Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Agricultural Economics 31 (2004) 81–96
AGRICULTURAL ECONOMICS
www.elsevier.com/locate/agecon

Calculating the cost of irrigation induced soil salinization in the tungabhadra project

John Janmaat*

Department of Economics, Acadia University, Wolfville, NS, Canada B0P 1X0

Received 25 October 2001; received in revised form 4 October 2002; accepted 1 June 2003

Abstract

Irrigation projects in developing countries have a history of poor performance. Inefficiencies result as water applications deviate from plans and induce greater than projected rates of soil degradation through water logging and salt accumulation. Over time, the collective impact of these forces will converge to an equilibrium with a level of output that may be far below the system's potential. The Tungabhadra Project in south west India is experiencing all of these problems. Integrating geographic, hydrologic, biologic and economic features, the lost production value is estimated for a range of equilibria to which this system may converge. For the lower left bank main canal of the Tungabhadra project, the total economic cost of soil degradation are approximately 14.5% of the system's productive potential while sub-optimal distribution losses may approach 37.1%.

JEL classification: N5; Q0; Q1

Keywords: Soil degradation; Economic cost; Irrigation; Semi-arid agriculture

1. Introduction

Soil degradation can cause serious productivity losses for irrigated agriculture. Two forms of soil degradation that are common in irrigated agriculture are salinisation–the accumulation of chemical salts in the root zone–and water logging–the presence of excessive water in the root zone. The amount of salt in the soil column is controlled through management of the balance between the removal of salts by leaching and the addition of salts in irrigation water. Insufficient volumes of water cannot remove salts.

Salinisation is a long term consequence of being too frugal with water in the short term. However, excessive water availability, either through heavy application or inflows from other fields, also depresses yields. When the water available to the plant exceeds the optimal amount, the plant 'drowns'. A continuous high water table can also lead to a breakdown in the structure of the soil. The water management problem is further complicated by the fact that salt bearing water can flow downhill between soil blocks. Thus, one farmer's solution to soil degradation becomes another farmer's problem.

The Tungabhadra project (TBP) in south west India provides a site for the estimation of the cost of soil degradation in an irrigation project. The TBP is a protective irrigation system, designed to provide limited supplemental irrigation over a large area (Hugar, 1997; Mollinga, 1998). The distribution of water in the TBP has implications both in terms of economic efficiency

* Tel.: +1-902-585-1461; fax: +1-902-585-1070.
E-mail address: jjanmaat@acadiau.ca (J. Janmaat).

0169-5150/$ – see front matter
doi:10.1016/j.agecon.2003.06.001

and equity. If the institutional instruments to optimally distribute the water supply are unavailable, substantial inefficiencies will result. With area based, rather than volume based water pricing, farmers face a marginal input cost for water that is near zero. Since the total supply is finite, excessive use of water by some limits the water available to others. Efficiency pricing of water (Vaux and Howitt, 1984; Vaux, 1986) could reduce these problems. Unfortunately, many barriers prevent the efficient pricing of water (Tsur and Dinar, 1995, 1997).

Soil degradation costs are not borne equally throughout the irrigation system. Downhill farmers receive saline water from uphill farms, suffering greater losses. The inequitable distribution of soil degradation costs often compounds the inequities that are generated by deviations from the planned water distribution pattern. Since these systems are often designed to increase equity between farmers, their performance is a failure on both efficiency and equity grounds. The objective of this paper is to develop a model of the TBP which can be used to estimate the economic costs of soil degradation separately from the costs of sub-optimal water distribution.

The paper is organised as follows. In Section 2 some background material concerning the problem of soil degradation in irrigated agriculture is reviewed. A model of the TBP is then developed, emphasising the linkage between the field where the water is applied, the channel which delivers water to the field, the distributary which delivers water to the channel, and the main canal which delivers water to the distributary from the reservoir. The output of this model is then described, starting with the field and moving up to the system as a whole. In the penultimate section some qualifications of these results are discussed. This discussion is followed by a conclusion that wraps up the paper.

2. Background

The role of agriculture in the development of the Indian economy is widely recognised (Bhatia, 1988; Mollinga, 1998; Rao, 1994). The importance of agricultural development, in particular irrigation, is highlighted by two main concerns. First, India cannot industrialise fast enough to absorb the potential outflow of labour from the rural economy. Irrigated agriculture enhances rural employment opportunities, offsetting the attraction of urban areas. Second, the rural poor are particularly vulnerable to drought, a constant risk in the monsoon-dependent regions of India (Rao, 1994; Bhatia, 1988). The primary justification for the TBP was the provision of protective irrigation.

Large irrigation projects have seldom approached design performance. In a comprehensive analysis, Chambers (1988) finds that most large irrigation projects in developing countries are performing far below design targets. He attributes the poor performance in part to a failure to understand the various linkages between the loci of control generated by the system. Perry and Narayanamurthy (1998) explore farmers' response to rationed water. Protective irrigation systems are designed to maximise return per unit of water. Perry and Narayanamurthy (1998) argue that farmers aim to maximise return per unit of land, and attempt to acquire the water they need to do so. Given the fact that farmer's incentives are inconsistent with the objectives of the project designers generates a potential for sub-optimal performance, a potential that is all too often realized.

Mollinga (1998) conducts a detailed analysis of the lower left bank main canal of the TBP. His perspective emphasises the many different interconnecting aspects of the irrigation project, and the tendency of analysts to suppress too many of the details. Mollinga's work emphasises the political, economic, and social relationships that are created by the technology of irrigation. In the TBP, legally sanctioned cropping patterns, known as 'localisations' specify what each parcel of land can be used for. The localisations are intended to maximise the overall productivity of the system. On paper, irrigation staff have the authority to impose severe sanction against those farmers who violate their localisation. Unfortunately, localisations seldom correspond with farmers' profit maximising crops. Given limited irrigation staff with little credible authority, farmers fail to follow the specified localisation pattern, generating a grossly inequitable distribution of water.

The pattern of water distribution in large irrigation projects is often very unequal. For a set of irrigation projects in India, Deshparde and Supe (1989) find that head end farmers tend to use about 1.5 times as much water as is required for the crops being grown, while those in the tail end are unable to acquire the water they need. For a sample of Indian irrigation projects, Saini et al. (1989) find that farmers in the head end of the system grow more water-intensive crops than those in the tail region. Farmers in the tail end 'insure' against the problems of irregular water supplies by adjusting inputs such as seeding density and fertiliser. Kijne (1996) finds that protective irrigation projects in Pakistan have similar problems. In some cases, as much as four times the planned area of water-intensive crops (rice in particular) is grown. For the TBP left bank main canal, Hugar (1997) finds that farmers in the head region are growing four times as much paddy as specified by their localisations. Even in the tail region, considerably more paddy rice is grown than specified, consuming water which is then not available for others in this region.

In a previous paper, Janmaat (2001) determines the optimal policy for a stylised irrigation system, given constraints on enforcement options. It is shown that the best way to manage irrigation water is to distribute it such that any distributary which receives water receives enough for all farmers to grow any crop they chose and have water left over to flush salts from the system. To enhance equity, water supplies should be alternated between distributaries. Although not attaining the first best solution, this approach contains soil degradation problems.

