The Australian Journal of
Agricultural and Resource Economics
Journal of the Australian Agricultural and Resource Economics Society

The Australian Journal of Agricultural and Resource Economics, 53, pp. 251–264

The optimal extraction of water along an arbitrarily configured river system

Alex Coram and Lyle Noakes†

The fundamental problem for any scheme of water management that tries to maximise welfare across a river system is that of determining the optimal allocation at every point. The problem cannot, in general, be avoided by trading water rights because the price will not account for the effect of extraction at any one point on all other points. This article interprets the problem in terms of the indeterminacy that results from missing information on the value of water at internal junctions between rivers. It is then solved in an optimal control theory framework.

Key words: optimal water usage; river systems; control theory; Pontryagin.

1. Introduction

Social planners and policy-makers sometimes have to make decisions about the allocation of water from an arbitrarily large number of interconnected rivers in situations where demand exceeds supply in all, or some, parts of the system. This raises a straightforward and important question. How should water be allocated to various users at every point?

The issues around this question can be illustrated by considering the problems faced in managing systems such as the Colorado River Basin in the US or the Murray-Darling Basin in Australia. The Murray-Darling Basin, for example, contains more than 34 main rivers all of which eventually flow into the Murray river to give one end point and 33 points of intersection or nodes. The stated objective of the Murray-Darling Basin Commission is to implement an 'Integrated Catchment Management' policy for this entire system (Murray-Darling Basin Commission 2001). This is to be done by setting targets for catchment health that are intended to protect urban users, to protect those who use the river for irrigation, to protect other environmental and economic assets and to protect the health of the system as a whole by imposing an end of catchment area target on flows at the terminal point, however, defined.

The difficulty with the Murray-Darling Basin Commission's approach to integrated catchment management is that the sorts of criteria proposed are more directed at the question of capacity building than developing a clear set

† Alex Coram (email: a.b.t.coram@rgu.ac.uk), Department of Political Science, University of Western Australia, Stirling Highway, Crawley 6009, Australia and Aberdeen Business School, Robert Gordon University, Aberdeen, Scotland. Lyle Noakes, Department of Mathematics and Statistics, University of Western Australia.

doi: 10.1111/j.1467-8489.2007.00437.x

of protocols or models of interaction that might help provide some principles that could be used to inform decision making and discussions between the different interests within catchments. In addition, it is not clear how the catchment approach can be extended to the system as a whole. There are, for example, considerable problems in setting targets for catchment health. Not only is it difficult to work out how environmental and economic demands are to be reconciled, but there is also the considerable difficulty presented by the fact that end point targets for one catchment area are starting point flows for a downstream catchment. This means that it is not only necessary to deal with different environmental and economic priorities within catchments. It is also necessary to deal with different priorities between catchment areas.

In a similar manner the National Water Commission in Australia has called for a 'nationally consistent approach to water planning' that is able to settle 'the trade-offs between competing outcomes for water systems, based on the 'best available science.' (National Water Commission 2008). This again raises questions about the criteria to be used to determine consistency between different sorts of users and across different geographical regions.

This article suggests one way to approach the problem of establishing some consistent principles for water allocation would be to start by determining the allocation of water at each point that would maximise welfare across the system, in some sense. This would, at the very least, provide a basis for the sort of integrated view and consistency that both Commissions require in that the welfare implications of extraction at any upstream point for all downstream points could be more clearly understood and interests at different stages of the system could be treated in a balanced fashion.

It must be stressed that we do not claim that a calculation of welfare would, in any sense, provide a set of criteria that would substitute for negotiation and discussion, or for the capacity building that is a central theme of the Murray-Darling Commission's report. Nor do we believe that the approach here substitutes for detailed policy work. We are aware that welfare is a contested concept. It is also likely that any attempt at using welfare as a means of reconciling different interests, and the welfare functions that are attributed to different interests, would all be the subject of conflict and debate. On the other hand conflicts over the appropriate utility functions are likely to be no more intense than those over direct allocation.

Offsetting these difficulties is the advantage that, unlike direct allocation approaches, a welfare based approach gives some idea of the implication of different allocations across the entire system based on a common yardstick. This should be of direct assistance to policy-makers and of assistance more generally in the process of debate and bargaining. In addition, it would help guide discussion by providing a basis for understanding what sort of claims can be justified on at least one set of fairly broad criteria.

