Discounted Linear Exponential Quadratic Gaussian Control

by

Lars Peter Hansen

University of Chicago, NBER and NORC

and

Thomas J. Sargent

Hoover Institution, Stanford University

and

University of Chicago

Revised: September 1994

Abstract:

We describe a recursive formulation of discounted costs for a linear quadratic exponential Gaussian linear regulator problem which implies time-invariant linear decision rules in the infinite horizon case. Time invariance in the discounted case is attained by surrendering state-separability of the risk-adjusted costs.

Address:

Thomas J. Sargent, Hoover Institution, Stanford, California 94305; e-mail address: `sargent@riffle.stanford.edu`.

This research was funded by grants from the National Science Foundation. Wen Fang Liu, Alex Taber, and Thomas Tallarini, Jr. provided helpful comments on a previous draft.

Discounted Linear Exponential Quadratic Gaussian Control

This paper formulates a version of a *discounted* Gaussian optimal linear regulator in which the return function is modified as suggested in [5], [6], [13]-[15] to incorporate a risk-adjustment. In [5], the problem is formulated for the *undiscounted* case. Contributions [2] and [15] described how recursions on a Riccati difference equation apply to a discounted version of the problem. In their formulation with discounting, the optimal decision rules fail to be time-invariant: over time the effects of the risk-parameter 'wear off,' and the decision rules eventually converge to what would prevail in the usual linear-quadratic case.

We propose an alternative discounted version of the problem that preserves time-invariance of the decision rules in the infinite-horizon problem. We attain this desirable outcome by specifying the cost function recursively, and by surrendering the assumption embedded in previous formulations that the risk-adjusted measure of cost is separable across states of the world.

I. Cost Formulation

Let $\{\mathcal{J}_t : t = 0, 1, \ldots\}$ denote an increasing sequence of information sets (sigma algebras); and $\{w_t : t = 1, 2 \ldots\}$ an m-dimensional sequence of independently and identically normally distributed random vectors with mean zero and covariance matrix I, where w_{t+1} is independent of $\mathcal{J}_t$ and w_t is measurable with respect to $\mathcal{J}_t$. Let $\{x_t : t = 0, 1, \ldots\}$ denote an n-dimensional sequence of state vectors that evolve according to:

$$x_{t+1} = Ax_t + Bu_t + Cw_{t+1}, \quad t = 0, 1, \ldots \ , \tag{1}$$

where x_0 is a given initial vector that can be random but is restricted to be measurable with respect to $\mathcal{J}_0$. In (1), $\{u_t : t = 0, 1, \ldots\}$ is a k-dimensional sequence of control vectors where u_t is restricted to be measurable with respect to $\mathcal{J}_t$. Let $\beta \in (0, 1)$ be a discount factor.

We use the following *risk-adjusted* measure of *cost* for each period $t = 0, \ldots, T-1$:

$$\begin{aligned} \mathcal{C}_{t,T} &= \{u_t' Q u_t + x_t' R x_t - (2\beta/\sigma) \log E[\exp(-\sigma \mathcal{C}_{t+1,T}/2) \mid \mathcal{J}_t]\} \\ \mathcal{C}_{T,T} &= 0 \ . \end{aligned} \tag{2}$$

When $\beta = 1$, the time zero cost $\mathcal{C}_{0,T}$ can be computed using only the time zero conditional expectation operator $E(\cdot|\mathcal{J}_0)$.[1] But when β is strictly less than one, computation of $\mathcal{C}_{0,T}$ with recursion (2) uses the conditional expectation operators $E(\cdot|\mathcal{J}_1), \ldots, E(\cdot|\mathcal{J}_{T-1})$ as well. Therefore, when $\beta < 1$, specification (2) relaxes the assumption of state-separability axiomatized by [10].[2] We shall show that the functional form used in (2) has the features that: (1) the value functions are quadratic functions of the state vector, as in the familiar *optimal linear regulator* problem and in its risk-adjusted version suggested by [5] and [13]; (2) the statistics of the noise process influence the optimal decision rules in a way that depends on the value of σ; and (3) the infinite time horizon problem is well posed and yields a time-invariant optimal linear control law.

