	1.0588	0	0.4706	0	0.4706	1	0	0.5000	0	0.5000
-1.50	1.1000	0	0.6000	0	0.6000	1	0	0.6667	0	0.6667
-1.00	1.2000	0	0.8000	0	0.8000	1	0	1.0000	0	1.0000
-0.50	1.5000	0	1.0000	0	1.0000	1	0	2.0000	0	2.0000
-0.10	1.9615	0	0.3846	0	0.3846	1	0	10.0000	0	10.0000
0.00	2.0000	0	0.0000	0	0.0000	n.a.	n.a.	n.a.	n.a.	n.a.
0.01	1.9996	0	-0.0400	0	-0.0400	1	0	-100.0000	0	-100.0000
0.10	1.9615	0	-0.3846	0	-0.3846	1	0	-10.0000	0	-10.0000
0.15	1.9174	0	-0.5505	0	-0.5505	1	0	-6.6667	0	-6.6667
0.20	1.8621	0	-0.6897	0	-0.6897	1	0	-5.0000	0	-5.0000
0.50	1.5000	0	-1.0000	0	-1.0000	1	0	-2.0000	0	-2.0000
1.00	1.2000	0	-0.8000	0	-0.8000	1	0	-1.0000	0	-1.0000

Table 2

Expected optimal feedback (EOF) control and optimal feedback ((OF) control for the MacRae (1972) parameter set at various values of b

<i>EOF</i> b	<i>CASE</i> K_1	p_1	G_0	g_0	u_0	<i>OF</i> K_1	<i>CASE</i> p_1	G_0	g_0	u_0
-2.60	1.0890	0.4449	0.2225	1.2426	1.2426	1.0631	0.3157	0.2363	1.2820	1.2820
-1.90	1.1438	0.7192	0.2668	1.5739	1.5739	1.1063	0.5315	0.2946	1.6754	1.6754
-1.20	1.2500	1.2500	0.3066	1.9708	1.9708	1.2008	1.0041	0.3696	2.2895	2.2895
-0.50	1.4200	2.1000	0.2407	1.7119	1.7119	1.3920	1.9600	0.3614	2.5341	2.5341
-0.25	1.4704	2.3520	0.1408	1.0260	1.0260	1.4612	2.3059	0.2343	1.6998	1.6998
-0.20	1.4773	2.3864	0.1150	0.8407	0.8407	1.4712	2.3558	0.1945	1.4175	1.4175
0.00	1.4900	2.4500	0.0000	0.0000	0.0000	1.4900	2.4500	0.0000	0.0000	0.0000
0.20	1.4773	2.3864	-0.1150	-0.8407	-0.8407	1.4712	2.3558	-0.1945	-1.4175	-1.4175
0.25	1.4704	2.3520	-0.1408	-1.0260	-1.0260	1.4612	2.3059	-0.2343	-1.6998	-1.6998
0.30	1.4623	2.3113	-0.1649	-1.1965	-1.1965	1.4495	2.2477	-0.2693	-1.9429	-1.9429
0.90	1.3182	1.5909	-0.3046	-2.0478	-2.0478	1.2707	1.3536	-0.3945	-2.5728	-2.5728
1.60	1.1810	0.9052	-0.2867	-1.7473	-1.7473	1.1376	0.6882	-0.3257	-1.9098	-1.9098

A Appendix: Intermediate Steps in the Derivation of the Cautionary and Deterministic Components of the Cost-To-Go

First derive equation (19) and equation (20) – the cautionary cost term. To do so begin with equation (18), namely

$$J_{C,2} = \frac{\sigma_b^2 w_2}{2} (\alpha u_0 + u_{01})^2 + \frac{\sigma_b^2}{2} \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2 u_0 + b w_2 u_{01} + w_2 x_{02} - w_2 \tilde{x}_2)^2 + \frac{q}{2} \left[\alpha^2 w_2 + w_2 + w_1 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2)^2 \right] + \frac{\sigma_b^2 w_1}{2} u_0^2, \quad (\text{A-1})$$

and replacing u_{01} and x_{02} by their definition, that is (8) and (9) respectively, it yields

