

Bellman's Equation and a Continuous Linear Programming Problem

by

W.R.S. Sutherland

Abstract: Bellman's optimality equation of dynamic programming is examined in the context of a discrete-time, continuous-state economic development model. The main focus of the paper is on the interpretation of this functional equation as a linear programming problem in an infinite-dimensional setting. The connection between this linear programming problem and Bellman's functional equation is developed using a theory of equivalent models. The discrete-state version of the problem is discussed by usual theory for finite-dimensional linear programming. Then the abstract theory for infinite-dimensional linear programming is applied to the continuous-state problem in order to obtain results on existence and strong duality. The paper concludes with several simple examples of the dual pair of continuous linear programming problems.

Mathematics Subject Classification Numbers: 90C48, 90C39, 49C15.

1. Introduction:

This paper examines the connection between Bellman's "optimality equation" of continuous-state space dynamic programming and the theory of linear programming problems in an infinite-dimensional setting. Our discussion will be based on a simple discrete-time economic development model. Bellman's functional equation has a particularly simple form for this model.

The economic development model is constructed as follows. Let I be a subset of n -dimensional space representing the set of possible states of some economy at some point in time. For each $x \in I$, let $S(x)$ denote the subset of states in I which can be attained in the following time period. The set of all feasible transitions in one time period is $T = \{(x,y) \mid y \in S(x), x \in I\}$. It will be convenient to assume that $(x,x) \in T$ for each $x \in I$. Finally assigning a cost $c(x,y)$ to each feasible transition (x,y) , and letting $0 < \delta < 1$ represent a fixed discount factor, we obtain the following statement of the economic development model:

Given an initial state x_0 , find a sequence of feasible transitions $(x_t, x_{t+1}) \in T$ for $t = 0, 1, \dots$ such that the value $\alpha(x_0) = \text{infimum}\{\sum_{t=0}^{\infty} \delta^t c(x_t, x_{t+1})\}$ is attained over the set of all feasible sequences from x_0 .

The existence of an optimal sequence, for each initial state $x_0 \in I$, is immediate under the usual assumptions that T is a compact set and that $c(x,y)$ is a continuous function on T . Bellman's "principle of optimality", in the context of this economic development model, asserts that the optimal value function $\alpha(x)$ is the unique solution of the functional equation:

$$\alpha(x) = \text{minimum}\{c(x,y) + \delta\alpha(y) \mid y \in S(x)\} \text{ for all } x \in I.$$

The existence of a unique continuous solution for Bellman's functional equation was originally proven by Radner [5] for this economic development model. It follows that the optimal sequences are generated by the recursion $x_{t+1} = p(x_t)$ for $t = 0, 1, \dots$, where the (time-independent) policy function p satisfies the equation $\alpha(x) = c(x,p(x)) + \delta\alpha(p(x))$ for all $x \in I$. The functions α and p can be computed, for any specific model, by using the well-known value improvement or policy improvement algorithms of dynamic programming.

Bellman's functional equation leads naturally to the consideration of those functions $\alpha(x)$ which satisfy the inequality $\alpha(x) \leq c(x,y) + \delta\alpha(y)$ for all $(x,y) \in T$. If we next replace the minimization operator in Bellman's equation by the requirement that the function $\alpha(x)$ be a maximal element (subject to the above inequality) in the space $C(I)$ of continuous real-valued functions on I , we obtain the following continuous linear programming problem:

Find a function $\alpha(x)$ which maximizes $\int_I \alpha(x) dx$ subject to
the constraints $\alpha(x) \leq c(x,y) + \delta\alpha(y)$ for all $(x,y) \in T$.

This continuous linear programming problem is the focal point of this paper. Although this problem is linear with respect to the function $\alpha(x)$, it should be noted that the inequality constraints are not in the usual Volterra equation form that is typical of most of the literature on continuous linear programming problems (see Anderson and Nash [1] for a recent exposition).

This paper is organized as follows. Section 2 considers economic development models, with different transition cost functions, that are equivalent in the sense that they share the same optimal sequences. This approach is used to provide an elementary proof that solutions of the continuous linear programming problem are solutions of Bellman's functional equation, and vice-versa. Section 3 considers a discretized version of the economic development model to which the usual finite-dimensional duality theory of linear programming can be applied. Section 4 then applies the duality theory of abstract linear programming to the infinite-dimensional continuous linear programming problem. The main result of the paper consists of showing that a strong duality result holds for this problem (although generalized functions are needed to attain the optimal value in the dual program). Three examples are presented in Section 5.