II. Cost Recursions and Aggregator Functions

To characterize some properties of our discounted, risk-adjusted costs we follow [7], [8], [9], [3] and use an *aggregator function* α to represent costs recursively. An aggregator function maps hypothetical controls, states, and next period costs into current period costs. Let $\mathcal{U}$ be a space of k-dimensional random control vectors, $\mathcal{X}$ an n-dimensional space of random state vectors, and $\mathcal{L}^+$ the set of nonnegative scalar random variables. For convenience, we let the random variables in $\mathcal{L}^+$ attain the value $+\infty$ for some states of the world. Our aggregator function maps $\mathcal{U} \times \mathcal{X} \times \mathcal{L}^+ \rightarrow \mathcal{L}^+$ and is indexed by a sigma algebra $\mathcal{J}$:

$$\begin{aligned} \alpha(u, x, \Gamma \mid \mathcal{J}) &= u' Q u + x' R x + \beta \rho(\Gamma \mid \mathcal{J}) \\ \rho(\Gamma \mid \mathcal{J}) &\equiv -(2/\sigma) \log E[\exp(-\sigma \Gamma/2) \mid \mathcal{J}]. \end{aligned} \tag{3}$$

[1] Alternatively, the discount factor β could be placed inside the exponential function in (2). This would result in an equivalent risk-adjusted measure of costs when the risk-adjustment parameter σ is replaced by σ/β.

[2] In abandoning state-separability to introduce independent adjustments for risk into objective functions, we are following [8], [3], [11], and [12].

The transformation ρ adjusts next period's costs for risk prior to discounting. If x and u are measurable with respect to $\mathcal{J}$, then $\alpha(u, x, \Gamma \mid \mathcal{J})$ is also. For a hypothetical control u, state x, and next period cost Γ, $\alpha(u, x, \Gamma \mid \mathcal{J})$ is the discounted risk-adjusted current period cost. Notice that (2) can be expressed in terms of α:

$$\mathcal{C}_{t,T} = \alpha(u_t, x_t, \mathcal{C}_{t+1,T} \mid \mathcal{J}_t) \ . \tag{4}$$

We establish some useful properties of the aggregator function. Our first result will be used when we extend the time horizon to infinity.

Lemma 2.1: The function $\rho(\cdot \mid \mathcal{J})$ is monotone increasing in Γ.

Proof: Suppose that $\Gamma_2 \geq \Gamma_1 \geq 0$, then

$$\begin{aligned}\rho(\Gamma_2 \mid \mathcal{J}) - \rho(\Gamma_1 \mid \mathcal{J}) &= -(2/\sigma)\{\log E[\exp(-\sigma\Gamma_2/2) \mid \mathcal{J}] - \log E[\exp(-\sigma\Gamma_1/2) \mid \mathcal{J}]\} \\ &= -(2/\sigma)\log\{E[\exp(-\sigma\Gamma_2/2) \mid \mathcal{J}]/E[\exp(-\sigma\Gamma_1/2) \mid \mathcal{J}]\} \\ &\geq 0 \ . \ \blacksquare\end{aligned}$$

Our next result can be applied to establish that the cost criterion is globally convex in the controls and states when σ is negative.

Lemma 2.2: $\rho(\cdot \mid \mathcal{J})$ is convex in Γ when $\sigma \leq 0$.

Proof: Consider any nonnegative random variables Γ_1 and Γ_2 and their convex combination $\omega\Gamma_1 + (1-\omega)\Gamma_2$ for some $0 < \omega < 1$. Then

$$\begin{aligned}&\log E\{\exp[-(\sigma/2)\omega\Gamma_1 - (\sigma/2)(1-\omega)\Gamma_2] \mid \mathcal{J}\} \\ &\quad = \log E\{[\exp(-\sigma\Gamma_1/2)]^{\omega}[\exp(-\sigma\Gamma_2/2)]^{1-\omega} \mid \mathcal{J}\} \\ &\quad \leq \log(\{E[\exp(-\sigma\Gamma_1/2) \mid \mathcal{J}]\}^{\omega}\{E[\exp(-\sigma\Gamma_2/2) \mid \mathcal{J}]\}^{1-\omega}) \\ &\quad = \omega \log E[\exp(-\sigma\Gamma_1/2) \mid \mathcal{J}] + (1-\omega)\log E\{\exp[-\sigma\Gamma_2/2] \mid \mathcal{J}\}\end{aligned}$$

where the inequality follows from a conditional version of the Hölder Inequality. $\blacksquare$