In what follows we make a start on considering how this welfare optimisation problem might be solved. It turns out that the solution for a system of interconnected rivers, or for any system with interconnected flows of this type, is not straightforward. This means that the paper is necessarily fairly abstract. It is mostly concerned with establishing that the problem can, in fact, be satisfactorily solved in the first place and with determining a method of solution. Unless these hurdles can be overcome there is no point in discussing a welfare based input into the policy discussion.

Some notes on the implementation of a programme for numerical calculations, are provided in the second last section of the paper. As is shown, the programme required to implement the approach is easy to run and requires little computing time.

A more difficult problem is in determining the type of functional relations between the various interests and water extraction that would be required to set up the model. On the other hand, given the low computational costs, it would be feasible to run a variety of models based on different sorts of assumptions about payoff functions as inputs to debates on allocation. This might be of some interest from a policy perspective and as an adjunct to public discussion.

Although acquisition, or collating, the required data at sufficiently large number of points is fairly straight-forward it is, in itself, a major project. This is well-beyond the scope of this paper and would require a separate, dedicated, research initiative. Such an initiative would meet the National Water Commission's desire for increased 'inputs from socio-economic analyses' and evidence based planning.

What we can show, however, is that calculation across an entire system is feasible. Moreover, it is also feasible from a public policy perspective since, by the standards of river management, the required inputs are inexpensive.

In order to set this out more precisely we begin with a statement of the theoretical structure of the problem in Section 2. A general model of the water allocation problem is developed in Section 3 and analysed in Section 4 where the main results of the study are presented. We give an example in Section 5. We discuss the procedure for a numerical solution to the optimisation problem in Section 6.

2. The problem

The problem of analysing access to water in a manner that maximises welfare across all users in a river system has so far not received much attention in the literature and has not been given a rigorous formal treatment. The few studies that have treated a river as a flow in order to consider allocation problems have only considered the problem under restricted conditions, or for flows between discrete points. The article by Quiggin (1988), for example, considers the problem of maximising welfare across a number of different activities and takes into account the impact of upstream users on downstream users. He does this by adopting an asset value approach in order to examine the way in which upstream agricultural activities may degrade water supplies for subsequent users and hence influence their patterns of production. He only studies the case where water is used at a small set of discrete points and the optimisation

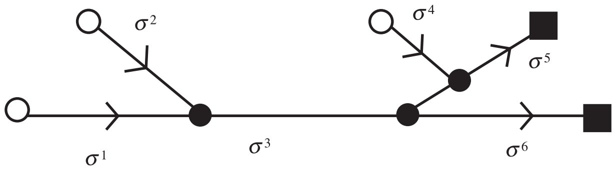

Figure 1 Example of a river system.

problem can be solved with standard optimal programming methods. In addition, his model avoids the problem that is of concern here by assuming that there are no river junctions between points of extraction. Other studies which have treated a river as a flow, such as those by Weber (2001) for example, have also been limited to discrete number of points with no intervening junctions. There are no detailed explorations of the optimal programme across a complete system with continuous inflows and outflows at every point and multiple points of intersection between rivers. This problem was raised and described in Coram (2006) but not solved.

In order to appreciate the problem begin by observing that the uni-directional ordering of agents along a set of flows means that assumptions about normal market conditions and independence of traders underlying much of the work on markets for water rights do not hold. This is because a river is a flow in which the agents have a uni-directional ordering.[1] It follows that, if a market for rights produces a uniform price, it cannot maximise efficiency and certainly cannot maximise welfare, so a different approach needs to be found. This point is illustrated in the example at the end of the article.

The natural way to think about the welfare problem along a flow is in terms of dynamic optimisation with a payoff function along the entire river system and the control being the amount of water extracted at each point. An alternative to this might be to try and deal with the problem in terms of network optimisation, and a glance at Figure 1 could suggest this approach. Although this idea is attractive it has a number of drawbacks. The most important of these, for our purposes, is that network optimisation can only deal with gains and losses at nodes where two or more arcs meet. This means that it cannot deal with gains and losses along any river in the system, or with extraction at any points other than at a junction between rivers (Trick 1996).