2. Equivalent Models:

This section describes the class of economic development models which have the same optimal sequences as the original model. This discussion also provides some additional motivation for the consideration of the continuous linear programming problem. In particular, it is shown how the solutions of Bellman's functional equation can be obtained from the linear programming problem.

Let p be a feasible (time-dependent) policy function. That is, $p_t(x) \in S(x)$ for all $x \in I$ and for all $t = 0, 1, \dots$. Given an initial state $x_0 \in I$, this policy generates the sequence $(x_t, y_t) = (x_t, x_{t+1}) = (x_t, p_t(x_t))$ of feasible transitions. The total discounted cost for this sequence will be denoted by the value $\alpha(c, p)(x_0)$, where

$$\alpha(c, p)(x_0) = \sum_{t=0}^{\infty} \delta^t c(x_t, p_t(x_t)).$$

The standard approach to finding a minimal total cost sequence for the economic development model would be to consider variations in the policy functions $p_t(x)$. However we will now pursue the alternate approach of considering variations in the transition cost function $c(x, y)$. The basic idea is that the economic development model can be solved more easily for some cost functions than for others. In particular, a model (c) is said to be trivial if (1) $c(x, y) \geq 0$ holds for all $(x, y) \in T$, and (2) for each $x \in I$ there exists some $y = \pi(x) \in S(x)$ such that $c(x, \pi(x)) = 0$. Clearly, if model (c) is trivial, then $y = \pi(x)$ is an optimal policy function, for every $x_0 \in I$. Consequently, our original model can be solved if its transition cost function can be transformed into one which is trivial without affecting the optimality of the sequences for the original model. The following definition (see [6] for other applications) leads to an appropriate restriction on the class of transformed models.

Definition: Two models (c) and (d) are said to be equivalent if the difference function $\alpha(c, p)(x_0) - \alpha(d, p)(x_0)$ is independent of the choice of the feasible policy p .

The usefulness of this definition depends on the following two lemmas.

Lemma 2.1: Equivalent models have the same optimal policy functions.

Proof: Let p, q be feasible policy functions. If models (c) and (d) are equivalent, then

$$\alpha(c, p)(x_0) - \alpha(d, p)(x_0) = \alpha(c, q)(x_0) - \alpha(d, q)(x_0)$$

holds for each $x_0 \in I$. Since this equation can be rearranged into the equation

$$\alpha(c, p)(x_0) - \alpha(c, q)(x_0) = \alpha(d, p)(x_0) - \alpha(d, q)(x_0),$$

it follows that p is optimal for model (c) if and only if it is optimal for model (d) . □

The next result provides an explicit algebraic restriction on the class of transformed models.

Lemma 2.2: Models (c) and (d) are equivalent if and only if there exists a continuous function $\alpha(x)$ such that $c(x,y) = d(x,y) + \alpha(x) - \delta\alpha(y)$ holds for all $(x,y) \in T$.

Proof: Suppose that there exists a continuous function $\alpha(x)$ satisfying this identity. Then, for any feasible policy function p , the difference function can be reduced as follows:

$$\begin{aligned} \alpha(c,p)(x_0) - \alpha(d,p)(x_0) &= \sum_{t=0}^{\infty} \delta^t \{c(x_t, x_{t+1}) - d(x_t, x_{t+1})\} \\ &= \sum_{t=0}^{\infty} \delta^t \{\alpha(x_t) - \delta\alpha(x_{t+1})\} = \lim_{\tau \rightarrow \infty} \sum_{t=0}^{\tau} \delta^t \{\alpha(x_t) - \delta\alpha(x_{t+1})\} \\ &= \lim_{\tau \rightarrow \infty} \{\alpha(x_0) - \delta^{\tau+1}\alpha(x_{\tau+1})\} = \alpha(x_0). \end{aligned}$$

Since this difference function is independent of the choice of the policy function p then models (c) and (d) are equivalent. Conversely, let (c) and (d) be equivalent models. Then the equation

$$\alpha(c,p)(x_0) - \alpha(d,p)(x_0) = \alpha(c,q)(x_0) - \alpha(d,q)(x_0)$$

holds for any feasible policies p and q for $x_0 \in I$. Consider the policies p and q which yield the sequences $\{x, x, x, \dots\}$ and $\{x, y, y, \dots\}$ respectively, from any $x = x_0 \in I$ where $(x,y) \in T$. Then, using the continuous function $\alpha(x) = (c(x,x) - d(x,x))/(1-\delta)$, this equation reduces to the identity in the statement of the lemma. \square

This theory of equivalent models can be used to show that Bellman's equation and the continuous linear programming problem have the same solutions.