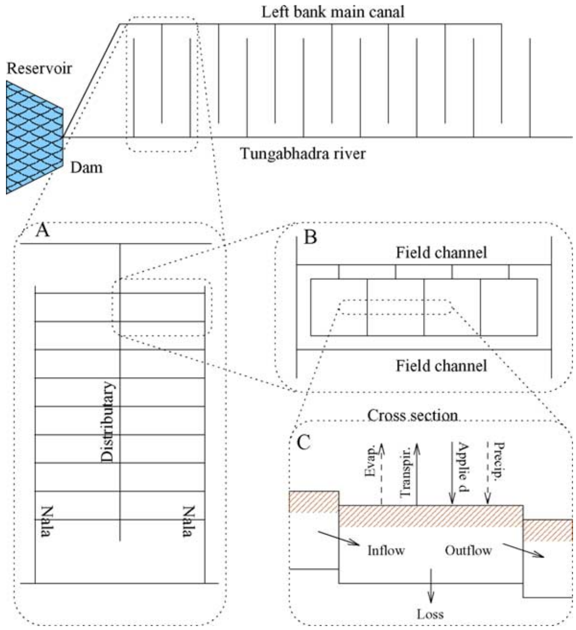

Fig. 1. Stylised representation of the TBP left bank main canal. The canal supplies water to a number of distributaries. The distributaries supply water to a number of field channels (A). Each field channel supplies water to a number of fields (B). The fields are related by a number of water sources and sinks (C).

3. The model

An irrigation system can be envisioned as a hierarchical system designed to distribute water from the main reservoir to each block of land through a network of canals. Fig. 1 shows a representation of the network of canals that distributes water from the main reservoir behind the Tungabhadra dam to the individual blocks of land. The backbone of the delivery system is the main canal. It transports water from the reservoir to a series of distributaries. These distributaries generally follow the height of land between natural drainage channels known as *nalas*. Water flows from the distributaries into field channels, from where the water flows to the actual fields.

The modelling exercise is essentially the reverse of the system description given above. The fundamental unit of action in the model system system is a block of land 100 m on a side, a one hectare square. The crop chosen determines the profit that will be generated by each cell, and the amount of excess water which must be accounted for. Individual cells are joined together as field channels. Field channels are collected into distributaries, and the distributaries are joined together as an irrigation system. At each higher level, changes in the patterns of water distribution can affect the performance of the system.

The total area that can receive irrigation from the left bank main canal (LBMC) of the Tungabhadra project is 243,915 ha. The total length of the canal is 227 km (Tungabhadra Board, 2000). The bulk of the irrigation water is released between the cities of Gangawati and Raichur. These cities are approximately 112 km apart. Gangavati lies approximately 32 km from the Tungabhadra dam (Bartholomew, 1995). There is little irrigable land between Gangavati and the dam, as the valley is quite narrow here. It is therefore assumed that the bulk of the irrigable land lies along a linear distance of about 120 km. Given the total area of the project, the implied average linear length of a distributary is about 20 km. There are 87 distributaries in the system (Mollinga, 1998). This implies an average distance between two distributaries of 1.38 km, which suggests that the average distance between a distributary and a drain is about 700 m. For the model, each field channel contains seven one hectare cells, each distributary has 400 field channels, 200 on each side of the distributary. The entire system has 87 distributaries. Together this generates an area of 243,600 ha.

The LBMC is allocated a total of about 26 M ha cm[1] of water (Tungabhadra Board, 2000). In principle, this system could provide almost 107 cm of water to each hectare of land in the command area, assuming no transmission losses. However, in most cases, this is not happening. For the 10 year period between 1984/85 and 1993/94, as reported by Hugar (1997), an average of just under 82,000 ha were irrigated per season. This severe under-performance is a result of water use patterns that deviate substantially from those planned for. In particular, water-intensive crops are commonly grown in violation of localisations. These high water crops are grown both in the kharif (monsoon) and rabi/summer (post-monsoon) seasons, further increasing the demand water. With paddy farmers utilising more than 300 cm of water for back to back paddy crops, and actual delivered supply coming in at close to 21 M ha cm, there is enough water distributed to produce paddy continuously on about 70,000 ha.

3.1. The soil column

3.1.1. Salt and water balance

The time unit for this model is 1 year. With this resolution, water logging is taken to be a within period phenomena, while salt accumulation and leaching occurs between periods. Salt and water balance are modelled in each soil column. The details of the water and salt balance approach can be found in van Hoorn and van Alphen (1994). The total available water in the soil column, A, is the sum of the net inflows (inflow from the uphill cell, L, precipitation, P, and applied water, I).

$$A = L + I + P.$$

[1] One ha cm of water is enough water to supply one centimetre of water to one hectare of land, one hundred cubic metres.

The salinity of this available water, measured as electrical conductivity, is,

$$\mathrm{EC}_A = \frac{(L \cdot \mathrm{EC_L} + I \cdot \mathrm{EC_I} + 2W_{\mathrm{fc}} \cdot \mathrm{EC_0})}{A},$$

where $\mathrm{EC_L}$, $\mathrm{EC_I}$, and $\mathrm{EC_0}$ are the electrical conductivities of the lateral inflow, the irrigation water, and the soil in this soil column before any irrigation occurs. W_{fc} is the water content of the relevant region of the soil column, the root zone, at field capacity. The electrical conductivity is measured using a saturation extract from a soil sample. The volume of the saturation extract is typically about twice the field capacity, which explains the factor 2.

The available water in the soil column leaves the column either as evaporation, transpiration, lateral outflow, or loss. Loss represents water that leaves by being flushed into a drainage canal or percolating to the groundwater table at a depth where it no longer has any impact on plant growth. The sinks for the available water can be partitioned as those which involve liquid outflows and those which involve gaseous outflows. If we let X represent the water that is excess to evaporation and plant transpiration, then;

$$X = A - T(A, \mathrm{EC_0}) - V(A - T(A, \mathrm{EC_0})).$$

$T(A, \mathrm{EC_0})$ represents the transpiration water use by the crop and $V(A - T(A, \mathrm{EC_0}))$ is the water lost to evaporation. The amount of water use depends on the salinity of the soil as well as the amount of available water. The amount of water that evaporates depends on the amount of water left after transpiration demand by the crop. This simplification ignores the impact of irrigation and precipitation timing on the salinity experienced by the crop. Judicious timing of irrigation and fortunate timing of precipitation can drive salts down out of the root zone, permitting crop growth in soils where it would not otherwise occur. The model construction also assumes a degree of separability between transpiration and evaporation that may be unrealistic.

Evaporation and transpiration do not remove appreciable amounts of salt from the soil column. Salt leaves dissolved in liquid water. The concentration of salt in any excess water is,

$$\mathrm{EC}_X = \frac{\mathrm{EC}_A A}{X},$$

provided that X is greater than zero. If X is equal to zero, then outflow is not an issue. The salinity of the root zone after the crop year is over is,

$$\widetilde{\mathrm{EC}}_0 = \frac{\mathrm{EC}_A A}{2W_{\mathrm{fc}} + X}.$$

If $A = 0$, then $\widetilde{\mathrm{EC}}_0 = \mathrm{EC}_0$. This relation defines the dynamics of the model. The salinity in the root zone at the end of the crop year, $\widetilde{\mathrm{EC}}_0$, is the salinity at the beginning of the next crop year. Thus, in steady state, for all levels of A, $\widetilde{\mathrm{EC}}_0 = \mathrm{EC}_0$. The convergence of the system to a (possibly cyclical) equilibrium is guaranteed by the shape of the crop response function developed below. The dynamics of the convergence path, and the relationship of this path to the starting conditions and parameter values, are not considered. An investigation of the impacts of varying these conditions is left to further research.

For simplicity, details of the water table level are ignored. Instead, water is tracked from cell to cell. Therefore, for the next cell downhill, the inflow L is equal to $(1 - \rho)X$, where ρ measures the share of the excess water in this cell that leaves the system. The smaller the value of ρ, the greater the water and the salt flowing into the downhill cell, generating an externality.

