

Timing Groundwater Protection: An Optimal Control Analysis of Annual Floors versus Terminal Targets for Irrigated Cotton in Lubbock County, Texas

Patrick Nyenkan

Patrick.nyenkan@ttu.edu

Texas Tech University

Research Article

Keywords: groundwater management, optimal control, irrigation, cotton, Ogallala Aquifer, Texas High Plains

Posted Date: October 12th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11311274/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Timing Groundwater Protection: An Optimal Control Analysis of Annual Floors versus Terminal Targets for Irrigated Cotton in Lubbock County, Texas

Department of Agricultural and Applied Economics, Texas Tech University, Lubbock, Texas
Patrick Nyenkan

Abstract

Groundwater protection rules can allow the same cumulative pumping yet differ in when water may be used. This study develops a discrete-time optimal control model of irrigated, dryland and fallow cotton in Lubbock County, Texas, to compare a year 20 stock target with an annual year-end floor. Production follows experimental irrigation-yield curves, costs follow 2024 extension budgets, and hydrologic benchmarks come from Texas Water Development Board modeling. The 20-year linear program is solved under perfect foresight for wet, dry and alternating production sequences. With a 90% stock target, the annual floor leaves cumulative pumping and mean acreage unchanged but lowers net present value by \$1.0 to \$13.5 million and raises the lowest modeled stock from 63–84% to 90% of the initial stock; before committed fixed costs the floor costs 0.2% to 2.7% of discounted returns. The shadow value of the stock constraint is about \$400 per acre-foot in the wet sequence and \$37 to \$39 in the dry and alternating sequences. County irrigated yields are about two-thirds of experimental plot yields, and county data cannot identify the yield response to water; rescaling yields removes irrigation from the dry sequence and, below a multiplier of about 0.68, from the other sequences. Across 2,000 random parameter sets, the gain from irrigation depends mainly on the yield multiplier, dryland yield and cotton price rather than on initial storage. The results are conditional scenario evidence, not a calibrated county forecast.

Keywords: groundwater management; optimal control; irrigation; cotton; Ogallala Aquifer; Texas High Plains.
JEL codes: C61, Q15, Q25.

Introduction

Irrigated cotton in Lubbock County depends on an aquifer that is drawn down faster than it is replenished. In the county series compiled for this study, observed Ogallala pumping averages about 166,000 acre-feet per year over 2000 to 2024, while the modeled net inflow benchmark is about 50,000 acre-feet per year [3, 8]. The TWDB reports 106,160 acres of irrigated cotton in the county in 2024 [8]. Whether farmers can keep irrigating, and what it costs them to protect the remaining stock, depends on how groundwater rules are written.

Most rules protect the stock in one of two ways. A terminal target requires the aquifer to reach a specified level at the end of a planning period and says nothing about the path. An annual floor requires the stock to stay above a threshold in every year. The two rules can permit the same cumulative pumping and the same ending stock, yet they imply different timing of irrigation, different discounted returns and different risk to wells in the middle of the horizon. The economics literature on groundwater has mostly compared optimal and competitive extraction [7, 10, 11], and has paid less attention to how the form of a stock constraint changes the use of water over time.

This paper studies cotton production in Lubbock County as a constrained optimal control problem. A planner chooses irrigated, dryland and fallow acreage and applied water each year for 20 years to maximize discounted returns, subject to a groundwater stock equation, an annual pumping ceiling, and either a year 20 stock target or an annual year-end floor. Production comes from experimental irrigation-yield curves [1], costs from 2024 extension budgets [5, 6, 9], and hydrologic benchmarks from Texas Water Development Board (TWDB) modeling [2, 3]. The model is a linear program solved under perfect foresight.

The paper makes four contributions. First, it compares terminal and annual-floor rules holding cumulative pumping fixed, so that differences come from timing alone. Second, it reports the shadow value of the stock constraint and the cost of conservation per acre-foot. Third, it tests how far the experimental curves can be tied to county data, finding that county irrigated yields are about two-thirds of plot yields and that county data cannot identify the response to water. Fourth, it propagates parameter uncertainty by re-solving the model for 2,000 random parameter sets and ranks the inputs by influence.

The results are conditional scenario evidence, not a county forecast or a pumping recommendation. They show that the annual floor lowers discounted returns by a small share relative to a terminal-only target, that the value of water depends on whether irrigation is profitable enough for water to be scarce, and that the yield transfer from plots to county fields matters more than the initial stock benchmark. The paper proceeds with the background literature, the methods, the results, a discussion and a conclusion.

