Multiperiod Optimization: Dynamic Programming vs. Optimal Control: Discussion

Hovav Talpaz

So, have the dynamic programming's day arrived or not? From Burt's comprehensive review, one may possibly draw both answers: yes and no. The positive response is only logical given the list of applications in his paper for the last 20 years or so, which is by no means a complete one. Burt's useful suggestions of how to apply dynamic programming in the multi-dimensional state variable problems illustrate the much wider areas of applications which may be tried with high chance of success. On the other hand, difficulties arise with reducing the system's dimension and number of stages. The need for experience and creative imagination in using the dynamic programming methodology still leaves it out of range for many. In my discussion, I first comment briefly on Burt's major points, then prerequisites for popularizing dynamic programming will be laid out, two additional algorithms will be suggested, followed by a few comments on Zilberman's paper and some closing remarks.

Comments on Dynamic Programming

Many applications of dynamic programming, like Burt's equation (3), do assume the Markovian property, namely, that the system is completely defined and expressed by the current level of its state variables. However, as we move to improve farm management practices of high quality crops and intensive livestock enterprises, this assumption may be too strong to make, and if it is made, then the chance of implementing the new model is significantly reduced. Without a good algorithm to solve the non-Markovian system, one may be inclined to further simplify the problem by solving a static proxy system.

Regarding Burt's equation (3), for G(·) and h(·) to converge to a limiting function, it is required that G(·) be finite and

$$\lim_{n\to\infty} \beta^n E \; V_{n-1}(\cdot) \to 0.$$

In mentioning the obstacles to implementation, I would like to stress a few more:

1. Data requirements are substantially greater and much more expensive to get. This seems so obvious, yet it is a very serious problem in many farm management problems. Furthermore, there is an interaction between the data collection and our ability to utilize it in solving the problem. For example, in the management of water and fertilizer application over time for high-yield cash crop production, the question of how often to sample the states of the system is a crucial one, and may influence the solution. Another example is the insect scouting in cotton — should a sequential sampling be applied or not?
2. Simplification which goes too far leads to poor implementation and a reduced demand for further analysis.
3. It is the *uniqueness* of many dynamic programming problems, and the consequent special programming required to solve them, which is probably the

Hovav Talpaz is an associate professor in the Department of Statistics, ARO, The Volcani Center, Bet-Dagan, ISRAEL and in the Department of Agricultural Economics at Texas A&M University.

The author is grateful to C. Richard Shumway for his constructive comments.

greatest obstacle to implementation. G(u, x) and h(u, x, ε) in equation (3) might be linear, at best, or non-linear, or given by a highly complex subsystem to be evaluated by a computer simulation. Then, to obtain an efficient algorithm, this system has to be carefully analyzed using expertise in mathematics and operations research as well as programming skills. Of course, a team approach could be employed where specialists participate and cooperate on the project. This requires coordination, communication, and good working habits, and above all, good financial arrangements in bringing all that together.

4. I agree completely with Burt's observations on teaching the subject in class. In my opinion, a great deal of additional time in the classroom and at the computer terminal is needed. Bellman's principle of optimality and its translation into recursive equation can be taught, analyzed and illustrated in a relatively short time. Yet, to conceptualize solution frameworks for different structural problems is time consuming. I'm convinced that by-and-large most of our students do receive much less time than the critical mass needed for independent pursuit of the subject, allowing for some evaporation process.
5. As in many other areas, but more so here, only highly skilled instructors can teach dynamic programming well without actually being involved with it in their own research or practice. Only then can well-designed problems be offered as homework with the teacher standing by to bail out the inexperienced with hints and suggestions. The inexperienced teacher tends to stick to the few textbook problems available without providing the students with solution techniques and some alternative methods.
6. A major obstacle is that the "curse of dimentionality" is alive and well as long as conventional methods of solution like the *value-iteration* are applied. Burt has discussed dimension reduction methods, generally by some sort of simplification. Although these ways are certainly valid, it again requires careful analysis of the system which is expensive in time and money. Later, I discuss possible ways of overcoming this problem differently.
7. Finally, but by no means least important, economists like to "economize" a system or a model. However, they prefer to receive such a system as complete as possible — a pure "black box." They usually hate to be bothered with the small details and even consider the cooperative work in building such a model, with the physical or biological relationships, as an unnecessary must. Such models, usually in a form of numerical simulation designed by the non-economists, are built without the intent to maximize an after-thought objective function which often prove ill-structured for such a purpose. Although it may seem to the economist a waste of his time, I am sure, based on my own experience [e.g., Talpaz and Hillel; Talpaz et al.], that the marginal benefit/cost ratio is high for a greater involvement of the economist in the design of such a simulation model or in the coordination and specification of the overall effort.

Making Dynamic Programming More Popular

Many conclusions on how to popularize dynamic programming have been presented in Burt's paper. However, the most crucial prerequisite is the availability of efficient, and as standard as possible, algorithms for solution. For example, many deterministic dynamic programming problems with linear recursive equations can be solved easily by linear programming. Most of the "traveling salesman" and decision tree (or network analysis) problems are solved efficiently by linear programming.