The results presented in this paper assume a root zone depth of 20 cm with a volumetric water content of 25% (Miller and Donahue, 1990). This generates a field capacity of $W_{\mathrm{fc}} = 5$ cm. Annual precipitation was 60 cm with annual evaporation also equal to 60 cm (Patil, 1997). The salinity of the irrigation water was taken to be equal to 0.2 dS/m (Patil, 2001).

3.1.2. Crop response functions

At low levels of soil salinity, as measured by the electrical conductivity, plants show no yield response to an increase in the level of salinity. However, after a threshold level is reached, most plants show an almost linear decline in yield with further increases in soil salinity (Dargan et al., 1982; Somani, 1991). Let the maximum yield attainable, assuming all other requirements are met, be $Y_{\max}$, a quantity that varies by plant. Let EC_T.be the threshold salinity level, beyond which yield is affected, and let m represent the percentage decrease in yield for a one unit increase in salinity. The yield response function can then be represented as,

$$Y_{\mathrm{pot}} = \begin{cases} Y_{\max} & \mathrm{EC}_0 \le \mathrm{EC}_T \\ Y_{\max}(1 - m(\mathrm{EC}_0 - \mathrm{EC}_T)) & \mathrm{EC}_0 > \mathrm{EC}_T. \end{cases}$$

The soil salinity is often measured at the end of the cropping season. The crop production component of this model is run to equilibrium, at which point the final salinity must equal the initial salinity, at least if the cell is subject to the same activity.

Application of irrigation water can act to dilute or flush salt out of the root zone. The electrical conductivity of the soil extract is approximately linearly related to the salt concentration of the soil extract (van Hoorn and van Alphen, 1994). This allows a dilution impact to be specified as $\mathrm{EC}_{\theta,\mathrm{eff}} = (A^*_\theta/A)\mathrm{EC}_0$ where A^*_θ is the recommended water application rate for crop θ. EC_0 is the electrical conductivity at the beginning of the crop season. $\mathrm{EC}_{\theta,\mathrm{eff}}$ is the effective electrical conductivity of the soil extract, as appropriate to crop θ, when A centimetres of water are available to the crop. However, this advantageous dilution effect must be balanced against the impact of excessive water on plant growth.

For crop water use, a linear relationship between yield and transpiration is commonly used (Doorenbos and Kassam, 1979; Ragab, 1996), up to the maximum yield. The response of a crop to applied water is taken to be curvilinear (Solomon, 1983; Perry and Narayanamurthy, 1998). Unfortunately, few parameterisations of curvilinear water response functions are available, especially for cases in which water is applied in excess of crop needs. Therefore, a piecewise linear water response function is used. The water response can be expressed in a function as,

$$Y = \begin{cases} 0 & A \le \underline{w}^{\mathrm{L}} \\ Y_{\mathrm{pot}}(A - \underline{w}^{\mathrm{L}})/(\bar{w}^{\mathrm{L}} - \underline{w}^{\mathrm{L}}) & \underline{w}^{\mathrm{L}} < A \le \bar{w}^{\mathrm{L}} \\ Y_{\mathrm{pot}} & \bar{w}^{\mathrm{L}} < A \le \bar{w}^{\mathrm{H}} \\ Y_{\mathrm{pot}}(1 - (A - \bar{w}^{\mathrm{H}})/ \\ \quad (\underline{w}^{\mathrm{H}} - \bar{w}^{\mathrm{H}})) & \bar{w}^{\mathrm{H}} < A \le \underline{w}^{\mathrm{H}} \\ 0 & \underline{w}^{\mathrm{H}} < A. \end{cases} \tag{1}$$

The potential yield of the crop is Y_{pot}, which is determined by the soil salinity. If applied water is below $\underline{w}^{\mathrm{L}}$, yield is zero. Yield increases linearly until

$\underline{w}^{L}$ is reached, at which point it plateaus. When water application increases beyond $\bar{w}^{L}$, yield falls linearly, attaining zero at $\underline{w}^{H}$.

Water excess to plant needs was calculated as,

$$X_C = \begin{cases} A & y = 0 \\ A - \bar{w}_{\theta}^{L}(Y/Y_{max}) & y > 0, \end{cases}$$

where X_C is the excess water relative to crop demands. When yield is zero, the crop does not use any water, so that all available water is excess. The amount of water available for leaching depends on the amount of evaporation which occurs. This excess is calculated as,

$$X = \begin{cases} 0 & X_C \leq e \\ X_C - e & X_C > e, \end{cases} \quad (2)$$

where e is the amount of evaporation from the soil and X is the excess water available for leaching. These relationships are shown in Fig. 2. The linearities in the model are strong assumptions. Crop water use is better modelled as curvilinear, as estimated by Solomon. Water evaporation better modelled as a non-linear process, dependent also on pore spacing and soil air flow. However, there was insufficient information available to meaningfully calibrate these more complex models. When calibrated to match the general results of the linear model, the more sophisticated forms added little to the final results.

The crop water use and salt leaching relationship combine as a system that will converge to an equilibrium for a constant amount of available water in each year. Suppose that the amount of water available to the crop is less than that which the crop can use in a particular crop year. Any salts contained in the water will accumulate, increasing the root zone salinity in the next crop year. With this accumulation, crop yield will fall. If the amount of water available to the crop remains constant, then water will come available for leaching as the amount used by the crop falls. Therefore, the rate increase in salinity from year to year will fall, and if leaching is sufficient, salinity will decline. The system will stabilise when the salt brought into the site exactly matches the salt that leaves by leaching.

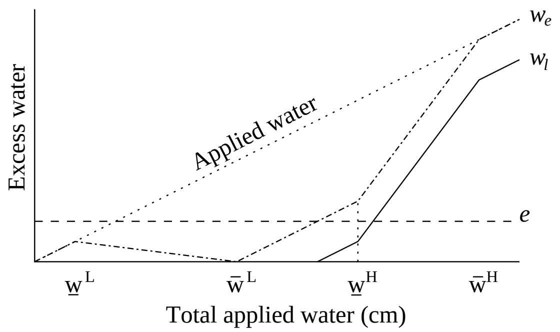

Fig. 2. Excess water and leaching water. w_e is the excess water that is not used by the plant. e is the amount of water that can be evaporated. w_l is the amount of water available for leaching, which is the excess water after evaporation.

Table 1
Critical values for salinity response function

Crop	Threshold (dS/m)	Slope (%/dS/m)
Cotton	7.7	5.2
Paddy	3.0	12
Jowar	6.8	16
Groundnut	3.2	29
Sunflower	2.5	10

Source: Dargan et al. (1982), Somani (1991).

For the principle crops grown in the TBP, the threshold and slope values are shown in Table 1, where salinity is measured as electrical conductivity in units of decisiemens per meter (dS/m). Among the listed crops, paddy, groundnut and sunflower are quite sensitive to salinity levels, while jowar and especially cotton are tolerant of higher salt levels. However, paddy is a bit of an anomaly. Although it is quite sensitive to salinity, this crop also has the ability to grow in ponded fields. The high water level can dilute the salts in the root zone and allow the crop to grow where it would otherwise be unsuitable.

The specific levels for the water thresholds in the plant growth model vary by plant species. The values used were approximated using the difference between projected yields with irrigation and under rain fed conditions[2]. In the TBP area, the normal annual rainfall is about 600 mm per year. Assuming about 500 mm of that falls during the kharif season, this was used as a water level for rain fed crops. The water available to irrigated crops in the kharif season was assumed to be this amount plus the recommended ir-

[2] Production recommendations were taken from an extension publication written in Kanada, the local language in the state of Karnataka. Translation for the purposes of referencing the document was not possible.