The aggregator α is convex in u and x because Q is positive definite and R is positive semidefinite. Because α is additively separable in its three arguments (u, x, Γ) and

$\rho(\cdot \mid \mathcal{J})$ is convex in Γ, when σ is negative, α is convex in all three arguments. Because a composition of convex functions is convex, the cost measure $\mathcal{C}_{t,T}$ is convex in the control/state sequence between time t and time T. Increasing $-\sigma$ makes the cost criterion *more* convex and hence risk adjustments more pronounced. For this reason, we are particularly interested in the $\sigma < 0$ case.

III. Infinite Horizon Costs

To formulate the infinite-horizon optimization problem, we first must verify that the associated costs are well defined. This leads us to study the time t cost criterion when the time horizon T increases. We use the result from Lemma 2.1 that the function ρ and hence the aggregator α is monotone increasing in Γ. It follows from this monotonicity that for any control-state sequence $\{(u_t, x_t) : t = 0, 1, \ldots\}, \{\mathcal{C}_{t,T} : T = t, t+1, \ldots\}$ converges almost surely to a limit cost $\mathcal{C}_t$, although this random variable might be infinite $(+\infty)$.

The infinite horizon control problem can be formalized as follows. A *feasible* control process $U = \{u_t : t = 0, 1, ...\}$ is a stochastic process of controls adapted to the sequence of sigma algebras $\{\mathcal{J}_t : t = 0, 1, ...\}$. For any such control process we define recursively via (1) a corresponding state vector process $X = \{x_t : t = 0, 1, ...\}$ adapted to the same sequence of information sets, where we take the initial x_0 as prespecified. Given U and hence X, compute the time zero cost $\mathcal{C}_0$ by evaluating the almost sure limit of $\{\mathcal{C}_{0,T} : T = 1, 2, ...\}$. Define the time zero infinite horizon cost associated with a feasible control process U to be $\mathcal{K}(U|x_0)$. The infinite horizon problem is to minimize $\mathcal{K}(U|x_0)$ by choice of a feasible control process U for each initialization x_0.

For some of our analysis we impose:

ASSUMPTION 1: There exists a matrix F for which the absolute values of all eigenvalues of $A - BF$ are less than $\beta^{-1/2}$.

In light of Lemmas 2.1 and 2.2 and the previous discussion, we have the following:

Theorem 3.1: When σ is negative, $\mathcal{K}(\cdot|x_0)$ is convex. When σ is positive and Assumption 1 is satisfied, there exists a feasible control process U such that $\mathcal{K}(U|x_0)$ is finite.

Proof: We have already argued that the cost criterion is convex in the control process for the finite horizon problem when $\sigma \leq 0$ (see Lemma 2.2 and the ensuing discussion). The convexity of the infinite horizon problem follows by taking almost sure limits when the horizon is extended.

To show that the limiting cost $\mathcal{C}_t$ can be made finite for some control process when σ is positive, we use the $\sigma = 0$ case as a benchmark. The aggregator when $\sigma = 0$ is given by $\alpha^*(u, x, \mathcal{C}|\mathcal{J}) = u'Qu + x'Rx + \beta E(\mathcal{C}|\mathcal{J})$. By the convexity of the exponential function, $E[\exp(-\sigma\mathcal{C}/2)|\mathcal{J}] \geq \exp[-(\sigma/2)E(\mathcal{C}|\mathcal{J})]$. Taking logarithms and multiplying by $(-2\beta/\sigma)$ it follows that

$$\alpha(u, x, \mathcal{C}|\mathcal{J}) \leq \alpha^*(u, x, \mathcal{C}|\mathcal{J}) \quad \text{for } \sigma \geq 0. \tag{5}$$

Given inequalities (5) and the monotonicity of α and α^* in $\mathcal{C}$, it can be established that

$$\mathcal{C}_t \leq \mathcal{C}_t^* \quad \text{for } \sigma \geq 0, \tag{6}$$

where $\mathcal{C}_t^*$ is the time t, $\sigma = 0$ cost. It follows from (6) that for $\sigma > 0$, the 'adjusted' cost will be finite whenever the 'unadjusted' (or *) cost is finite. In light of Assumption 1, there exists a time invariant control law that makes the 'unadjusted' cost and hence the adjusted cost finite. ∎

Unfortunately, the inequalities in the proof of Theorem 3.1 work in the wrong direction for the case in which $\sigma < 0$. In that case, the 'risk adjusted' cost can be infinite when the 'unadjusted' cost is finite. The finding that the cost criterion can be infinite when σ is strictly negative closely parallels results in [5].