Unfortunately, most problems in agriculture are non-linear. However, potentially relevant developments have been made recently in a wide range of disciplines in operation research, which can be of help. Two approaches are briefly discussed next.

A Dynamic Markov Chain Decision Model

Consider again the system under uncertainty as defined by Burt's equation (5) but as a finite state with infinite stages ($n \rightarrow \infty$) Markov decision process. This equation can be restated as

$$(1) \qquad h(i, k) = q_i^k + \beta \sum_{j=1}^{m} p_{ij}^k V_j;$$

$$i, j = 1, \ldots, m \text{ states } k = 1, \ldots, K.$$

Kislev and Amiad have formulated this problem as a linear programming problem some fifteen years ago. Recently, Denardo provided a generalized and simplified formulation, again as a linear programming problem, provided some regularity conditions hold (which is usually the case). With $a_1, \ldots, a_m$ any fixed constants:

$$\text{Max} \sum_{i=1}^{m} \sum_{k=1}^{K} q_i^k x_i^k$$

subject to

$$(2) \qquad \sum_{k=1}^{i} x_i^k - \beta \sum_{j=1}^{m} \sum_{k=1}^{k} x_j^k p_{ij}^k = a_i$$

$$\text{for } i = 1, \ldots, m$$

$$x_i^k \geqslant 0 \text{ for all } i, k.$$

This problem is known to be feasible and bounded, and any optimal basis for it has the following properties. For each i, the variable x_i^k is positive for exactly one k. Moreover,

$$\pi(i) = k \text{ whenever } x_i^k > 0.$$

Policy $\pi(i)$ is an optimal policy. That is, $\pi(i)$ satisfies $V^{\pi}(i) \geqslant V^{\delta}(i)$ for all i and all δ possible policies. π also satisfies the second version of the principle of optimality, namely: there exists a policy that is optimal for every state at every stage.

The alternative algorithms discussed in Burt which belong to the families of "policy interation" and "value iteration" as well as the method based on a "Newton Raphson iteration" [Blackwell] may for certain cases compete well with the above linear programming formulation. However, the linear programming codes are widely available and generally well understood by economists. They provide an excellent tool which is simple to use and allow for an easy sensitivity analysis.

Dynamic Programming with Non-Linear Functions

When a dynamic system is characterized by non-linearities in the objective function and/or the constraints, then in most cases, special algorithms are necessary for solving the problem, making it expensive. What is needed is a standard code, very much like linear programming or regression codes, where the user defines his problem and the code takes it from there. Major advances in non-linear constraint optimization were made in the 1970's, preparing the ground for such algorithms. The large software houses are busy building and improving their packages to handle these problems. For example, Murtagh and Saunders at the Systems Optimization Laboratory, Stanford University, have developed the MINOS/AUGMENTED package. Following is a brief description of this algorithm.

The problem must be expressed in the standard form:

$$\text{minimize} \qquad F(x) + C'x + D'Y$$

$$\text{subject to} \qquad f(x) + A_1 Y = B_1$$

$$A_2 x + A_3 Y = B_2$$

$$L \leqslant \begin{bmatrix} x \\ y \end{bmatrix} \leqslant U$$

where

$$f(x) = \begin{bmatrix} f^1(x) \\ \vdots \\ f^{m1}(x) \end{bmatrix}$$

and the functions $f^i(x)$ are smooth and have known gradients. x is the vector of the nonlinear functions in the objective and constraint functions, respectively. The equal signs of the constraints can be replaced by $\leq$ or $\geq$ inequalities. The A's are known constant matrices; C, D, B_1, B_2, L, and U are known vectors. The solution process consists of a sequence of "major iterations." At the start of each major iteration, F(x) and f(x) are linearized by the first two terms of the Taylor's approximation, then the specified problem is augmented as a Lagrangean objective function with linear constraints, given by

$$\begin{aligned} \text{minimize } & F^\circ(x) + C'x + D'Y \\ & - \lambda_k' [f(x) - \tilde{f}(x)] \\ & + \tfrac{1}{2} P[f(x) - \tilde{f}(x)]' [f(x) - \tilde{f}(x)]. \end{aligned}$$

Subject to the same constraints as above, with $\tilde{f}(x)$ replacing f(x) in the constraint equation and given by

$$\tilde{f}(x) = f(x_k) + J(x_k)\,(x - x_k),$$

where $J(x_k)$ is the Jacobian at the k^{th} step. λ_k are the dual values of the previous step and P is a penalty parameter which multiplies the quadratic penalty function to assure convergence. For other details, see Mortagh and Saunders.

Judging from test problems already tried with this package, I see no problem in solving optimal control problems such as maximizing present value of profits for the next 100 periods of a monoplist or oligopolist. Other problems such as firm's economic growth with a concave non-linear discounted utility function, non-linear consumption-investment constraints, and capital-investment difference equations with non-linear bounds on borrowing capacities, can be solved by this algorithm.