$$J_{C,2} = \frac{\sigma_b^2 w_2}{2} \left(\alpha u_0 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) [\alpha b^2 w_2 u_0 + \alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 (\gamma - \tilde{x}_2) - \lambda_1 \tilde{u}_1] \right)^2 + \frac{\sigma_b^2}{2} \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) \left\{ \alpha b w_2 u_0 + b w_2 \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) [\alpha b^2 w_2 u_0 + \alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 (\gamma - \tilde{x}_2) - \lambda_1 \tilde{u}_1] + w_2 \left[b \left(\alpha - \frac{\alpha b^2 w_2}{\lambda_1 + b^2 w_2} \right) u_0 + \alpha^2 x_0 + \alpha \gamma + \gamma + b \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 \gamma - b w_2 \tilde{x}_2 - \lambda_1 \tilde{u}_1) \right] - w_2 \tilde{x}_2 \right\}^2 + \frac{q}{2} \left[\alpha^2 w_2 + w_2 + w_1 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2)^2 \right] + \frac{\sigma_b^2 w_1}{2} u_0^2. \quad (\text{A-2})$$

or

$$J_{C,2} = \frac{\sigma_b^2 w_2}{2} \left(\alpha u_0 - \left\{ \frac{1}{\lambda_1 + b^2 w_2} \right\} \alpha b^2 w_2 u_0 - \left\{ \frac{1}{\lambda_1 + b^2 w_2} \right\} [\alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 (\gamma - \tilde{x}_2) - \lambda_1 \tilde{u}_1] \right)^2 + \frac{\sigma_b^2}{2} \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) \left\{ \alpha b w_2 u_0 + b w_2 \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) \alpha b^2 w_2 u_0 + b w_2 \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) [\alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 (\gamma - \tilde{x}_2) - \lambda_1 \tilde{u}_1] + w_2 \left[b \left(\alpha - \frac{\alpha b^2 w_2}{\lambda_1 + b^2 w_2} \right) u_0 + \alpha^2 x_0 + \alpha \gamma + \gamma + b \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 \gamma - b w_2 \tilde{x}_2 - \lambda_1 \tilde{u}_1) \right] - w_2 \tilde{x}_2 \right\}^2 + \frac{q}{2} \left[\alpha^2 w_2 + w_2 + w_1 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2)^2 \right] + \frac{\sigma_b^2 w_1}{2} u_0^2. \quad (\text{A-3})$$

Use the quantities ν_1 , ν_2 and ν_3 as defined in equation (21), i.e.

$$\begin{aligned}
\nu_1 &= \alpha \left(1 - \frac{b^2 w_2}{\lambda_1 + b^2 w_2} \right) \\
\nu_2 &= \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) [\alpha^2 b w_2 x_0 + \alpha b \gamma w_2 + b w_2 (\gamma - \tilde{x}_2) - \lambda_1 \tilde{u}_1] \\
\nu_3 &= \alpha^2 x_0 + \alpha \gamma + \gamma - \tilde{x}_2
\end{aligned} \tag{A-4}$$

and rewrite equation (A-3) as

$$\begin{aligned}
J_{C,2} &= \frac{\sigma_b^2 w_2}{2} (\nu_1 u_0 + \nu_2)^2 + \frac{\sigma_b^2}{2} \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) \{ b w_2 \nu_1 u_0 + b w_2 \nu_2 + w_2 [b \nu_1 u_0 + \nu_3 + b \nu_2] \}^2 \\
&+ \frac{q}{2} \left[\alpha^2 w_2 + w_2 + w_1 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2)^2 \right] + \frac{\sigma_b^2 w_1}{2} u_0^2 \\
&= \frac{\sigma_b^2 w_2}{2} (\nu_1 u_0 + \nu_2)^2 + \frac{\sigma_b^2}{2} \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) \{ 2 b w_2 \nu_1 u_0 + 2 b w_2 \nu_2 + w_2 \nu_3 \}^2 \\
&+ \frac{q}{2} \left[\alpha^2 w_2 + w_2 + w_1 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2)^2 \right] + \frac{\sigma_b^2 w_1}{2} u_0^2.
\end{aligned} \tag{A-5}$$