The general problem that we face at this point is that standard dynamic optimisation approaches cannot be applied to this type of system in a straightforward way. There are at least two reasons for this, which present independent sub-problems, simply referred to as problems in what follows:

1. The first of these concerns the situation where there are dams and other means of storing water, and where there are fluctuations in the water level

[1] The literature on rights is substantial. See, for example, Burness and Quirk (1979); Kanazawa (1991); Booker and Young (1994); Weber (2001) and Freebairn and Quiggin (2006).

over time at each point. In this case time of extraction matters and we have to deal with some sort of distributed parameter system.
2. Second, and, we think, most pressingly, is that dynamic optimisation techniques normally deal with problems in which there is a continuous single flow of some variable, such as time, in the dynamics. Where there are multiple rivers, however, we have something like a problem of optimising across several different flows simultaneously. In addition, there will be jumps in the amount of water at points where rivers intersect and some means has to be found of getting information about the value of the water at these points. Unless this type of problem can be dealt with, it is difficult to see how it is possible to determine the optimal rate of extraction across the entire system.

This article is concerned with the second of these problems. It considers this for the special case where there are sufficiently many points of extraction that the returns from use along the system can be approximated by continuous functions.[2]

One preliminary question that should be dealt with is, why can't this problem be overcome by dealing with jumps in a state variable as purchases or sales at a point? This is an established technique in the literature, see Seierstad and Sydsaeter (1987, pp. 194–210). If the price were given, as well as the function that determined the change in the amount of stock for a given outlay, it would be possible to calculate the optimum amount to be bought or sold at any junction between rivers from the shadow price for the resource. The amounts derived from this fictitious transaction would then correspond to the desired inflows or outflows.

This technique won't work for a system of rivers, however, since price is not exogenous. This means that the payoff function for the purchase of water from a tributary at a junction and the function for the change in stock, are not well-defined and there is insufficient information to solve the problem in this way.

In what follows we specify the problem for applying dynamic optimisation theory in terms of the indeterminacy in the differential equations that results from missing information on the end points at internal junctions between rivers. We show how the Pontryagin principle can be used to provide this information.[3]

3. The river system

The system is made up of n interconnected rivers and is written as S. Any continuous flow of water that terminates at the sea, or in a lake, or marsh, or

[2] This gives considerable advantages in the analysis and should provide a reasonable approximation to the discrete problem. The technique developed in this article will hold for either case.

[3] For an accessible introduction to the Pontryagin principle see Macki and Strauss (1982).

wherever, and any flow that terminates on another path, is thought of as a separate river. Each river is labelled $\sigma^i \in S$ where i is the index $i = 1, \ldots, n$. An internal point in the system where rivers terminate on another river or commence from another river is called a node. Where rivers flow into and out of a lake, the lake could be treated as a node for present purposes. A point where a river starts that is not on another river, say in the mountains is called a starting point. A point where a river terminates with the property that no other rivers in the system leave that point is called a final termination point. Examples would be where a river flows into the sea, or a lake with no outflowing rivers in S, or into another political jurisdiction. See Figure 1 for an example. Nodes are represented as solid circles and starting points as open circles. Final termination points are solid squares. It is assumed that rivers collect or lose, water as the result of rain, evaporation or seepage along their way in the normal manner.

There is an arbitrarily large number of agents who wish to extract water in different intervals throughout the entire system. It is assumed that extraction imposes a cost on other users, or the community and floods and other situations where it might be desirable to reduce the level of the river are ignored. Other users would also include individuals who use the river for recreation, or derive an income from recreational use, those who derive aesthetic and other forms of enjoyment, and environmental and community interests. In order to keep the presentation simple, it is assumed that each of these interests can be represented by a single function, in general different for each interest, that covers the payoff from extraction and from the quantity of water in the river at any point. The quantity of water might also be taken as a rough proxy for the effects of pollution and return flows. Inflow of a pollutant, for example, could be taken as a reduction in the quantity of water available. If the problem is extended to multiple users, the reduction in the quantity of water may be different for users of different types.[4]

The planner's problem is to find the amount that should be extracted at each point in order to maximise an additive welfare function that depends on the payoffs to all individuals affected by extraction from the river. In order to see how this might be done we need to specify the problem in detail.