IV. Arbitrary Time-invariant Linear Control Laws

As a prelude to showing that the solution to the infinite horizon optimal control problem is linear and time invariant, we characterize the cost associated with any such law. Suppose that $u_t = -Fx_t$. We can compute our cost criterion under this control law by doing recursions on the aggregator function α evaluated at the linear control and a quadratic representation of next period's costs. Let V be a positive semidefinite matrix and c be a nonnegative real number. Note that

$$\alpha(-Fx, x, y'Vy + c \mid \mathcal{J}) = x'R^*x + \beta\rho(y'Vy + c \mid \mathcal{J}) \tag{7}$$

where

$$y = A^*x + Cw \ , \tag{8}$$

$A^* \equiv A - BF, R^* \equiv F'QF + R$, and w is normally distributed conditioned on $\mathcal{J}$ with mean zero and covariance matrix I. We use the following formula from [5] to compute the right side of (7):[3]

$$\begin{aligned}\rho(y'Vy + c \mid \mathcal{J}) =& x'A^{*\prime}[V - \sigma VC(I + \sigma C'VC)^{-1}C'V]A^*x \\ &+ c + (1/\sigma)\log\det(I + \sigma C'VC).\end{aligned} \tag{9}$$

This formula only works when $(I + \sigma C'VC)$ is positive definite. If $(I + \sigma C'VC)$ has one nonpositive eigenvalue, then the left side of (7) is infinite (where we interpret $\log(+\infty) = +\infty$). Substituting (9) into (7), we obtain

$$\alpha(-Fx, x, y'Vy + c \mid \mathcal{J}) = x'\mathcal{S}(V)x + \mathcal{U}(V, c) \tag{10}$$

where

$$\mathcal{S}(V) \equiv R^* + \beta A^{*\prime}[V - \sigma VC(I + \sigma C'VC)^{-1}C'V]A^* \tag{11}$$

and

$$\mathcal{U}(V, c) \equiv \beta c + (\beta/\sigma)[\log\det(I + \sigma C'VC)]. \tag{12}$$

Consequently with the Gaussian information structure, the aggregator α maps a translated quadratic cost measure for next period's costs into a translated quadratic cost measure today. Since the infinite-horizon cost is computed by iterating on α as in (4), the resulting initial period cost measure, when finite, is quadratic in the initial state vector plus a constant, say, $x_0'V^*x_0 + c^*$.

[3] This is a special case of formulas (33) and (38) in [5], where P_k in Jacobson's formulas is taken to be the identity matrix.

Theorem 4.1: Suppose that $u_t = -Fx_t$ for $t = 0, 1, ...$ and that the corresponding time zero cost $\mathcal{C}_0$ is finite. Then $\mathcal{C}_0 = x_0' V^* x_0 + c^*$ for some positive semidefinite matrix V^* and some nonnegative scalar c^* that are the smallest solutions to:

$$V^* = \mathcal{S}(V^*), \qquad c^* = \mathcal{U}(V^*, c^*). \tag{13}$$

Proof: Recall that $\mathcal{C}_0$ is defined to be the almost sure limit point of the monotone sequence $\{\mathcal{C}_0, j : j = 1, 2, ...\}$. Let $\mathcal{S}^j(V)$ denote the jth iterate of the transformation $\mathcal{S}$ evaluated at a matrix V. Similarly, define

$$\mathcal{U}^j(V, c) \equiv \mathcal{U}[\mathcal{S}^{j-1}(V), \mathcal{U}^{j-1}(V, c)]. \tag{14}$$