Table 2
Critical values for the water response function

Crop	Maximum yield	w^{L}	w^{H}	$\bar{w}^{\mathrm{H}}$	$\bar{w}^{\mathrm{L}}$
Cotton	30	30	125	155	245
Paddy	80	80	120	160	480
Jowar	60	0	90	130	170
Groundnut	30	10	110	130	170
Sunflower	12	40	80	110	140

Table 3
Crop price and cost function parameters

Crop	Price	Marginal cost	Fixed cost
Cotton	1651.91	399.52	5719.33
Paddy	371.31	100.00	4500.00
Jowar	270.00	70.26	3290.93
Groundnut	920.03	200.00	6000.00
Sunflower	1009.98	333.33	2000.00

rigation volume. The recommended amount was assumed to generate the maximum possible yield. Excepting paddy, which is grown on a flooded field, the yield potential was assumed to fall back to zero before the water volume recommended for paddy was realized. The threshold values and maximum yield for the crops modelled are shown in Table 2.

The precise details of farm level water management have been ignored. It is acknowledged that method and timing of water application can strongly influence crop productivity and the salinity conditions. These effects are beyond the scope of this analysis.

3.1.3. Profit function

The dynamic nature of the profit maximisation decision is not explicitly modelled, and risk issues are also ignored. The farmer in this model is a myopic profit maximiser. Further, each site is taken to be controlled by an individual who has no control over decisions at other sites. To maintain tractability and limit the shortcomings of the myopic nature of the model, the water application levels available were limited to three, 300 cm/ha, 90 cm/ha and rain fed. The high water application level is in line with the quantity of water farmers use when growing two paddy crops per year. 90 cm/ha is approximately equal to the amount of water that would be available if all irrigable land actually received water. Rain fed is the fallback decision when there is no water available. The farmer's optimisation problem is to choose the crop(s) to grow and how to distribute the available water between these crops. This problem can be written as

$$\pi(A, EC_0) = \max_{\theta_1,\theta_2,A_1} \left\{ \begin{array}{l} p_{\theta_1} y_{\theta_1}(A_1|EC_0) - c_{\theta_1}(A_1|EC_0) \\ +p_{\theta_2} y_{\theta_2}(A - A_1|EC_0) - c_{\theta_2}(A - A_1|EC_0) \end{array} \right\}$$

where p_{θ_i} is the price for crop θ_i, $y_{\theta_i}(A_i|EC_0)$ is the yield function, incorporating the salinity and water availability issues described above. The cost function is assumed linear, so that $c_{\theta_i}(A_i|EC_0) = FC_{\theta_i} + VC_{\theta_i} \cdot Y_{\theta_i}$, where FC_{θ_i} and VC_{θ_i} are fixed and variable cost components respectively. This specification models the farmer as choosing the best cropping plan each year, given the amount of water that is available. What control the farmer does have over water availability consists of moving water between the rabi and kharif seasons. The farmer is also assumed to know the salinity of the soil, ensuring that the crop chosen is optimal for both the amount of water available and the salinity of the soil.

For the crop price, the gross return divided by the yield was used, as reported by Hugar (1997). The use of gross return, in contrast to grain price, is an attempt to incorporate the value of other crop products not normally reported as part of the yield. The costs presented by Hugar are generally consistent with the expected direction of a linear cost function. The reported costs were totals. However, totals were presented for different seasons and regions, and for different yields. A linear relationship was extracted from Hugar's numbers. When this failed to produce a meaningful result, rough estimates were made using the change in recommended input quantities corresponding to different expected yields. The price and cost estimates are shown in Table 3.

Farmer responses to soil salinity are more sophisticated than the model suggests. Thiruchelvam and Pathmarajah (1997) document some of the control options that the farmers have, and their attitude towards these remedial practices. These measures are beyond the scope of the current analysis.

In the TBP, from the listed crops, cotton requires both the kharif and rabi seasons for one crop. The other crops can be grown in combination. Ten possible cropping combinations were considered. They are:

(1) Fallow–fallow,
(2) Cotton,
(3) Paddy–paddy,
(4) Paddy–sunflower,
(5) Jowar–sunflower,
(6) Groundnut–sunflower,
(7) Paddy–fallow,
(8) Jowar–fallow,
(9) Groundnut–fallow, and
(10) Sunflower–fallow.

Except for the maize–sunflower pair, this list contains the main cropping patterns reported by Hugar (1997). Using the prices and costs along with the yield functions permits the calculation of profits for each salinity level and each single crop. The profit for each cropping combination over a range of water and salinity levels was then determined. Once profits were calculated for each cropping combination at each salinity and water level, they were compared to determine which cropping pattern generated the highest profit at each salinity and level of applied water. This profit/water/salinity relationship then became the basis of the externality cost estimation.

3.2. Field channel

The field channel links together a number of soil columns. It was assumed that lateral flow of water only occurs between soil columns that are along the same field channel. In effect, the land that is irrigated by one field channel is assumed to lie along the fall line for the region, and little flow is assumed to occur between adjacent field channels.

Water available in the field channel was used first by the uphill farmers. Crops that were grown with heavier volumes of water were grown above those that used lesser volumes of water, and these were grown above those that relied on rainfall. This particular allocation reflects that which is commonly observed. It is also the allocation that has the potential to generate the greatest externality problems.

3.3. Distributary

The distributary canal transfers water from the main canal to a series of field channels. The distributary generally follows the height of land between two natural drains. Since the distributary is normally much larger than the field channels, water is able to travel along its length faster than it can travel along the field channels. Together with the tendency for more water to be used at the top of a field channel that at the bottom, this sets up a pattern of water use as shown in Fig. 3. The area devoted to heavily irrigated crops occupies a triangular wedge the base of which is parallel to the main canal, and the length follows the distributary. This wedge is contained inside a wedge of lightly irrigated crops. The remainder of the land is devoted to rain fed crops.

The parameter δ is used to represent the share of the total water available in the system which is used to grow lightly irrigated crops. Let I_l and I_h represent the number of centimetres of water per hectare per year applied to lightly and heavily irrigated crops, respectively. If a_l and a_h are the areas devoted to lightly and heavily irrigated crops, then the total water used in the distributary is $I_h a_h + I_l a_l = I_t$, where I_t is the total water released into the distributary. Since $\delta = I_l a_l / I_t$,

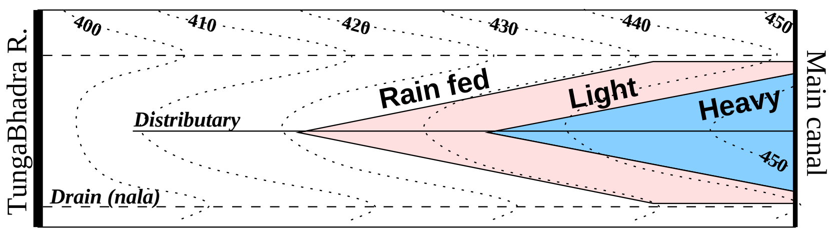

Fig. 3. Stylised distributary. Contour elevations are representative of those in the neighbourhood of Gangawati, in the head region of the area serviced by the TBP left bank main canal. The 'Heavy' region represents area devoted to water-intensive crop. The 'Light' region represents area devoted to lightly irrigated crops. The remainder of the area is devoted to rain fed crops.

the total area devoted to heavily irrigated crops is calculated as $a_h = I_t(1-\delta)/I_h$.

The parameter γ is used to represent the ratio between the distance along the distributary and the distance along the field channels that water flows. The wedge containing the heavily irrigated crops is therefore a triangle with height given by $(2a_h/\gamma)^{1/2}$. The total volume of water available to grow heavily irrigated crops was determined by building a 'staircase' of one hectare cells that rested on the distributary. Each step in the staircase was γ blocks wide. The lightly irrigated area was determined by building another staircase on top of the water-intensive staircase, using the same step widths.