The flow of water along a river σ^i is approximated by the almost everywhere continuously differentiable function $x^i(t)$ where t is a point on the river and the length of each river is normalised to give $t \in [0,1]$. It is assumed that water can be extracted in any interval along the river and, taking advantage of the properties of flows, that we can approximate the amount extracted in an arbitrarily large number of intervals by the almost everywhere continuously differentiable function $u^i: [0,1] \rightarrow \mathbf{R}$. Since u^i is to be chosen it is

[4] It would be possible to extend the analysis to the case where there is one or more types of pollutant that cannot be represented in terms of the volume of water in the river. In this case it would be necessary to specify the problem in terms of the appropriate state and control vectors. This is left for subsequent research.

considered the control variable. A control vector is given by $u = (u^1, \ldots, u^n)$ where $u \in U$.

The change in the volume of water per unit of time at every point along a river σ^i can now be written as

$$\dot{x}^i := g^i \tag{1}$$

for $i = 1, \ldots, n$ where $g^i := g^i\,(u^i, a^i)$ and $a^i(t)$ is the gains and losses due to rainfall or evaporation, or other inflows and outflows that are not part of the control programme in the interval around t.

The payoff function for extracting water in an interval around t in σ^i can be written as

$$f^i := f^i(c^i u^i, \hat{c}^i x)$$

where $f^i = \sum f^i_q$ can be considered as the composite function made up of the sum of all interests we might wish to include in the model and the constants c^i, $\hat{c}^i$ take account of the normalisation of length to give $t \in [0,1]$.

It is also possible that a scrap value might be attributed to a final termination point. This water may have some environmental or commercial value, or value to the next jurisdiction. This is written as $\psi^i(x^i(1))$. It is assumed that $(\mathrm{d}\psi^i/\mathrm{d}x^i)|_{t=1} > 0$ for $x^i(1) < \bar{x}^i$ where $\bar{x}^i(1)$ is the natural level.

In addition, it may also be desirable to add constraints on the amount to be extracted at each point. These might take the form $x^i(t) \geq b^i(t)$ or $u^i(t) \leq b^i(t)x^i(t)$, for example, where $b^i(t) \geq 0$ is defined at all t along σ^i.

It follows that the planner wants to find the rate of extraction at every point that maximises the function

$$J := \sum_{i=1}^{n} \left(\int_0^1 f^i \mathrm{d}t + \psi^i(x^i(1)) \right) \tag{2}$$

subject to the dynamics in Equation (1), the constraints on extraction, and the endpoint conditions.

It is assumed that this problem has a unique solution and that this solution can be found in the sense that all the required boundary condition problems can be solved for the differential equations. It is also assumed that locally optimal solutions are also unique. We also assume that the optimal solutions are continuous with respect to variations in their parameters. These assumptions are reasonable for practical problems of this type.[5]

As is obvious, the task of finding this solution is not straightforward. This is because multiple rivers may flow into and away from a node, and there are

[5] This assumption would be expected to hold, except in pathological cases where the two point boundary condition problems generated may not be solvable. See Noakes (1998) for a discussion of this and for a development of computational algorithms.

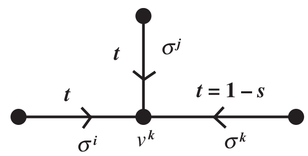

Figure 2 Example of a subsystem $\sigma^i \cup \sigma^j \cup \sigma^k = \langle v^k \rangle$.

multiple termination and commencing points. It follows that the amount of water at any node will depend on the solution to optimisation problems at other parts of the system. This means that it is not possible to apply the Pontryagin conditions across the whole system simultaneously, or start at a terminal point and use backward induction as might be done in a standard dynamic programming problem.

In the next section we show how the conditions required for a solution can, in principle, be found.

4. The optimal allocation

The optimal allocation at every point across the entire system will be determined by reducing the problem to a set of sub-problems that have the form required to obtain a local solution using the Pontryagin maximum principle, and then using these local solutions to get the global allocation. This is done by using the parameter values from each local solution and progressively iterating across the entire system with the new values as data in each iteration. This process is discussed again below when we have demonstrated that it will work as claimed.