Background and literature

Optimal control of groundwater

The economic analysis of groundwater extraction as a dynamic resource problem starts with Burt [10], who treated the aquifer as a stock and showed that optimal pumping equates the marginal return to water with its pumping cost plus the value of leaving water in the ground. Gisser and Sánchez [7] compared optimal control with competitive extraction for the Pecos Basin and found small gains from optimal management, a result that shaped the literature and is often called the Gisser-Sánchez effect. Feinerman and Knapp [11] examined the benefits of groundwater management under varying assumptions, and Koundouri [12] surveys the field. These studies mostly contrast optimal and competitive paths with stock-dependent pumping costs. The present model differs in imposing a stock requirement exogenously, so its shadow values measure the cost of a rule rather than the gain from coordination.

Risk and stochastic supply

Another strand adds uncertainty to the problem. Tsur and Graham-Tomasi [13] show that groundwater acts as a buffer against variable surface supplies, and later work on conjunctive management extends this to stochastic settings. Perfect foresight, as used here, ignores that role: it treats the wet and dry sequences as known in advance, so the results are best read as an upper bound on what a rule can achieve.

The High Plains and the Texas Southern High Plains

A large literature documents depletion of the Ogallala Aquifer and its economic consequences. Scanlon et al. [14] assess groundwater depletion and the sustainability of irrigation in the US High Plains. Hornbeck and Keskin [15] study how the aquifer has shaped agricultural adjustment over time, and Terrell, Johnson and Segarra [16] estimate the economic effects of depletion in the Texas High Plains. Recharge in the Southern High Plains is small relative to pumping, which is consistent with the inflow benchmark used here, and local management is carried out by groundwater conservation districts such as the High Plains Underground Water Conservation District, whose rules set annual per-acre limits [4]. This paper uses that district rule as a screening ceiling, not as a verified description of property-level rights.

Stock constraints, timing and spatial externalities

Most groundwater rules are written as stock targets or pumping limits. A terminal target, such as retaining a share of storage over a planning horizon, constrains cumulative pumping but allows temporary depletion, whereas an annual floor constrains the path. The distinction matters because pumping is a common-pool problem: Provencher and Burt [17] analyze the externalities of common-property extraction, and Pfeiffer and Lin [18] show that spatial externalities in groundwater use are empirically important. A single-tank model, as used here, abstracts from those local drawdown effects, which is why the annual floor in this paper protects modeled year-end stocks and not conditions at individual wells.

Irrigation response, measurement and yield transfer

Empirical work on irrigation demand and yield response faces data problems that apply directly to county series. Hendricks and Peterson [19] estimate irrigation water demand on intensive and extensive margins, and Mieno and Brozović [20] show that incomplete information about prices and quantities biases elasticity estimates. In this paper, county irrigation depth is an agency estimate, with most annual values falling on whole acre-inches, so a regression on it is biased toward zero and cannot identify the response curve. Experimental plot data, such as the cotton cultivar trials of Snowden, Ritchie and Thompson [1], identify the shape of the response but come from research plots, and the literature on yield gaps [21] shows that farm yields commonly fall below plot yields. That motivates the plot-to-county yield multiplier used in the calibration here.

Positioning of this paper

The paper sits between these literatures. It borrows the stock-and-control structure of the optimal control tradition, applies it to a county-scale cotton system in the Texas High Plains, and focuses on the form of the stock constraint and on how sensitive conclusions are to calibration. It does not estimate a structural model of farmer behavior, represent spatial drawdown, or include stochastic weather, and those are the main extensions the literature suggests.

Methods

Study design and calibration

The planning horizon contains 20 annual periods. Available land is fixed at 106,160 acres, using the 2024 irrigated cotton acreage supplied from the Texas Water Development Board (TWDB) as a scenario proxy. The model represents that eligible land pool rather than all county cropland. It uses selected historical hydrologic benchmarks and 2024 budgets; it is not an econometric model fitted to the complete 2000–2024 county data series.

Irrigated production is calibrated to the three irrigation-treatment means for each of the 2010 and 2011 New Deal experiments reported by Snowden et al. (2013). Irrigation depth is converted from millimeters to acre-feet per acre, and lint yield from kilograms per hectare to pounds per acre. The curves summarize experimental production conditions and are transferred to the county scenario without claiming county-level causal identification [1].