Then by applying formula (10) repeatedly, it follows that

$$\mathcal{C}_{0,j} = x_0' \mathcal{S}^j(0) x_0 + \mathcal{U}^j(0, 0)$$

where we have imposed the terminal cost restriction $\mathcal{C}_{j,j} = 0$. Since the cost sequence is finite by assumption and converges almost surely for any initialization x_0, it must be that the sequence of positive semidefinite matrices $\{\mathcal{S}^j(0) : j = 1, 2, ...\}$ converges (entry by entry) as does the scalar sequence $\{\mathcal{U}^j(0, 0) : j = 1, 2, ...\}$ of nonnegative numbers. Let V^* and c^* denote the respective limit points. From (12) and (14), it follows that $\mathcal{U}^j(0, 0) \geq (\beta/\sigma)[\log\det(I + \sigma C' \mathcal{S}^{j-1}(0) C)]$. Taking limits it follows that $c^* \geq (\beta/\sigma)[\log\det(I + \sigma C' V^* C)]$, and hence $(I + \sigma C' V^* C)$ is nonsingular. It follows that the operator $\mathcal{S}$ is continuous, which in turn implies that V^* is a fixed point of $\mathcal{S}$. The convergence of $\{\mathcal{U}^j(0, 0)\}$ to c^* together with the continuity of $\mathcal{U}$ at (V^*, c^*) ensures that $c^* = \mathcal{U}(V^*, c^*)$. Therefore, (V^*, c^*) solves (13).

Let $(\hat{V}, \hat{c})$ be any other solution of (13) where $\hat{V}$ is a positive semidefinite matrix and $\hat{c}$ is a nonnegative real number. By the monotonicity of α in its third argument, it follows that $x_0' \hat{V} x_0 + \hat{c} \geq x_0' \mathcal{S}^j(0) x_0 + \mathcal{U}^j(0, 0)$. Taking limits of the right side as $j \to \infty$ it follows that $x_0' \hat{V} x_0 + \hat{c} \geq x_0' V^* x_0 + c^*$. Since this inequality holds for any initialization of x_0, $\hat{V} \geq V^*, \quad \hat{c} \geq c^*$. The inequality involving the V's can be established by making x_0

large in a several different directions, while the inequality entailing the c's can be verified by setting x_0 to zero. ∎

REMARK: Since the $\mathcal{S}$ transformation has the form of the operator associated with a matrix Riccati equation, we can construct an optimal control problem which has the smallest fixed point of $\mathcal{S}$ for its solution. Associated with this problem are discounted, infinite-horizon versions of notions of 'pessimism' and 'optimism' in [5] and [13]. See [4] for details.

V. Solution to the Infinite Horizon Discounted Problem

We solve the infinite-horizon problem by first deriving discounted versions of the recursions obtained by [5]. Consider the two-period problem:

$$\min_{u,y} \alpha(u, x, y'Wy + d \mid \mathcal{J}) \quad \text{subject to} \quad y = Ax + Bu + Cw \tag{15}$$

where w is an m-dimensional normally distributed random vector with mean zero and covariance matrix I that is independent of the state vector x, and where $(I + \sigma C'WC)$ is positive definite. Using the previously mentioned formula from [5], we are led to solve a corresponding unconstrained quadratic optimization problem:

$$\min_u \{u'Qu + x'Rx + \beta(Ax + Bu)'\mathcal{D}(W)(Ax + Bu) + \mathcal{U}(W, d)\} \tag{16}$$

where

$$\mathcal{D}(W) \equiv W - \sigma WC(I + \sigma C'WC)^{-1}C'W. \tag{17}$$

The optimal control law is

$$u = -\mathcal{F} \circ \mathcal{D}(W)x, \quad \mathcal{F}(V) = \beta[Q + \beta B'VB]^{-1}B'VA, \tag{18}$$

and the minimized value of the criterion is $x'\mathcal{T} \circ \mathcal{D}(W)x + \mathcal{U}(W, d)$, where

$$\begin{aligned} \mathcal{T}(V) &\equiv R + \mathcal{F}(V)'Q\mathcal{F}(V) + \beta[A - B\mathcal{F}(V)]'V[A - B\mathcal{F}(V)] \\ &= R + A'(\beta V - \beta^2 VB(Q + \beta B'VB)^{-1}B'V)A. \end{aligned} \tag{19}$$

Consequently, if next period's value function is quadratic in the state plus a constant term, the current period value function will have the same functional form. When $\sigma = 0, \mathcal{D}$ is the identity operator so that $\mathcal{T} \circ \mathcal{D}$ is just $\mathcal{T}$. Note that the $\mathcal{T}$ is the operator associated with the matrix Riccati equation for the ordinary discounted version of the optimal linear regulator problem.