In order to reduce the problem in the required manner it is assumed that there are K nodes in the system and specify a subsystem $\langle v^k \rangle \subseteq S$ for $k \in \{1, \ldots, K\}$ as follows. Let $\langle v^k \rangle$ be made up of the $m \leq n$ rivers that terminate on, or commence from, the node v^k. Each river in $\langle v^k \rangle$ commences or ends either at some other node, a starting point or a final terminating point. With the substitution $t = 1 - s$ in all outflowing rivers $\langle v^k \rangle$ can be rearranged so that v^k is the terminating point for all m rivers in $\langle v^k \rangle$. This means that, with suitably defined reverse flows, the only difference between nodes in the system is given by the number of terminating rivers. See Figure 2 for an example where all endpoints are nodes.

We need to show that it is possible to get the information required to solve for each sub-systems $\langle v^k \rangle$ where $\cup_k \langle v^k \rangle = S$ and that the solutions for the sub-systems will converge to the global solution by iteration.

Let the m state variables x^{ki} correspond to the flows along the rivers, σ^i, that enter or leave v^k and the m control variables u^{ki} correspond to extraction from these rivers where $i = 1, \ldots, m$. Define reverse flows so that all x^{ki} terminate on v^k. The Lagrangean for the reverse flow problem is

$$L^k := \sum_i \alpha_0^{ki} f^{ki} + \sum_i \alpha^{ki} g^{ki} + \sum_j \sum_i \mu^{kji} h^{kji}$$

where $\alpha_0^{ki}(t)$ and $\alpha^{ki}(t)$ for $i = 1, \ldots, m$ are the costate variables, h^{kji} are the constraint equations, expressed in the usual form, and μ^{kji} are the appropriately specified Lagrangian multipliers.

The Pontryagin principle allows us to set $\alpha_0^{ki} := 1$ under usual conditions. This means that we have m costate variables and, under the conditions required for a solution, these enter into the necessary conditions as differential equations.

$$\dot{\alpha}^{ki} = -\frac{\partial L^k}{\partial x^{ki}}$$

If the terminal values of any state variable, x^{ki}, is not specified the associated costate is required to satisfy a transversality condition which gives its value at the termination point of the river. Since we have m differential equations for the x^{ki} from Equation 1 the system now has $2m$ differential equations. In order to determine the solutions for these we require $2m$ boundary values.

In any subsystem $\langle v^k \rangle$ the functions x^{ki} for $i = 1, \ldots, m$ will have initial values at $t = 0$ or $s = 0$ in any river where the flow is reversed. In the case where a river $\sigma^i \in \langle v^k \rangle$ has a final termination point either the value of $x^{ki}(1)$ is specified, or the value of $\alpha^{ki}(1)$ is given by the appropriate transversality condition in the original problem. The value $\alpha^{ki}(1)$ is imposed on $\alpha^{ki}(0)$ in the reverse flow problem. This means that we have information on m initial values. Write the set of initial values for $\langle v^k \rangle$ in the reverse flow problem as $w^k(0) := \{x^{k1}(0), \ldots, x^{km}(0), a^{k1}(0), \ldots, a^{km}(0)\}$.

It is now necessary to show that the information on m of the values in $w^k(0)$ is sufficient to determine the $2m$ values needed to solve the differential equations. The needed values may, of course be in $w^k(1)$. It is proven that this is possible in the following theorem.

Theorem 1. *Suppose m-values from the set $w^k(0)$ are given. Then there is sufficient information to provide the 2m conditions required to solve the reverse flow problem for $\langle v^k \rangle$.*

Proof. Assume that, in the original problem, the number of rivers flowing into the node v^k is r and that $m - r$ rivers flow out. Since the sum of flows in and out at the node v^k is zero in the reverse flow problem the sub-system $\langle v^k \rangle$ must terminate on the manifold given by

$$M = \{x(1) : \gamma(x(1)) = 0\}$$

where $\gamma(x(1)) := x^{k1}(1) + \ldots + x^{kr}(1) - x^{k(r+1)}(1) \ldots - x^{km}(1)$. This gives one of the additional conditions that are needed leaving $m - 1$ to be found. The Pontryagin principle tells us that the costates are orthogonal to M in the optimal programme. The orthogonal vector is given by $\nabla\gamma = p(1, \ldots, 1, \ldots, -1)$

where there are r positive terms and $m - r$ negative terms for p some arbitrary constant. Hence the costate vector for the reverse flow problem is $\alpha^k(1) = p\nabla\gamma^T$ and this gives m additional terminal conditions, with one additional unknown, as required.