When the amount of water released into the distributary was more than enough to irrigate all the land, for a given value of δ, the water was redistributed to minimise waste. The water-intensive wedge was expanded until the combined area for both heavily and lightly irrigated crops was equal to the total area in the distributary. In this case, the area for each type of crop is calculated by solving the pair of equations $a_l + a_h = a_t$ and $I_l a_l + I_h a_h = I_t$, where a_t is the total land area in the distributary.

Hugar (1997) regressed paddy stress days on distance from the dam along the main canal, distance from the main canal along the distributary, and distance from the distributary along a field channel. He found that the distance from the distributary has a much larger impact on the number of stress days than either the distance from the dam or the distance along the distributary from the main canal. For every 100 m along a field channel, the likelihood of experiencing a stress day grows by a factor of 23.8. After rounding, the value of 24 is used for γ in the model calculations.

3.4. The system

Adding the main canal completes the model. The main canal delivers water from the dam to the distributaries. The pattern of water releases into the individual distributaries selects a particular productivity for that distributary, given values for δ, ρ, γ and the biophysical characteristics.

The pattern of water use across distributaries is represented by the parameter ϕ, shown in Fig. 4. As the angle ϕ increases from 0 to 90°, the distribution of water goes from exactly equal over all distributaries to a

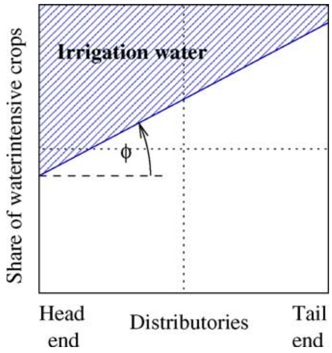

Fig. 4. Distribution of water-intensive crops between the head end and the tail end of the irrigation system.

pattern where all the water-intensive crops are grown in the head region and only rain fed crops are grown in the tail. The amount of water that is received by a particular distributary is determined by the interaction of this angle with the share, δ, of the applied water which is used to grow water-intensive crops. The maximum amount of water a distributary can receive is set such that the water use is consistent with the value of δ.

Fig. 5 shows the water distribution pattern for four different combinations of the share devoted to water-intensive crops, δ, and distributary allocation angle, ϕ. If $\delta = 1$ or $\phi = 0$, then each distributary receives an equal share of the water. When $\delta = 1$, water is being used exclusively for lightly irrigated crops. Since there is enough water available to irrigate all the land in this way, the angle is irrelevant. Likewise, when $\phi = 0°$, equal volumes of water are being released into each distributary. When $\delta = 0$, this water

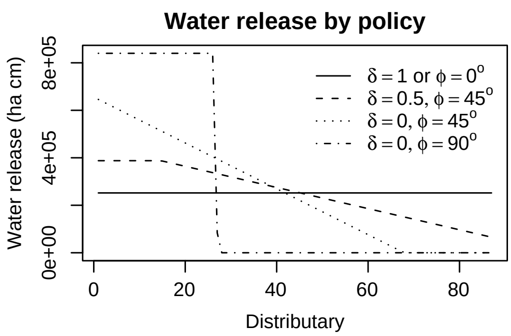

Fig. 5. Distribution of water over distributaries as a function of the share of water devoted to the lightly irrigated crop, δ, and the angle determining how water is distributed across the distributaries, ϕ.

will be used to grow water-intensive crops in that distributary, rather than lightly irrigated crops. Thus, much of the distributary will not be irrigated. However, each distributary will be equal. When $\delta = 0.5$ and $\phi = 45^\circ$, the first 15 distributaries receive enough water to devote half of it to water-intensive crops and the remainder to lightly irrigated crops. The total water required to fully irrigate these distributaries is greater than the amount needed to supply lightly irrigated crops on all the land, which is indicated by the height of the land. Beyond the first 15 distributaries, the amount of water released into each distributary declines as the distance from the source increases. When $\delta = 0$ and $\phi = 45^\circ$, less than 70 of the 87 distributaries receive water. No distributary receives enough water to grow the water-intensive crop on all the land in the distributary. As we increase the angle, the water release becomes more concentrated in the first distributaries. When $\delta = 0$ and $\phi = 90^\circ$, all the water is released into the first 27 distributaries. None of the other distributaries receive any water. Note that the area under each line is the same, reflecting the fact that adjusting δ and ϕ is a matter of adjusting the distribution of a fixed total amount of water.

The parameters δ and ϕ can be interpreted as policy options. Choosing ϕ involves determining how much water is to be released into each distributary. Choosing δ reflects the control that the command authority has below the distributary outlet. The localisation pattern that underlies the design of the system assumes that $\delta = 1$. The authority is able to distribute the water to achieve the goal of spreading the water thinly over a large area. As δ moves towards zero, the degree of control that the authority has over the water use at the field level declines. As a result of this loss of control, the authority can compensate by choosing an optimal value for ϕ. The results below illustrate the value of different policy options under a range of system conditions.

4. Results

The simulation was built using a nested set of objects coded in C++ and executed on a dual Pentium 933 MHz system running Linux. Four cases, $(EC_{IRR}, \rho) \in \{(0, 1), (0, 0.75), (0.2, 1), (0.2, 0.75)\}$, were evaluated. Each of the ranges $\delta \in [0, 1]$ and $\phi \in [0^\circ, 90^\circ]$ were divided into 100 intervals, generating a grid with 10,000 points. Total system profit was calculated at each of these points. Total computer time was approximately 60 h. To reduce the computation time to 60 h, two main shortcuts were used. First, the crop production function was calculated once for a range of values. All subsequent crop production values were drawn from this initial calculation. Linear interpolation was used to generate values when one had not been calculated. Second, field channel results were stored in an indexed collection such that if the same cropping pattern was used for a channel, it was not recalculated. This latter step saved considerable time, as on occasion a field channel required up to 1000 simulated years worth of iterations to reach a reasonably stable equilibrium, in terms of soil salinity and profit.

Each object was individually tested before being bundled into the next object in the hierarchy. This permitted the nature of the relationships which conspired to yield the final results to be illustrated as they were integrated. Graphing of the results was performed using the statistical and graphing language R (Ihaka and Gentleman, 1996).

4.1. Production function

The cropping pattern generated by the model is dominated by two principle choices, cotton and paddy on paddy. Given the model parameters, cotton is the optimal crop when water supplies are limited. It is able to tolerate high salt concentrations compared with other cropping options. The second most popular crop is paddy. When salinity levels are low, it can be double cropped on the same land. When salinity levels are high, it is grown as a single crop using a very high volume of water to dilute the salts. The remaining crops, groundnut and jowar, make their presence felt at salinity and water availabilities that fall between the two dominant crops. For the revenues and costs reported by Hugar (1997), cotton and paddy after paddy are the most profitable crops. Paddy on paddy is slightly more profitable, if soil conditions are excellent. When soil conditions are less than ideal, cotton becomes the most profitable crop. However, cotton is sensitive to water availability. If there is too much water available, perhaps as a result of water logging, the returns to growing cotton fall rapidly.

4.2. The system

As modelled, the irrigation system is a collection of 87 distributaries. Total profit was determined as a function of the share of water devoted to lightly irrigated crops and the angle of the water release pattern into the distributaries. The total water available could provide 90 cm of water to all the land in the system. In the absence of any salinity or water flow issues, total profit declines as the share of water devoted to lightly irrigated crops declines. The maximum profit of 7.76×10^9 Rs occurs when all the land is irrigated at 90 cm. When all the water is used to grow the heavily irrigated crop, total profit falls to 3.36×10^9 Rs, a value that is 43.4%of the maximum possible. Therefore, almost 57% of the profit that could be generated by the irrigation system is lost as a result of the inefficient use of the irrigation water, without any contribution from salinity or water logging problems. This degree of loss is in line with the performance observed by Chambers (1988) (Fig. 6).