Squaring the terms in equation (A-5) yields

$$\begin{aligned}
J_{C,2} &= \frac{\sigma_b^2 w_2}{2} (\nu_1^2 u_0^2 + \nu_2^2 + 2 \nu_1 u_0 \nu_2) + \frac{\sigma_b^2}{2} \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) \\
&\times \left[(2 b w_2 \nu_1)^2 u_0^2 + (2 b w_2 \nu_2 + w_2 \nu_3)^2 + 2 (2 b w_2 \nu_1) (2 b w_2 \nu_2 + w_2 \nu_3) u_0 \right] \\
&+ \frac{q}{2} \left[\alpha^2 w_2 + w_2 + w_1 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2)^2 \right] + \frac{\sigma_b^2 w_1}{2} u_0^2 \\
&= \frac{\sigma_b^2}{2} \left[w_2 \nu_1^2 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (2 b w_2 \nu_1)^2 \right] u_0^2 \\
&+ \frac{\sigma_b^2}{2} \left[2 \nu_1 \nu_2 w_2 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) 2 (2 b w_2 \nu_1) (2 b w_2 \nu_2 + w_2 \nu_3) \right] u_0 \\
&+ \frac{\sigma_b^2}{2} \left[\nu_2^2 w_2 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (2 b w_2 \nu_2 + w_2 \nu_3)^2 \right] \\
&+ \frac{q}{2} \left[\alpha^2 w_2 + w_2 + w_1 + \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) (\alpha b w_2)^2 \right] + \frac{\sigma_b^2 w_1}{2} u_0^2.
\end{aligned} \tag{A-6}$$

or, alternatively,

$$\begin{aligned}
J_{C,2} &= \frac{\sigma_b^2}{2} \left[w_2 \nu_1^2 \left(1 - \frac{4 b^2 w_2}{\lambda_1 + b^2 w_2} \right) + w_1 \right] u_0^2 + \sigma_b^2 \nu_1 w_2 \left[\nu_2 \left(1 - \frac{4 b^2 w_2}{\lambda_1 + b^2 w_2} \right) - \frac{2 b w_2 \nu_3}{\lambda_1 + b^2 w_2} \right] u_0 \\
&+ \frac{\sigma_b^2}{2} \left[\nu_2^2 w_2 - \frac{(2 b w_2 \nu_2 + w_2 \nu_3)^2}{\lambda_1 + b^2 w_2} \right] + \frac{q}{2} \left[\alpha^2 w_2 + w_2 + w_1 - \frac{(\alpha b w_2)^2}{\lambda_1 + b^2 w_2} \right].
\end{aligned} \tag{A-7}$$

which is identical to equation (19) when the δ 's are defined as in (20).

Next derive equation (24) and equation (25) – the deterministic component. Begin

with equation (23) i.e.

$$J_{D,2} = \frac{\lambda_0}{2} (u_0 - \tilde{u}_0)^2 + \frac{w_2}{2} (x_{02} - \tilde{x}_2)^2 + \frac{w_1}{2} (x_{01} - \tilde{x}_1)^2 + \frac{\lambda_1}{2} (u_{01} - \tilde{u}_1)^2 \quad (\text{A-8})$$

Again by replacing u_{01} and x_{02} by their definition, that is (8) and (9) respectively, using $x_{01} = \alpha x_0 + b u_0 + \gamma$, and equation (21) one obtains

$$J_{D,2} = \frac{\lambda_0}{2} (u_0 - \tilde{u}_0)^2 + \frac{w_2}{2} \{b\nu_1 u_0 + b\nu_2 + \nu_3\}^2 + \frac{w_1}{2} (\alpha x_0 + b u_0 + \gamma - \tilde{x}_1)^2 + \frac{\lambda_1}{2} \left\{ \left(-\frac{1}{\lambda_1 + b^2 w_2} \right) \alpha b^2 w_2 u_0 + \nu_2 - \tilde{u}_1 \right\}^2 \quad (\text{A-9})$$

Squaring and rearranging the terms in equation (A-9) yields

$$J_{D,2} = \frac{\lambda_0}{2} (u_0^2 - 2u_0\tilde{u}_0 + \tilde{u}_0^2) + \frac{w_2}{2} \left\{ (b\nu_1)^2 u_0^2 + 2(b\nu_1)(b\nu_2 + \nu_3) u_0 + (b\nu_2 + \nu_3)^2 \right\} + \frac{w_1}{2} \left[b^2 u_0^2 + 2b(\alpha x_0 + \gamma - \tilde{x}_1) u_0 + (\alpha x_0 + \gamma - \tilde{x}_1)^2 \right] + \frac{\lambda_1}{2} \left\{ \left(-\frac{1}{\lambda_1 + b^2 w_2} \right)^2 (\alpha b^2 w_2)^2 u_0^2 - \frac{2}{\lambda_1 + b^2 w_2} \alpha b^2 w_2 (\nu_2 - \tilde{u}_1) u_0 + (\nu_2 - \tilde{u}_1)^2 \right\} \quad (\text{A-10})$$