Few Remarks on Zilberman's Paper

This is a good review on the use of optimal control in agricultural economics with diverse application areas, along with literature references providing a guide for anyone who wishes to pursue the subject. Hence, I would like to limit my remarks on the subject of what Zilberman called the "lackluster performance" by the agricultural economists.

The main reason for low involvement in optimal control, in my opinion, is neither the immense data needs, which is a contributing factor, nor is it the lack of background which is a contributing factor too, both could be overcome in time. The main reason is that optimal control models rarely provide one with any better algorithm for a numerical solution. Quite the contrary, often a solution to the set of equations derived from the necessary conditions is more difficult than the application of a direct method as a constraint nonlinear optimization. Optimal control models do provide some excellent tools for qualitative sensitivity analysis, as well as, in generalizing economic theories on resource allocation, product distribution to markets, and have been used heavily in developing new economic theories through the examination of Pontryagin's necessary conditions.

Unfortunately, since optimal control in most cases does not provide one with an efficient algorithm, severe compromises and approximations are made, leading to exclusion from practical use. Zilberman's examples illustrate this point. In the irrigation problems, we see how fast the problem becomes complex as a new technology (e.g., drip) is considered. Only a single crop is considered, with a uniform production technique aside from irrigation. Having done all that, what we are left with is a set of equations (13)-(16) which have to be solved by special programming. Using other (more simplified) methods, agricultural economists, historically, considered heterogenous inputs in their analysis and practice. They have used heterogenous labor, water quality, soil types, various environmental conditions, as well as time. Furthermore, in many studies, all these "domains" have been considered simultaneously and solved for by efficient algorithms such as linear or quadratic programming. Since very often we owe the decision makers practical solutions, optimal control techniques have been exchanged from practical use. What we have been failing to do, though, is to generalize our findings by developing theories for better understanding and insight into the nature of the solution. That was well advocated by Zilberman in illustrating the role optimal control theory can play.

In a study aimed at improving aquaculture management [Talpaz and Tsur], an attempt was made to combine these approaches. First, optimal control was used to obtain optimal levels of biomass — extending Clark's work, hence, determining optimal harvesting, cycle length, and selection of fish size for harvest. Here, optimal control was very helpful in providing insights. Then, a multiple phase, nonlinear optimization algorithm, was employed to obtain the numerical solution.

An additional minor point should be made. When the objective function of any of Zilberman's examples is in discrete form, with Σ sign replacing $\int$, an efficient algorithm like MINOS, or in some cases separable programming, can handle the solution nicely.

Concluding Remarks

My own response to the question: Has the day arrived for dynamic programming and optimal control in agricultural economics? is a flat yes! Yet, it will require much more of our involvement in studying dynamic systems and in the development of the corresponding mathematical models while engaging with experts from disciplines other than our own. Recent technological advances made in non-linear optimization provide more usable ways to solve dynamic programming problems. These improvements are by no means the last word. Better convergence rates will be available in the near future. For example, adding another term of the Taylor's expansion to the non-linear objective function in the MINOS algorithm, should improve convergence, although a quadratic programming algorithm will replace the linear programming as the core code. Adding improved computer and telecommunication hardware, along with reduction in costs, dynamic programming and optimal control are bound to become more popular among us in the future.

References

Blackwell, D. "Discrete Dynamic Programming." *Annals of Mathematics and Statistics* 33(1962):718-26.

Clark, C. W. *Mathematical Bioeconomics: The Optimal Management of Renewable Resources*. New York: John Wiley & Sons, 1976.

Denardo, E. V. "Dynamic Programming." *Handbook of Operation Research Foundations and Fundamentals*, (eds.) J. J. Modor and S. E. Elmaghraby, New York: van Norstand Reinhold Co., 1978.

Kislev, Y. and A. Amiad. "Linear and Dynamic Programming in Markov Chains." *Journal of Farm Economics*, 50(1968):111-29.

Murtagh, B. A. and M. A. Saunders. "A Projected Lagrangean Algorithm and Its Implementation for Sparse Nonlinear Constraints." Technical Report SOL-80-1R, Department of O. R., Stanford University, 1981.

Talpaz, H. G. L. Curry, P. J. H. Sharpe, D. W. DeMichele and R. E. Frisbie. "Optimum Pesticide Allocation for Controlling the Boll Weevil on Cotton." *American Journal of Agricultural Economics*. 60(1978):469-75.

Talpaz, H. and D. Hillel. "Optimizing Soil Moisture Effectiveness by Mulch of Hydrophobic Aggregates Using Simulation and Numerical Search." *Southern Journal of Agricultural Economics* 8(1976):199-204.

Talpaz, H. and Y. Tsur. "Optimizing Aquaculture Management of a Single-Species Fish Population." *Agriculture Systems* (in press).