The percent of the maximum possible profit that is lost to water logging is shown in Fig. 7. The maximum water logging cost is 3.03%, which occurs when each distributary receives an equal amount of water and approximately 85% of the water is used to grow the heavily irrigated crop. Under these conditions, the lost profit due to excessive water availability in some

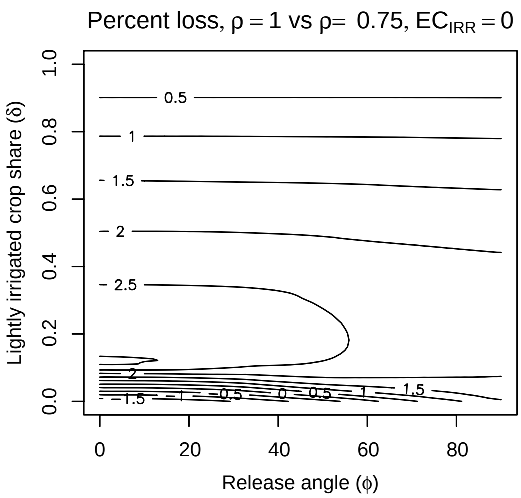

Fig. 6. Percent of maximum possible profit lost when $\rho = 0.75$.

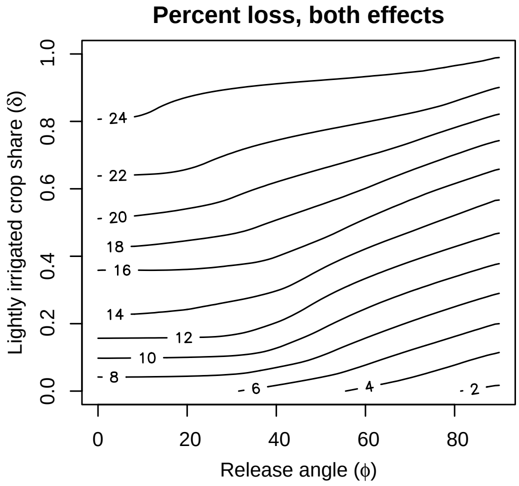

Fig. 7. Percent profit loss, relative to maximum possible profit, when irrigation water contains dissolved salts and all excess water does not leave the system from the sight of application ($\delta = 0.2$, $\rho = 0.75$).

fields exceeds the increased profit resulting from subsurface irrigation in other fields. As the share of the water used to grow heavily irrigated crops further increases, the extra profit resulting from subsurface irrigation increase. Rather than flowing to land growing the lightly irrigated crop, which is susceptible to excessive water application rates, it is flowing to land that receives no irrigation water. When all the water is used to grow heavily irrigated crops, and the amount of water released into each distributary is equal, then total profit is increased by 1.93%, relative to the profit which would be earned when $\phi = 0$ and $\delta = 0$ without lateral flow. If we increase the angle, then the amount of land receiving unintended subsurface irrigation falls, and the water logging cost being borne by the heavily irrigated crops comes to dominate. With $\phi = 90$ and $\delta = 1$, the profit loss is 1.41%. When the applied irrigation water contains no dissolved salts and farmers are able to optimise in the face of lateral inflow, the cost of water logging is quite small.

Adding salinity to the model leads to a large increase in the profit lost. Fig. 7 shows the percent profit loss, relative to the maximum possible, when both salinity and water logging problems occur. With both salinity and water logging problems, maximum profit falls to 5.89×10^9 from 7.76×10^9 Rs per year, for a loss of 24.11%. When the applied water contains dissolved salts, salts accumulate until yield for the lightly irrigated crop falls enough for leaching to bring the cell to equilibrium. The excess water, with its enhanced salt concentration, flows into the next downhill cell. There it combines with the salt in the applied water, causing further profit losses. With salinity increasing as one moves down the hill, losses suffered are greater than when local salt accumulation is the only issue. The combined effect is a substantial reduction in total system profit when the irrigation water is being distributed equally.

When the irrigation water is all used for heavily irrigated crops and the entire quantity of water is released in the first distributaries ($\delta = 0$, $\phi = 90°$), the combined water logging and salinity costs are minimal. The total profit in the absence of these problems is 3.36×10^9 Rs per year, while with these problems it falls to 3.25×10^9 Rs per year. The loss is 1.48% of the maximum possible profit. This cost occurs because some fields receive such high quantities of water that even the high water crop suffers. In all cases, salt accumulation is not a problem on irrigated land. For the land that is not irrigated, there are no uphill fields growing water-intensive crops which can send excess water with high salt levels into the fields growing rain fed crops. When this is the case ($\delta = 0$, $\phi = 0°$), the total profit falls from 3.36×10^9 Rs per year to 2.86×10^9 Rs per year, a loss of 6.51%. This loss represents a total abandonment of all the un-irrigated land. However, since the rain fed land earns a much smaller profit than the irrigated land, this loss is small in percentage terms.

Even in the presence of water logging and salinity problems, the profit that results from distributing the available water equally between the farmers is greater. Fig. 8 shows the actual profit that is generated by the system, in the presence of water logging and salinity problems. The difference between the worst case and the best case is now 51.4% compared to 56.6%, indicating that in percentage terms, the returns have fallen by almost the same amount at both the high and low ends.

The results shown in Fig. 8 support the conclusions of Janmaat (2001). In that analysis, the author considered a case in which farmers are able to choose the amount of water to use. The farmers near the outlet were able to take water before those further down the hill. The result was that the uphill farmers consistently

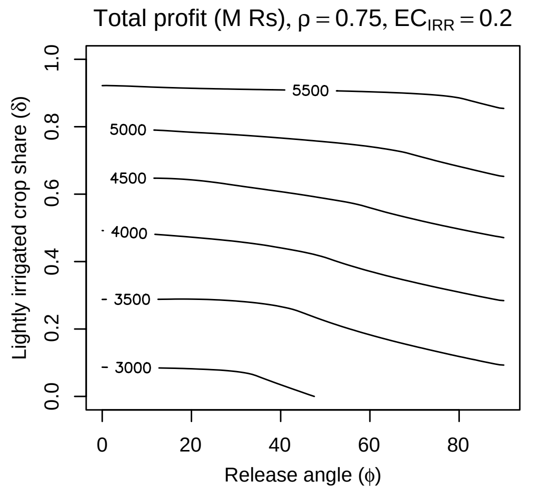

Fig. 8. Total profit in the presence of salinity and water logging problems.

use large quantities of water while the downhill farmers suffered the external costs. Janmaat's analysis implicitly assumed that $\delta = 0$ and investigated the effect of changing ϕ. That analysis considered the dynamic decision process of the farmer and permitted the command authority to choose a sequence of ϕ values that repeated over a number of years, rather than one policy for all time. However, Janmaat's earlier analysis used a highly stylised model from which quantitative cost estimates were difficult to generated. In this model, the details of the eternal cost process have been modelled with greater precision. The qualitative aspect of the earlier result carries through. If only constant policies are considered, the externality cost is minimised when the irrigation water is released such that those distributaries which receive water receive enough to grow the water-intensive crop on all the land within the distributary. Equity issues can be dealt with by rotating which distributaries receive water.