Following the usual backward induction argument, iterating on the composite transformation $\mathcal{T} \circ \mathcal{D}$ corresponds to increasing the time horizon. The minimized value $\mathcal{M}_{0,j}$ of the time 0, j-period optimization problem is

$$\mathcal{M}_{0,j} = x_0'(\mathcal{T} \circ \mathcal{D})^j(0)x_0 + \mathcal{V}^j(0,0) \tag{20}$$

where

$$\mathcal{V}^j(W,d) = \mathcal{U}[(\mathcal{T} \circ \mathcal{D})^{j-1}, \mathcal{V}^{j-1}(W,d)] \tag{21}$$

The infinite-horizon problem involves the limiting behavior of the sequence $\{\mathcal{M}_{0,j} : j = 1,2,...\}$. In light of Lemma 2.1, this sequence of minimized values is increasing and hence converges (almost surely) to a limit value $\mathcal{M}_0$, although the limit might be infinite.

Theorem 5.1: Suppose that $\mathcal{M}_0$ is finite. Then

$$\mathcal{M}_0 = x_0'W^*x_0 + d^*$$

for some positive semidefinite matrix W^* and some nonnegative scalar d^* that are the smallest solutions to: $W^* = \mathcal{T} \circ \mathcal{D}(W^*), \quad d^* = \mathcal{V}(d^*)$.

Proof: This theorem can be proved by imitating the proof of Theorem 4.1 where $\mathcal{T} \circ \mathcal{D}$ takes the place of $\mathcal{S}$ and $\mathcal{V}$ takes the place of $\mathcal{U}$. ∎

Theorem 5.2: Suppose that $\mathcal{M}_0$ is finite. Then the process U associated with the linear control law $u_t = -\mathcal{F} \circ \mathcal{D}(W^*)x_t$ minimizes $\mathcal{K}(U|x_0)$ by choice of a feasible control process U for every initialization x_0.

Proof: First note that

$$x_0'W^*x_0 + d^* \leq \mathcal{K}(U|x_0). \tag{22}$$

Since W^* is a fixed point of $\mathcal{T} \circ \mathcal{D}$, it is also a fixed point for the transformation $\mathcal{T}$ constructed using the candidate time-invariant linear control law $[F = \mathcal{F} \circ \mathcal{D}(W^*)]$. Moreover, $d^* = \mathcal{U}(W^*, d^*)$. Let (V^*, d^*) be the positive semidefinite matrix and nonnegative scalar given by Theorem 4.1 for the candidate control law. Then $V^* \leq W^*$, $c^* \leq d^*$. In light of relation (22), these inequalities can be replaced by equalities, and the conclusion follows. ∎

Both Theorems 5.1 and 5.2 presume that $\mathcal{M}_0$ is finite. From Theorem 3.1 we know that $\mathcal{M}_0$ is finite whenever Assumption 1 is satisfied and σ is positive. When σ is strictly negative, Assumption 1 is no longer sufficient. However, $\mathcal{M}_0$ can be computed as the limit of $\{\mathcal{M}_{0,j} : j = 1, 2, ...\}$ as in formulas (20) and (21). As long as the matrices $\{(I + \sigma C'(\mathcal{T} \circ \mathcal{D})^j C) : j = 1, 2, ...\}$ are positive definite and have a positive definite limit, $\mathcal{M}_0$ will be finite.

It is also of interest to know when the optimal control law stabilizes the state vector process (relative to β). As in the usual linear-quadratic ($\sigma = 0$) problem, factor the positive semidefinite matrix $R = L'L$.

Theorem 5.3: Suppose that σ is negative, $\mathcal{M}_0$ is finite and $(\beta^{1/2}A, L)$ is detectable. Then the absolute values of the eigenvalues of $A - B\mathcal{F}(W^*)$ are less than $\beta^{-1/2}$.