or, alternatively,

$$J_{D,2} = \left[\frac{\lambda_0}{2} + \frac{w_2}{2} (b\nu_1)^2 + \frac{w_1}{2} b^2 + \frac{\lambda_1}{2} \left(-\frac{1}{\lambda_1 + b^2 w_2} \right)^2 (\alpha b^2 w_2)^2 \right] u_0^2 + \left\{ -\lambda_0 \tilde{u}_0 + w_2 b \nu_1 (b\nu_2 + \nu_3) + w_1 b (\alpha x_0 + \gamma - \tilde{x}_1) - \frac{\lambda_1 \alpha b^2 w_2}{\lambda_1 + b^2 w_2} (\nu_2 - \tilde{u}_1) \right\} u_0 + \frac{\lambda_0}{2} (\tilde{u}_0^2) + \frac{w_2}{2} (b\nu_2 + \nu_3)^2 + \frac{w_1}{2} (\alpha x_0 + \gamma - \tilde{x}_1)^2 + \frac{\lambda_1}{2} (\nu_2 - \tilde{u}_1)^2 . \quad (\text{A-11})$$

Equation (A-11) is identical to (24) when the ψ 's are defined as in (25).

B Appendix: The Coefficients of the Cost-To-Go Function for Selected Values of b

This appendix contains some of the coefficients of the cost-to-go function, at selected values of b , used to derive Figure 8, page 20, and Figure 13, page 23. Table B.1, page 32, contains the coefficients for the MacRae (1972) problem. It is worthwhile to point out that the sign of b affects only the sign of the *linear coefficients* in the approximate cost-to-go, namely ϕ_2 , δ_2 and ψ_2 . Table B.2, page 33, contains the coefficients for the Beck and Wieland (2002) problem. For $b = 0$ most of the relevant coefficients of the approximate total cost-to go are not available because they include a *division by zero* term. When the parameter b happens to be around 0, say 0.02 as in Figure 9, page 21, is extremely high and the approximate total cost-to go takes on almost completely the shape of its probing component.

Accepted manuscript

Table B.1

Coefficients of the cost-to-go function for the MacRae (1972) problem at various values of b

b	ν_1	ν_2	ν_3	ϕ_1	ϕ_2	ϕ_3	δ_1	δ_2	δ_3	ψ_1	ψ_2	ψ_3
-2.60	0.090	1.994	5.95	0.129	1.351	-4.416	0.245	-0.044	0.571	4.093	-	8.406
-1.90	0.152	2.452	5.95	0.217	0.753	-3.369	0.238	-0.025	1.099	2.497	-8.367	9.965
-1.20	0.287	2.926	5.95	0.410	0.151	-1.073	0.222	0.268	2.243	1.365	-6.248	13.380
-0.50	0.560	2.380	5.95	0.800	-0.210	3.570	0.266	1.466	-0.894	0.674	-3.416	20.286
-0.25	0.659	1.400	5.95	0.941	-0.154	5.250	0.333	1.275	-5.749	0.546	-1.855	22.785
-0.20	0.673	1.144	5.95	0.962	-0.129	5.492	0.346	1.096	-6.677	0.529	-1.501	23.145
0.00	0.7	0	5.95	1	0	5.950	0.373	0	-8.602	0.500	0	23.826
0.20	0.673	-1.144	5.95	0.962	0.129	5.492	0.346	-1.096	-6.677	0.529	1.501	23.145
0.30	0.642	-1.638	5.95	0.917	0.175	4.967	0.319	-1.404	-4.744	0.565	2.196	22.365
0.40	0.603	-2.052	5.95	0.862	0.203	4.309	0.291	-1.516	-2.706	0.614	2.836	21.385
0.90	0.387	-2.959	5.95	0.552	0.066	0.625	0.220	-0.692	2.361	1.015	5.221	15.905
1.60	0.197	-2.674	5.95	0.281	-0.491	-2.607	0.232	-0.032	1.524	1.956	7.472	11.097

Table B.2

Coefficients of the cost-to-go function for the Beck and Wieland (2002) problem with $x_0 = 1$ at various values of b
 (Note: n.a. = not available because *dividing by zero*).