Theorem 2.3: A continuous function α satisfies Bellman's functional equation if and only if it is an optimal solution of the continuous linear programming problem.

Proof: Suppose that $\alpha(x)$ is feasible for the constraints of the linear programming problem. Then $\alpha(x)$ defines a model (d) which is equivalent to the original model (c). Furthermore, condition (1) for a trivial model also holds for (d). If $\alpha(x)$ is an optimal solution then condition (2) holds and it follows that $\alpha(x)$ solves Bellman's equation. Conversely, using the identity of Lemma 2.2, any solution $\alpha(x)$ of Bellman's equation yields an equivalent trivial model (d). \square

3. The Discrete Problem:

This section considers the special case of economic development models with discrete-state spaces, such as $I = \{1, 2, \dots, n\}$. The set of feasible transitions is $T = \{(i,j) \mid j \in S(i), i \in I\}$. Recall that we assume that $(i,i) \in T$ for all $i \in I$. If $j = p_t(i)$ is a feasible (time-dependent) policy function, then the total discounted cost $\alpha(c,p)(i_0)$ from an initial state i_0 is given by

$$\alpha(c,p)(i_0) = \sum_{t=0}^{\infty} \delta^t c_{i_t, i_{t+1}}.$$

The equivalent model approach of the previous section requires little modification for the discrete-state space problem. A necessary and sufficient condition for two models (c) and (d) to be equivalent is that there exists an n -vector α satisfying the equation

$$c_{i,j} = d_{i,j} + \alpha_i - \delta\alpha_j \text{ for all } (i,j) \in T.$$

The continuous linear programming problem is a finite-dimensional problem (originally stated by D'Epenoux [3]):

$$\begin{aligned} &\text{Find } \alpha_i \text{ which maximize } \sum_{i \in I} \alpha_i \text{ subject to the constraints} \\ &\alpha_i - \delta\alpha_j \leq c_{i,j} \text{ for all } (i,j) \in T. \end{aligned}$$

An equivalent model (d) is defined by the slack variables $d_{i,j} \geq 0$ for these constraints. The dual linear program is given by the problem:

$$\begin{aligned} &\text{Find } \beta_{i,j} \geq 0 \text{ which minimize } \sum_{i \in I} \sum_{j \in S(i)} c_{i,j} \beta_{i,j}; \text{ subject to the constraints} \\ &\sum_{j \in S(k)} \beta_{k,j} - \delta \sum_{i \in S^{-1}(k)} \beta_{i,k} = 1 \text{ for } k = 1, 2, \dots, n. \end{aligned}$$

The duality theory of linear programming yields a strong result for the discrete problem.

Lemma 3.1: The discrete primal and dual programs have optimal solutions of equal value.

Proof: Since these linear programs are finite-dimensional then it is sufficient to show that each of the dual pair of programs is feasible. In the primal program, let M be some constant such that $c_{i,j} \geq M$ and set $\alpha_i = M/(1 - \delta)$. In the dual program, set $\beta_{i,j} = 1/(1 - \delta)$ for $i = j$ and $\beta_{i,j} = 0$ otherwise. The feasibility of these solutions is easily verified. \square

Lemma 3.2: Let (c) be a discrete model, then there exists an equivalent trivial model (d).

Proof: Let α, β be an optimal pair of solutions. Let model (d) be defined by the slack variables of the primal solution, then (d) is equivalent to (c). Since $d_{i,j} \geq 0$ then condition (1) for a trivial model holds. Assume that condition (2) fails to hold for some $k \in I$, that is $d_{k,j} > 0$ holds for all $j \in S(k)$. Then by the complementary slackness condition $d_{i,j}\beta_{i,j} = 0$ for all $(i,j) \in T$ it follows that $\beta_{k,j} = 0$ for all $j \in S(k)$. But this would imply that β fails to satisfy the k^{th} constraint of the dual program. Since β is feasible then condition (2) must hold and (d) is a trivial model. \square

In summary, the existence of optimal solutions for the dual pair of linear programs of the discrete development model is a direct consequence of standard linear programming theory. The simplex algorithm of linear programming can be seen, in terms of the equivalent model approach, as transforming the original primal model (c) through a sequence of equivalent models until a trivial model is obtained. This process is analogous to the way in which Kuhn's "Hungarian Method" solves the optimal assignment problem.