Proof: As in the proof of Theorem 3.1, compare the risk adjusted costs with $\sigma < 0$ to the $\sigma = 0$ costs. The counterpart to inequalities (5) and (6) for $\sigma \leq 0$ are

$$\alpha(u, x, \mathcal{C}|\mathcal{J}) \geq \alpha^*(u, x, \mathcal{C}|\mathcal{J}) \quad \text{for} \quad \sigma \leq 0$$

and

$$\mathcal{C}_0 \geq \mathcal{C}_0^* \quad \text{for} \quad \sigma \leq 0. \tag{23}$$

Apply these inequalities for the control law $u_t = -\mathcal{F}(W^*)x_t$. Since the $(\beta^{1/2}A, L)$ is detectable, the unadjusted $\mathcal{C}_0^*$ cost will be finite only if $A - B\mathcal{F}(W^*)$ satisfies the eigenvalue restriction. Given inequality (23), $A - B\mathcal{F}(W^*)$ must satisfy the eigenvalue restriction because $\mathcal{C}_0$ and hence $\mathcal{C}_0^*$ is finite. ∎

Remark: Modifications of standard computational methods can be used to accelerate computational speed relative to iterating on the transformation $\mathcal{T} \circ \mathcal{D}$. See [4] for details.

VI. Summary

In economic applications, it is very useful to have formulations of dynamic programming problems that incorporate discounting, imply time-invariant decision rules for infinite horizon-problems, and are readily calculable.[4] For that reason, the discounted optimal linear regulator has been a workhorse in economic dynamics generally, and in particular in the development of real business cycle theory and the formulation of linear rational expectations models as applied in real business cycle theory and finance. The modified linear regulator described in this paper preserves the desirable features of the linear regulator, and adds aspects of risk-sensitivity.

References

[1] B.D.O. Anderson and J. B. Moore (1979). *Optimal Filtering*. Prentice Hall, New Jersey.

[2] Bouakiz and Sobel (1984). 'Nonstationary policies are optimal for risk-sensitive Markov decision processes'. *Report, College of Management*, Georgia Institute of Technology.

[3] Epstein, L.G., and S.E. Zin (1989). 'Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework'. *Econometrica*, 57:4, 937–969.

[4] L.P. Hansen and T. J. Sargent (1994). *Recursive Models of Dynamic Linear Economies*. Princeton University Press, forthcoming.

[4] Our recursive cost specification preserves the 'time consistency' property that is satisfied in the time and state-separable $\sigma = 0$ specification of the problem.

[5] D.H. Jacobson (1973). 'Optimal Stochastic Linear Systems with Exponential Performance Criteria and Their Relation to Deterministic Differential Games'. *IEEE Transactions on Automatic Control*, Vol. $AC - 18$, No. 2, 124–131.

[6] D.H. Jacobson (1977). *Extensions of Linear-Quadratic Control, Optimization and Matrix Theory*. Academic Press, New York.

[7] T.C. Koopmans (1960). 'Stationary Ordinal Utility and Impatience'. *Econometrica*, 28, 287–309.

[8] D.M. Kreps and E.L. Porteus (1978). 'Temporal Resolution of Uncertainty and Dynamic Choice Theory'. *Econometrica*, 46, 185–200.

[9] R.E. Lucas, Jr., and N.L. Stokey (1984). 'Optimal Growth with Many Consumers'. *Journal of Economic Theory*, 32, 139–171.

[10] J. Von Neumann and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton, N.J.: Princeton University Press.

[11] P. Weil (1990). 'Nonexpected Utility in Macroeconomics'. *Quarterly Journal of Economics*, Vol. 105, pp. 29–42.

[12] P. Weil (1993). 'Precautionary Savings and the Permanent Income Hypothesis'. *Review of Economic Studies*, Vol. 60, 367–383.

[13] P. Whittle (1981). 'Risk-sensitive Linear/Quadratic/Gaussian Control'. *Advances in Applied Probability*, 13, 764–777.

[14] P. Whittle (1989). 'Entropy-minimizing and Risk-sensitive Control Rules'. *System and Control Letters*, 13:4, 1–7.

[15] P. Whittle (1990). *Risk-sensitive Optimal Control*. John Wiley & Sons, New York.