In the TBP, approximately half of the water delivered by the system is being used to grow water-intensive crops. This suggests that a value of 0.5 is appropriate for δ. In the head region, enough water is released into the distributaries that it is being flushed into the drains. However, all the land is not being used to grow heavily irrigated crops. In general, some water does reach the end of the system. This suggests that ϕ does not lie at the extremes of 0° or 90°. If we select a value of $\phi = 45°$, then the total

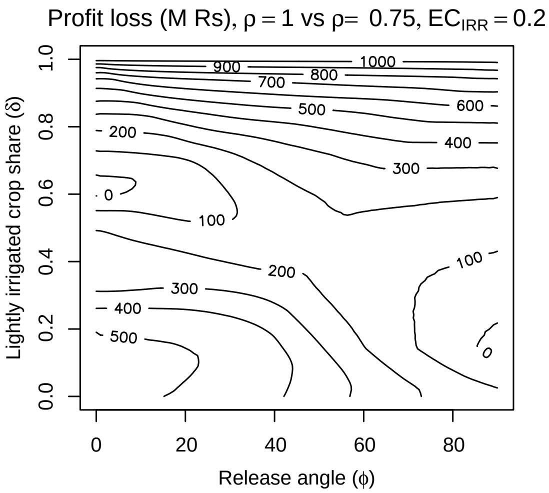

Fig. 9. Profit loss with saline irrigation water when lateral flow is introduced.

profit that could be realized if there was no salinity or water logging is 5.54×10^9 Rs per year. Relative to the maximum possible profit, this represents a loss of 28.6%. Adding salinity and water logging effects reduces total profit to 4.19×10^9 Rs per year, a reduction of 24.4%. Relative to the best possible case with salinity and water logging effects, profit is reduced by 28.8%.

In summary, adding salinity and water logging effects to the model has a large impact on aggregate profit. By comparing the situation with salinity but no lateral flow, the externality cost can be assessed. Fig. 9 shows the change in profit that results when the irrigation water is saline and lateral flow is introduced. These are the costs that could be avoided if drainage was installed to prevent the lateral flow of excess water between fields. The external cost is maximised when all the water is used to grow the lightly irrigated crop. If the release angle is held fixed at $\phi = 0^\circ$, costs fall with share until a minimum is reached near $\delta = 0.6$. The extra profit earned from unintended subsurface irrigation and leaching is sufficient to offset the cost of salinity-induced yield depression. As the share is further reduced, externality costs climb and reach a local maximum when $\delta = 0$. When the angle is fixed at $\phi = 90^\circ$, the effect is similar. However, reducing δ corresponds with a release of more water into the distributaries near the reservoir. This reduces

Table 4
Breakdown of costs by salinity, inefficient distribution, and water logging

Source	System profit (Rs)	Change (%)
$\delta = 0,\ \phi = 0, \mathrm{EC_{IRR}} = 0,\ \rho = 0$	7.76×10^9	n/a
$EC_{IRR} = 0.0 \rightarrow \mathrm{EC_{IRR}} = 0.2$	6.93×10^9	−10.7
$\delta = 0,\ \phi = 0^\circ \rightarrow \delta = 0.5,\ \phi = 45^\circ$	4.36×10^9	−37.1
$\rho = 1.0 \rightarrow \rho = 0.75$	4.19×10^9	−3.8

the area receiving subsurface irrigation and increases the volume received. The externality costs fall below zero close to $\delta = 0.2$, after which they again climb.

Table 4 reports the breakdown of the costs according to the different factors. Salinity of the irrigation water reduces equilibrium profit by almost 11%. The inefficient distribution pattern reduces equilibrium profit by more than 37%. The physical externality, lateral water flow, accounts for slightly less than 4% of the inefficiency, at the indicated distribution pattern. The representative policy for the TBP falls conveniently in a region where externality costs are low.

4.3. *Equity effects*

Fig. 10 shows the Gini coefficient calculated for the different water distribution patterns. The Gini coefficient ranges between zero and one, with zero representing equally distributed income. The maximum

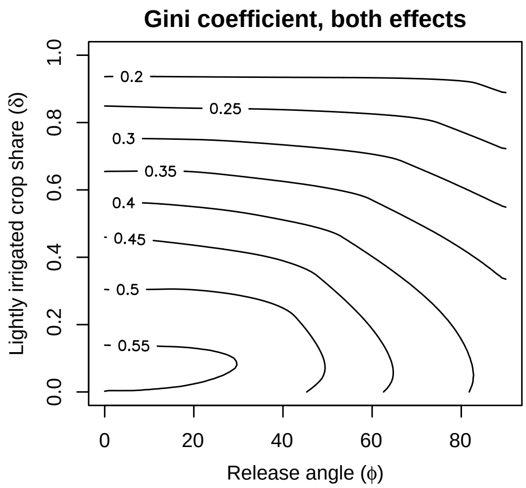

Fig. 10. Water distribution pattern equity effects as measured by the Gini coefficient.

inequality occurs where the distribution of water is equal between distributaries and almost all the water is used to grow heavily irrigated crops. There are two sources for this inequality. The largest contributor is the fact that as more water is devoted to growing water-intensive crops, more farmers are denied irrigation water. The second contributor is the presence of the salt accumulation externality. Salinity and water logging effects drive down the profit for lightly irrigated and rain fed land, while leaving heavily irrigated land virtually unaffected. This result is also in accord with earlier work by Janmaat. When there are no effective controls over farmer cropping decisions at the farm level, then δ is essentially zero. In such a situation, releasing the water equally between the distributaries reduces equity, as it forces all farmers who do not have access to water to suffer the damage of the externality. The best way to enhance overall equity is to rotate water availability between distributaries from year to year, rather than trying to ensure equal water releases into each distributary.

An interesting result of this analysis is the positive correlation between equity and efficiency. In general, cropping patterns that are less economically efficient are also cropping patterns that generate greater inequality. This is a consequence of the protective nature of the TBP. The design of a protective irrigation system aims to maximise the return per unit of water, not the return per unit of land. In the present context, return per unit water is maximised by distributing the water equally over the irrigable land. The greater the deviation from this equal distribution of water, the greater the efficiency cost and the greater the inequality.

5. Discussion

The analysis has relied on an integration of several different models and a number of assumptions to permit the models to be linked. Changing the assumptions and parameterisations may have a substantial influence on the final result. The impact of each of the main assumptions is considered in turn in the following:

The physical details of the area serviced by the TBP lower left bank main canal were simplified into a series of 87 equal sized strips of land that were 20 km long and 1.36 km wide. The actual system bears little resemblance to the simplified case. The length and width of the land serviced by a distributary interact with the pattern of water use to determine the area of land affected by salinity. Holding the general shape of the water use pattern constant, making the distributaries wider and shorter can substantially increase the area of land affected and the productivity cost.

The characteristics of the soil, the irrigation water, and the precipitation, determine how salt is accumulated in the system. An increase in the salinity of the irrigation water will require more leaching water to keep the system at equilibrium. This implies that the equilibria will occur at lower crop yields, generating a larger decline in productivity. A reduction in precipitation is akin to an increase in irrigation water salinity. Less precipitation induces less dilution of the irrigation water and less leaching of salts when there is no crop grown. The depth of the root zone and the field capacity of the soil interact to determine the nature of the solution into which the plant roots grow. Reducing the depth of the root zone and/or reducing the field capacity will increase the rate at which salts accumulate and the effective salinity of the soil solution that plants draw from.

The crop response functions establish how the crop responds to changes in salinity and the amount of water applied. Changing these characteristics of the biological system will change the shape of the yield functions and the amount of excess water available for leaching. The crop yield interacts with the crop price and the cost function to determine the profits to the farmer, and the crop choice which is optimal.

The economic characteristics determine the choices made by the farmer. Changing the price and cost function will change the optimal choice. Under the prices used in this study, two crops dominate in terms of profitability to the farmer. These crops are paddy and cotton. The prevalence of the first is consistent with the observed behaviour in the head region of the TBP (Hugar, 1997). The importance of cotton is consistent with the description of this crop as India's 'white gold' (The East India Cotton Association, 2001). However, the dominance of these crops in the model goes somewhat beyond the actual observed cropping pattern.