4. The Continuous Problem:

We now turn to the continuous-state space model. In order to discuss the duality theory for this infinite-dimensional problem it is necessary to introduce two sets of paired ordered topological vector spaces. The first set of paired spaces consists of the space $C(I)$ of continuous functions on the compact set I and its dual space $M(I)$, the regular Borel measures on I . A bilinear form $\langle \cdot, \cdot \rangle$ on these paired spaces is defined by $\langle \alpha, \nu \rangle = \int_I \alpha(x) d\nu(x)$. The positive cone P in $C(I)$ is defined to be the whole space; the dual (negative polar) cone is simply $P^* = \{0\}$. The second set of paired spaces consists of the space $C(T)$ of continuous functions on the compact set T and its dual space $M(T)$, the regular Borel measures on T . The bilinear form $\langle \cdot, \cdot \rangle$ on these paired spaces is defined by $\langle \beta, \mu \rangle = \int_T \beta(x,y) d\mu(x,y)$. The positive cone Q of $C(T)$ is defined to be the non-negative continuous functions on T ; the positive cone Q^* in $M(T)$ is the set of non-negative Borel measures on T . These bilinear forms will be continuous if the sup-norm topology is used in both function spaces.

The continuous linear programming problem of Section 1 can now be expressed in the following form (where λ denotes Lebesgue measure on I):

Find $\alpha(x)$ in $C(I)$ which maximizes $\langle \alpha, \lambda \rangle$
subject to the constraint $A\alpha + c \geq 0$.

The map $A: C(I) \rightarrow C(T)$ defined by $(A\alpha)(x,y) = -\alpha(x) + \delta\alpha(y)$ for $(x,y) \in T$ is a continuous linear map on these spaces. Its adjoint map $A^*: M(T) \rightarrow M(I)$ is implicitly defined by the identity $\langle A\alpha, \nu \rangle = \langle \alpha, A^*\nu \rangle$. A more explicit formula can be derived as follows. Let P_x and P_y denote the projections of a measure ν in $M(T)$ into $M(I)$ where we assume that the measure ν is extended to all of $I \times I$ by setting $\nu = 0$ on the complement of the set T . Then the adjoint map can be represented as $A^*\nu = -P_x(\nu) + \delta P_y(\nu)$. The dual linear program takes the form:

Find $\beta \in M(T)$ which minimizes $\langle c, \beta \rangle$
subject to the constraints $A^*\beta + \lambda = 0$ and $\beta \geq 0$.

The next result shows that the optimum values are attained in each of the above linear programs.

Theorem 4.1: Each of the primal and dual programs has an optimal solution.

Proof: The attainment of the maximum value for the primal program can be shown (e.g., Denardo [4]) by noting that Bellman's function equation is a contraction operator on $C(I)$. The attainment of the minimum value for the dual program follows from Alaoglu's theorem since the continuous function $c(x,y)$ can be assumed to be uniformly positive on T (by the addition of a constant) and the constraints of the dual are feasible for $\beta = \lambda/(1-\delta)$ as $(x,x) \in T$ for all $x \in I$. \square

The values of feasible solutions to the dual pair of linear programs always satisfy a weak duality inequality of the form $\langle \alpha, \lambda \rangle \leq \langle c, \beta \rangle$. The lack of a duality gap for this dual pair of continuous linear programs is assured by the following strong duality result.

Theorem 4.2: There exist optimal solutions α, β of the dual pair of linear programs such that $\langle \alpha, \lambda \rangle = \langle c, \beta \rangle$.

Proof: Let N be any strict lower bound for $c(x, y)$ on T . Since $A\alpha + c > 0$ holds for the function $\alpha(x) = N/(1-\delta)$, then $A\alpha + c$ belongs to the Mackey interior of the positive cone Q . The conclusion of the theorem then follows from Theorem 3.13 in [1, p. 55] and Theorem 4.1 above. \square

The optimal solutions for the dual pair of linear programs satisfy the complementary slackness condition $\langle A\alpha + c, \beta \rangle = 0$. Suppose that $y = \pi(x)$ is the optimal policy function determined by α and that the primal constraint is a strict inequality for $y \neq \pi(x)$, then the optimal β for the dual program must have all of its mass concentrated on the subset $\{(x, \pi(x)) \mid x \in I\}$ of T . Since this subset has Lebesgue measure zero then it appears that the dual program for the continuous development model is less attractive for computation than is the case for a discretized version of the same model (as in Section 3). The computation of optimal policies is illustrated in the next section.