It now needs to be shown that we can get as close as we wish to the optimal rate of extraction for the whole system by solving locally across all K nodes in any order and then iterating this process a sufficient number of times. Roughly speaking, this is because whenever the optimum is calculated for $\langle v^k \rangle$ using optimal values of $x^{ki}(0)$ from a previous calculation there must be a local, and hence global, improvement.

Write the iteration algorithm as F where F is constructed such that $J(F(u)) \geq J(u)$ with equality only when u is everywhere locally optimal. This gives:

Theorem 2. *If it is possible to solve the optimisation problem for each $\langle v^k \rangle \subseteq S$ then using F to iterate across all nodes will cause $F(u) \rightarrow \bar{u}$ where $\bar{u}$ is the unique solution to the optimisation problem for S.*

Proof. The proof is immediate if the river has zero nodes or one node since it has been assumed that a solution is always possible. Assume the river has more than one node and let v^k be a node with r inflowing rivers and $m - r$ outflowing rivers. Take known values of $x^{ki}(0)$ for $i = 1, \ldots, r$ and $\alpha^{kj}(1)$ or $x^{kj}(1)$ depending on the problem for $j = r + 1, \ldots, m$ and solve the optimisation problem for the first round for each $\langle v^k \rangle$ for all k. What must be shown is that iteration using $\alpha^{ki}(1)$ or $x^{ki}(1)$ and optimal values of $x^{ki}(0)$ and or $x^{ki}(1)$ from the previous calculation converges to some $u = \bar{u}$ and that $\bar{u}$ is the optimal control.

With U given the uniform norm the function $J : U \rightarrow \mathbf{R}$ is continuous and we can define

$$X \subseteq U := \{u : J(u) \geq J(u_0)\}$$

for some u_0. Iteration gives $F\colon U \rightarrow U_0$ where $U_0 \subseteq X$. Since $u \in U_0$ is piece wise optimal it can be parametrised by $K + 1$ endpoints. This means that we can consider $U_0 \subseteq X \subseteq \mathbf{R}^{K+1}$. Since X is bounded its closure is compact and this means that the sequence

$$u^{(0)}, u^{(1)}, \ldots, u^{(k)} = F(u^{k-1})$$

on $\bar{X}$ has a convergent subsequence

$$u^{kj}, u^{kj+1}, \ldots \rightarrow u^{kj\infty}$$

To show that every subsequence converges to the same $u^{kj\infty} = u^{\infty}$ note that $J(u^{kj}) \geq J(u^{kj-1})$ and hence $J(u^{kj\infty}) = \sup J(u^{kj\infty}) = V^{kj\infty}$. This gives

$$J(F(u^{k_j})) \geq J(u^{k_j}) \Rightarrow J(F(u^{k_{j\infty}})) \geq J(u^{k_{j\infty}}) = V^{k_{j\infty}}$$

and

$$J(F(u^{k_j})) = J(u^{k_{j+1}}) \leq V^{k_{j\infty}} \Rightarrow J(F(u^{k_{j\infty}})) \leq V^{k_{j\infty}}$$

which means that

$$J(F(u^{k_{j\infty}})) = J(u^{k_{j\infty}})$$

From the construction on F, this is only possible when $u^{kj\infty}$ is everywhere locally optimal. It follows that $u^{kj\infty} = u^{\infty} = \bar{u}$ for any subsequence.
□

It will be noted that the constraints in L^k on the absolute amount, or the proportion of, water to be extracted are included in each local, and hence in the global, solution. A constraint of the type $x^i(t) \geq \bar{b}^i$ imposed on σ^i starting from node v^k also affects extraction from all rivers that feed into σ^i. These will be accounted for by the permissible volumes of water at v^k.