The assumptions that farmers are myopic profit maximisers and ignorant of risk is a gross simplification that can explain part of the discrepancy between the model cropping pattern and that actually observed.

Crops such as jowar can play an important part in feeding livestock and providing an assured yield if water is not delivered as expected. Water deliveries, particularly in the tail region of the project, can be quite erratic. Thus, farmers may hedge against this risk by growing jowar. The fact that jowar is a crop which is used on the farm also implies that the market price understates the value of this crop to the farmer. To the extent that the non-cotton crops are more sensitive to salinity than cotton, the calculated productivity loss is likely to underestimate the true cost.

The distribution of water within the TBP is strongly affected by social interactions and the history of those interactions. These aspects are not considered in any detail here. Mollinga (1998) explores these issues. An important aspect of this history is how it interacts with the economic conditions to accentuate inequality. To grow paddy, land must be levelled and bunds installed to permit maintenance of a flooded field. This is costly. Before the TBP project was constructed, this part of Karnataka used to grow rain fed crops. Crop yields, being dependent on unreliable monsoons, were highly variable. Farmers were mostly poor. Most of these farmers did not have the financial means to take advantage of the water provided by the TBP. Many of them sold their land to the few farmers who where wealthy, and to migrants from the delta of the Krishna river to the east. These farmers had experience with irrigated agriculture, and the financial resources to acquire the land that the local farmers were willing to sell. These farmers were able to improve the land, making them able to grow the more lucrative high water crops. Thus, the TBP project ended up providing a greater benefit to those farmers that were already well off than to the poor farmers which the project was ostensibly supposed to help.

6. Conclusion

The pattern of water use in the region serviced by the lower left bank of the Tungabhadra project suffers from salinity related soil degradation. To estimate the long run cost of this degradation, a simplified model of the interactions between farmers, crops, soil and topography was constructed. Depending on the pattern of water distribution throughout the system, inefficiencies may be as high as 51.6% of potential system productivity, in the absence of any salinity or water logging related externalities. For the TBP left bank main canal, where approximately half of the water is being used to grow water-intensive crops, the inefficient water distribution pattern is the largest contributor to the low system performance. This factor accounts for 37.1% of the maximum potential profit being lost. Irrigation water salinity reduces profit by a further 10.7%, while lateral water flow accounts for a relatively minor 3.8% profit cost, given the water distribution pattern. The vast majority of the inefficiencies observed in the TBP are the result of a system that is unable to enforce a pattern of water use that maximises overall profitability. A particularly striking result is that moves which would increase the efficiency of the TBP wold also increase the equity of income in the system. This is a case in which the pursuit of efficiency is not at odds with a principle objective of the TBP project, a fortunate coincidence which it is hoped will spur action.

Acknowledgements

This research was supported with funding from the Shastri Institute. I am particularly grateful to Dr. L.B. Hugar of the Agricultural University of Dharwad, Raichur Campus, and Dr. S.G. Patil of the Agricultural University of Dharwad, Gangawati Research Station. Dr. Gary Van Loon and Dr. Clive Chapple, of Queen's University, were also instrumental in supporting this research. Two anonymous reviewers provided comments that helped refine the presentation. Any remaining oversights and errors are solely mine.

References

Bartholomew J., (Ed.), 1955. The Times Atlas of the World, vol. 2. Houghton Mifflin, Boston.
Bhatia B.M., 1988. Indian Agriculture: A Policy Perspective. Sage Publications, New Delhi, Chapter 8, pp. 167–182.
Chambers, R., 1988. Managing Canal Irrigation: Practical Analysis for South Asia, Cambridge University Press.
Dargan K.S., Singh O.P., Gupta I.C., 1982. Crop Production in Salt-Affected Soils. Oxford Press, New Delhi.
Deshparde, D.D., Supe, S.V., 1989. Problems of farmers in canal irrigation. Indian Journal of Agricultural Economics 44 (3), 290.
Doorenbos J., Kassam A.H., 1979. Yield Response to Water. No. 33 Irrigation and Drainage Paper. Food and Agriculture Organization of the United Nations. Rome.

Hugar L.B., August 1997. Performance of canal irrigation system in tungabhadra project–an economic analysis. Ph.D. Thesis, University of Agricultural Sciences, Dharwad.

Ihaka, R., Gentleman, R., 1996. R: a language for data analysis and graphics. J. Comput. Graph. Stat. 5, 299–314.

Janmaat J., 2001. Water Distribution to Contain Soil Degradation: Sustainability and Equity Implications. Interdisciplinary Environmental Review.

Kijne J.W., 1996. Water and Salinity Balances for Irrigated Agriculture in Pakistan. Research Report 6, International Irrigation Management Institute, Colombo, Sri Lanka.

Miller R.W., Donahue R.L., 1990. Soils: An Introduction to Soils and Plant Growth, sixth ed., Prentice-Hall, Englewood Cliffs, NJ.

Mollinga P.P., 1998. On The Waterfront: Water Distribution, Technology and Agrarian Change in a South Indian Canal Irrigation System. Ponsen en Looijen, Wageningen, The Netherlands.

Patil S.G., 1997. Annual report of all india co-ordinated research project on management of salt-affected soils and use of saline water in agriculture. Tech. rep., University of Agricultural Sciences, Dharwad: Agricultural Research Station, Gangavati, Gangavati, Karnataka, India.

Patil S.G., 2001. Personal Communication.

Perry C.J., Narayanamurthy S.G., 1998. Farmer Response to Rationed and Uncertain Irrigation Supplies. Research Report 24, International Water Management Institute, Colombo, Sri Lanka.

Ragab R., 1996. Constraints and applicability of irrigation scheduling under limited water resources, variable rainfall and saline conditions. In: Proceedings of the ICID/FAO Workshop on Irrigation Scheduling. No. 8 in Water Report. Food and Agriculture Organization of the United Nations, FAO, Rome.

Rao C.H.H., 1994. Agricultural Growth, Rural Poverty and Environmental Degradation in India. Oxford University Press, Delhi.

Saini, A.S., Thakur, D.R., Moorti, T.V., 1989. Spatial differentials in canal water use: a case study. Indian Journal of Agricultural Economics 44 (3), 296.

Solomon K.H., 1983. Irrigation Univormity and Yield Theory. Ph.D. Thesis, Utah State University, Logan, Utah.

Somani L.L., 1991. Crop Production with Saline Water. Agro Botanical Publishers (India).

The East India Cotton Association, 2001. History of Cotton in India. http://www.ecottonindia.com/history.htm.

Thiruchelvam S., Pathmarajah S., 1997. An Economic Analysis of Salinity Problems in the Mahaweli River System h Irrigation Scheme in Sri Lanka. Research Report, International Development Research Centre, Ottawa, Canada.

Tsur Y., Dinar A., May 1995. Efficiency and equity considerations in pricing and allocating irrigation water. Policy Research Working Paper 1460, The World Bank.

Tsur, Y., Dinar, A., 1997. The relative efficiency and implementation costs of alternative methods for pricing irrigation water. The World Bank Economic Review 11 (2), 243–262.

Tungabhadra Board, 2000. Salient features of canals in tungabhadra irrigation system. Web page, http://www.tbboard.com.

van Hoorn J.W., van Alphen J.G., 1994. Salinity control. In: Ritzema, H. (Ed.), Drainage Principles and Applications, second ed., International Institute for Land Reclamation and Improvement, Chapter 15, pp. 533–600.

Vaux Jr., H.J., 1986. Economic factors shaping western water allocation. American Journal of Agricultural Economics 66, 1135–1142.

Vaux Jr., H.J., Howitt, R.E., 1984. Managing Water Scarcity: An Evaluation of Interregional Transfers. Water Resources Research 20 (7), 785–792.