5. Some Examples:

This section presents three simple examples of continuous-state economic development models in which the state space I is a subset of the real line. In each example, the optimal solution α of the primal program can be explicitly stated in terms of elementary functions, but the corresponding optimal solution β of the dual can only be represented by a generalized function.

For the first example let $c(x, y) = mx + ny$ on $T = \{(x, y) \mid 0 \leq x, y \leq 1\}$. For simplicity we will consider only the case $m < 0$ and $n > 0$. If $m + \delta n > 0$ then the optimal solution of the primal program is given by $\alpha(x) = mx$ with the optimal policy function being $y = \pi(x) = 0$. A trivial equivalent model is given by the cost function $d(x, y) = (\delta m + n)y$. (If $m + \delta n < 0$ then $\alpha(x) = mx + (m + \delta n)/(1-\delta)$ and $y = \pi(x) = 1$). The optimal solution for the dual program is approximated by solving a discretized version in which the interval $I = [0, 1]$ is subdivided into k equal subintervals. Note that the complementary slackness condition requires that $\beta_{i,j} = 0$ for all i , all $j \neq 1$. The dual constraints then can be solved to yield $\beta_{1,1} = k(1 + (k-1)\delta)/(1-\delta)$ and $\beta_{i,1} = k$ for $i > 1$. The value of the dual objective function can be shown to converge to the value $\int_I \alpha(x) dx = m/2$ of the primal program as $k \rightarrow \infty$. This illustrates the lack of a duality gap

in this example. This construction also shows that the optimal solution of the dual linear program will not be given by a completely continuous measure (i.e., by an integrable function on T).

The second example is similar to the first except that the cost $c(x,y) = ax^2 + by^2$ is a convex quadratic function on the unit square. The optimal solution of the primal linear program is given by $\alpha(x) = ax^2$ where the optimal policy is again given by $\pi(x) = 0$ for all $x \in I$. A trivial equivalent model is defined by the cost function $d(x,y) = (b + \delta a)y^2$. The optimal value of the primal program, $\int_I \alpha(x) dx = a/6$, is again approached by the value of the feasible solutions β_{i_j} for the discretized dual program as $k \rightarrow \infty$, where β is the same as in the first example.

The third example (due to Tilquin [7,p.44]) involves a typical economic development model. Let $T = \{(x,y) \mid 0 < y < 3x^{1/2}, 0 < x < 9\}$, $\delta = 2/3$, and $c(x,y) = -\ln(3x^{1/2} - y)$. The solution of Bellman's equation is given by the optimal value function $\alpha(x) = -3 \ln(2x^{1/4})$ and the optimal policy function is $y = \pi(x) = x^{1/2}$. An equivalent trivial model is given by the transition cost function $d(x,y) = \ln(2x^{3/4} / (y^{1/2}(3x^{1/2} - y)))$. The value of the dual program is $\int_I \alpha(x) dx = 9(3 - 4\ln(2) - 6\ln(3))/4$. This value is attained by the feasible solution of the dual program represented by the generalized function $\beta(x,y)$ which equals zero if $y \neq x^{1/2}$, and otherwise is given by $\beta(x,y) = 3$ for $0 < x \leq 1$, and $\beta(x,y) = 3(1 - (2/3)^{n+1})$ where n and x satisfy the inequalities $n-1 \leq -\ln(x) / \ln(2) < n$. The necessity of using generalized functions in the dual program was anticipated by Bellman [2, p. 211] in his discussion of the "bottleneck problem" of dynamic programming.

References:

- [1] Anderson, E.J. and P. Nash, "Linear Programming in Infinite-Dimensional Spaces", Wiley-Interscience, Chichester, 1987.
- [2] Bellman, R., "Dynamic Programming", Princeton University Press, Princeton, N.J., 1957.
- [3] D'Epenoux, F., "Sur un problème de production et de stockage dans l'aléatoire", Revue Française de Recherche Opérationnelle, Vol. 14, 1960, pp. 3-16.
- [4] Denardo, E.V., "Contraction mappings in the theory underlying dynamic programming", SIAM Review, Vol. 9, 1968, pp. 165-177.
- [5] Radner, R., "Dynamic programming of economic growth", Report 17, Center for Research in Management Science, University of California, Berkeley, 1964.
- [6] Sutherland, W.R.S., "Equivalent linear optimization problems", Methods of Operations Research, Vol. 45, 1983, pp. 345-356.
- [7] Tilquin, C.A.J., "On One-Sector Development Models", M.Sc. thesis, Department of Mathematics, Dalhousie University, Halifax, Nova Scotia, 1971.