5. Example

Consider the optimisation problem for a system made up of three rivers with a single point of intersection. It is assumed that $f^i(c^i u^i, \hat{c}^i x)$ is concave in x^i and u^i for $i = 1, 2, 3$ and that the value of f^i and the value of extraction increases with the amount of water in the river so that $(\partial f^i/\partial x^i) > 0$ and $(\partial^2 f^i/\partial u^i \partial x^i) > 0$. It is also assumed $(\partial g^i/\partial u^i) = -1$. Let σ^1 and σ^2 be inflowing and assume $x^1(0)$ and $x^2(0)$ are known. σ^3 terminates with a scrap value $k(x^3(1))$. There are no inequality constraints on extraction.

The Hamiltonian for the system is

$$H: = \sum_i (f^i + \alpha^i g^i)$$

and the necessary conditions are

$$c^i \frac{\partial f^i}{\partial u^i} = \alpha^i \text{ for } i = 1, 2 \text{ and } c^3 \frac{\partial f^3}{\partial u^3} = -\alpha^3$$

and

$$\dot{\alpha}^i = -\hat{c}^i \frac{\partial f^i}{\partial x^i}$$

with $\alpha^3(0) = k$ in the reverse flow problem.

It follows immediately that, if the volume of water in the river is increasing or remains the same prior to extraction, the amount of water to be extracted increases in the optimal programme as we move downstream along any river.

If, on the other hand, the volume of water is decreasing owing to evaporation or other losses independent of extraction, that the amount extracted may need to decrease to satisfy the second partial derivative on f^i. This means that, in any pricing system, the price of water should either decline as we move downstream, or vary according to the flow in the river.

In order to complete the solution note that $x^1(1) + x^2(1) - x^3(1) = 0$ in the reverse flow problem. From Theorem 1 the costates are given by $p(1, 1, -1)^T$, for p a constant. Hence

$$\alpha^1(1) = \alpha^2(1) \text{ and } \alpha^1(1) = -\alpha^3(1)$$

Solving for α^3 for $s = 1 - t$ gives

$$\alpha^3(s) = -\int_0^s \hat{c}^3 \frac{\partial f^3}{\partial x^3} \mathrm{d}r - k$$

and this gives

$$\alpha^1(t) = \int_t^1 \hat{c}^1 \frac{\partial f^1}{\partial x^1} \mathrm{d}r + \int_0^1 \hat{c}^3 \frac{\partial f^3}{\partial x^3} \mathrm{d}s + k \text{ and } \alpha^2(t) = \int_t^1 \hat{c}^2 \frac{\partial f^2}{\partial x^2} \mathrm{d}r + \int_0^1 \hat{c}^3 \frac{\partial f^3}{\partial x^3} \mathrm{d}s + k$$

It is now possible to use this and the initial conditions on $x^i(0)$, together with the equation for $x^i(1)$, to solve the system of differential equations for all x^i and α^i.

6. Discussion of the iteration procedure

The iteration procedure can be described informally as follows. By starting with any node chosen at random and holding the end points for each river flowing into or out of that node constant it is possible to use the Pontryagin principle, in order to get a new parameter value for the flow at that node which improves the payoff. This parameter value can now be used in some subsequent calculation. We can think of this in terms of adjacent nodes, but in practice nodes can be chosen in any manner. This process is repeated across the entire system with checks to ensure that the iteration covers all nodes, and does not repeat a node until a parameter has been updated. In this way each local solution to the optimisation problem affects subsequent optimisations by feeding in updated parameter values.

This updating procedure can be thought of in a similar way to Gauss-Seidel or Jacobi relaxation methods for systems of linear equations. In this case the procedure is to guess an initial set of values for the unknowns. These are then used to solve part of the system to produce a new set of values which can again be fed into the system for other iteration. If the system is diagonally dominant the process it can be shown that the process will converge to the correct solution.

In our problem the system is not linear but it is straightforward to see that the payoff, that is the value of J, must increase at each stage in the iteration since the parameter values must improve. This constant improvement is not, however, in itself, sufficient to prove that iteration would produce convergence to the optimum rate of extraction and we needed Theorem 2 to guarantee that this would happen.

It will be noted that, since every local problems is small scale, each step in the iteration procedure is capable of rapid implementation. The only physical data required to solve the problem at a node are the volume of water at the commencement point for a river flowing in, or the terminal point for a river flowing out. Once this initial data is acquired it can be fed into any programme that gives a numerical solution to the optimal control problem specified in terms of the payoff function f. In order to get the iteration all that is necessary is that the output values for the parameters from the previous step are fed into the next step using an appropriate updating routine.