b	ν_1	ν_2	ν_3	ϕ_1	ϕ_2	ϕ_3	δ_1	δ_2	δ_3	ψ_1	ψ_2	ψ_3
-2.00	0	0.50	1	0.25	2	-1	0.125	0	1	2	-2	0.5
-1.50	0	0.67	1	0.44	1.50	-1	0.125	0	1	1.125	-1.50	0.5
-1.00	0	1	1	1	1	-1	0.125	0	1	0.500	-1	0.5
-0.50	0	2	1	4	0.50	-1	0.125	0	1	0.125	-0.50	0.5
-0.10	0	10	1	100	0.10	-1	0.125	0	1	0.005	-0.10	0.5
0.00	n.a.	n.a.	1	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
0.01	0	-100	1	10000	-0.01	-1	0.125	0	1	0.000	0.01	0.5
0.10	0	-10	1	100	-0.10	-1	0.125	0	1	0.005	0.10	0.5
0.15	0	-6.67	1	44.44	-0.15	-1	0.125	0	1	0.011	0.15	0.5
0.20	0	-5	1	25	-0.20	-1	0.125	0	1	0.020	0.20	0.5
0.50	0	-2	1	4	-0.50	-1	0.125	0	1	0.125	0.50	0.5
1.00	0	-1	1	1	-1	-1	0.125	0	1	0.500	1	0.5

References

- Amman, H. M. and Kendrick, D. A.: 1994, Nonconvexities in stochastic control problems: An analysis, in W. W. Cooper and A. B. Whinston (eds), *New Directions in Computational Economics*, Vol. 4 of *Advances in Computational Economics*, Kluwer Academic Publishers, Dordrecht, the Netherlands, pp. 57–94.
- Amman, H. M. and Kendrick, D. A.: 1995, Nonconvexities in stochastic control models, *International Economic Review* **36**, 455–475.
- Amman, H. M., Kendrick, D. A. and Tucci, M. P.: 2007, Solving the Beck and Wieland model with optimal experimentation in DUALPC, *Working paper*, Utrecht University, Heidelberglaan 8, 3584 CS Utrecht, the Netherlands.
- Beck, G. and Wieland, V.: 2002, Learning and control in a changing economic environment, *Journal of Economic Dynamics and Control* **26**, 1359–1378.
- Cosimano, T. F.: 2008, Optimal experimentation and the perturbation method in the neighborhood of the augmented linear regulator problem, *Journal of Economic Dynamics and Control* **32**, 1857–1894.
- Cosimano, T. F. and Gapen, M. T.: 2006, An algorithm for approximating optimal experimentation problems using the perturbation method, *Working paper*, Department of Finance, University of Notre Dame, Notre Dame, Indiana, USA.
- Harvey, A.: 1981, *Time series models*, Philip Allan, Oxford, UK.
- Kendrick, D. A.: 1978, Non-convexities from probing in adaptive control problems, *Economics Letters* **1**, 347–351.
- Kendrick, D. A.: 1981, *Stochastic control for economic models*, first edn, McGraw-Hill Book Company, New York, New York, USA. See also Kendrick (2002).
- Kendrick, D. A.: 2002, *Stochastic control for economic models*. 2nd edition available at
url: <http://www.eco.utexas.edu/faculty/Kendrick>.
- Kendrick, D. A. and Amman, H. M.: 2006, A classification system for economic stochastic control models, *Computational Economics* **27**, 453–481.
- MacRae, E. C.: 1972, Linear decision with experimentation, *Annals of Economic and Social Measurement* **1**, 437–448.
- MacRae, E. C.: 1975, An adaptive learning role for multiperiod decision problems, *Econometrica* **43**, 893–906.
- Mizrach, B.: 1991, Non-convexities in an stochastic control problem with learning, *Journal of Economic Dynamics and Control* **15**, 515–538.
- Norman, A. L., Norman, M. R. and Palash, C. J.: 1979, Multiple Relative Maxima in Optimal Macroeconomic Policy: An Illustration, *Southern Economic Journal* **46**, 274–729.
- Tse, E. and Bar-Shalom, Y.: 1973, An actively adaptive control for linear systems with random parameters, *IEEE Transaction on Automatic Control* **18**, 109–117.
- Tucci, M. P.: 1997, Adaptive control in the presence of time-varying parameters, *Journal of Economic Dynamics and Control* **22**, 39–47.
- Tucci, M. P.: 1998, The nonconvexities problem in adaptive control models: A simple computational solution, *Computational Economics* **12**, 203–222.
- Tucci, M. P.: 2004, *The Rational Expectation Hypothesis, Time-varying Parameters*

and Adaptive Control, Springer, Dordrecht, the Netherlands.

Tucci, M. P., Kendrick, D. A. and Amman, H. M.: 2007, Expected optimal feedback with time-varying parameters, *Quaderni del Dipartimento di Economia Politica* 497, Università di Siena, Siena, Italy.

Document name: main.tex

Date: 22 June 2010

Time: 21:39

Accepted manuscript