It is possible, of course that the routine might not terminate. It will be noted, however, that convergent sequences are Cauchy and that the difference between any two successive values given by $|u^{(k)} - u^{(k-1)}|$ is decreasing. This means that we can put a stop on the iteration when $|u^{(k)} - u^{(k-1)}|$ becomes sufficiently small.

Given the size, and routine nature, of each local problem the computational demands are small for a reasonably specified f. The appropriate routine could be implemented using any commercially available package that deals with optimal control such as MISER 3 or MATLAB.

7. Conclusion

This article has argued that, to provide the basic criteria for making the sort of consistent decisions that are required by the Murray-Darling and the National Water Commissions, a solution to the problem of calculating the optimal rate of extraction at every point across a river system made up of an indeterminate number of rivers would be useful. It is believed that such calculations may be of direct use in allocation or, more reasonably, in the process of negotiation between interests and in providing some common measure of the impact of the actions of some interests on others. It is also believed that they would go some way to meeting the requirements of the National Water Commission for increased inputs for socio-economic analysis.

A solution to the problem was provided for a system with an indeterminate number of rivers and it was shown how this could be implemented by a basic iteration procedure. It was observed that the solution could be extended to cover the case where there is more than one type of payoff, including that to environmental groups and recreational users who do not extract water from the river. In addition it could also be extended to the case where there are different types of pollutants or water qualities running back into the river from users by extending the dimension of the problem. An obvious example of such an extension would be salinity related problems.

It was observed that the most significant conceptual problem in applying the solution would come in estimating utility functions for various users. A more practical problem is that, to get the required data on flows at all relevant nodes resources, well-beyond those available to the authors of this paper, would be needed. The costs of getting this data would be relatively low, however, and could easily be funded.

It is also worthwhile noting that the solution could be implemented at little computational cost. This means that, once the initial data on the river system were given, it would be practical to run a variety of solutions with different payoff functions. Alternatively a model could be built sequentially for increasingly complicated assumptions about payoffs, or additional sets of payoffs, to test different assumptions.

References

Booker, J.F. and Young, R. (1994). Modeling intrastate and interstate markets for Colorado River water resources, *Journal of Environmental Economics and Management* 26, 66–87.

Burness, H. and Quirk, J. (1979). Appropriate water rights and the efficient allocation of resoures, *American Economic Review* 69, 25–37.

Coram, A. (2006). The optimal allocation of water along a system of rivers: a continuous model with sequential bidding, *Australian Journal of Agricultural and Resource Economics* 50, 313–26 and *Corrigendum* 51, 112.

Freebairn, J. and Quiggin, J. (2006). Water rights for variable supplies, *Australian Journal of Agricultural and Resource Economics* 50, 295–312.

Kanazawa, M. (1991). Water quality and the economic efficiency of appropriate water rights, in Dinar, A. and Zilberman, D. (eds), *The Economics and Management of Water and Drainage in Agriculture*. Kluwer Academic, Boston.

Macki, J. and Strauss, A. (1982). *Introduction to Optimal Control Theory*. Springer-Verlag, New York.

Murray-Darling Basin Commission (2001) *Integrated Catchment Management in the Murray-Darling Basin 2001–10*. Murray-Darling Basin Ministerial Council, Canberra, ACT.

National Water Commission (2008). Water planning in Australia position statement. Available from URL: http://www.nwc.gov.au/. [accessed 7 April 2008]

Noakes, L. (1998). A global algorithm, *For Geodesies Journal of Australian Mathematical Society of (Series A)* 64, 37–50.

Quiggin, J. (1988). Murray river salinity – an illustrative model, *American Journal of Agricultural Economics* 70, 635–645.

Seierstad, A. and Sydsaeter, K. (1987). *Optimal Control Theory with Applications*. Elsevier, North Holland.

Trick, M. (1996). *Network Optimization*. Available from URL: http://mat.gsia.cmu.edu/classes/networks/networks.html. [accessed 3 July 2008]

Weber, M. (2001). Markets for water rights under environmental constraints, *Journal of Environmental Economics and Management* 42, 53–56.