Table 1. Principal baseline parameters

Parameter	Value	Interpretation
Available land	106,160 acres	Fixed eligible land proxy
Horizon and discount rate	20 years; 3%	First year undiscounted
Initial usable stock	1,750,000 AF	Hypothetical historic benchmark [2]
Net inflow before other pumping	50,265 AF/year	Historic modeled-flow benchmark [3]
Other groundwater withdrawal	16,768 AF/year	Exogenous scenario assumption
Groundwater supply share	80%	Remaining supply unverified
Cotton withdrawal ceiling	159,240 AF/year	$1.50 \text{ AF} \times \text{eligible acres proxy}$ [4]
Applied-water cost	120 dollars/AF	Budget-derived cost benchmark [5]
Committed fixed cost	255.25 dollars/acre/year	Charged on all eligible land [5]
Dryland lint yield	400 lb/acre	Budget benchmark in both conditions [6]
Groundwater stock target	90% or 95% of initial stock	Year 20 target or annual year-end floor

Calibration checks and county-data diagnostics

Three checks assess how far the experimental curves and benchmark inputs can be tied to county observations. First, annual irrigated cotton yield is regressed on TWDB irrigation depth, annual precipitation and a linear trend for the 18 years with data (2011 and 2019–2024 are missing). The linear model is reported with HC3 and Newey-West (two-lag) standard errors and repeated without the trend, in first differences and in logarithms; a pairs bootstrap (5,000 draws) and leave-one-out estimates describe sensitivity to individual years. Irrigation depth is cotton water from all sources divided by TWDB cotton acreage, so it is an agency estimate rather than a measured field application [8].

Second, the ratio of county irrigated yield to the experimental yield at similar depth (2010) and to the drip-budget yield (2012–2018) is compared with the ratio of pivot to drip budget yields to select a plot-to-county yield multiplier [5, 9]. Third, observed Ogallala pumping for 2000–2024 is compared with the modeled inflow benchmark, and TWDB cotton water is compared with Ogallala irrigation for all crops, to check the internal consistency of the storage, inflow and groundwater-share assumptions. These checks do not estimate the production function, which remains identified only by the experimental treatment means.

Production and economic objective

Annual controls are irrigated acreage a_t , dryland acreage d_t , and applied-water increments z_{1t} and z_{2t} above the lowest observed irrigation depth. Fallow acreage is the remaining land. For treatment points $j = 0, 1, 2$, q_{jt} denotes applied depth and y_{jt} denotes lint yield per acre. Total irrigated lint Y_{it} and applied irrigation W_t are determined by concave piecewise linear interpolation:

$$W_t = q_{0t}a_t + z_{1t} + z_{2t}, \quad P_t = gW_t$$

$$Y_{it} = y_{0t}a_t + m_{1t}z_{1t} + m_{2t}z_{2t}$$

$$m_{jt} = \frac{y_{jt} - y_{j-1,t}}{q_{jt} - q_{j-1,t}}, \quad j = 1, 2$$

The interpolated curve has declining segment slopes in each experimental year. Segment water volumes are bounded by irrigated acreage. Irrigated depth therefore stays within the observed treatment range; zero irrigated acreage is permitted, and no production response is extrapolated below the lowest irrigated treatment.

$$a_t, d_t \geq 0, \quad a_t + d_t \leq A$$

$$0 \leq z_{1t} \leq (q_{1t} - q_{0t})a_t$$

$$0 \leq z_{2t} \leq (q_{2t} - q_{1t})a_t$$

The objective maximizes county-level discounted net returns over the horizon. Annual returns π_t include irrigated cotton margin, dryland cotton margin, and fixed costs on all available acres A :

$$\max_{a_t, d_t, z_{1t}, z_{2t}} \sum_{t=1}^{20} \frac{\pi_t}{(1+r)^{t-1}}$$

$$\pi_t = u_I Y_{it} - c_I a_t - c_W W_t + u_D y_{Dt} d_t - c_D d_t - fA$$

Here u_I and u_D are net receipts per pound after cottonseed revenue, stripping and ginning; c_I and c_D are other variable costs per planted acre; c_W is the cost per applied acre-foot; f is committed fixed cost; and y_{Dt} is dryland lint yield per acre. The baseline lint price is \$0.72/lb, giving $u_I = \$0.711873/\text{lb}$ and $u_D = \$0.703075/\text{lb}$. Budget totals imply $c_I = \$530.54/\text{acre}$ and $c_D = \$198.62/\text{acre}$ [5, 6].

Dryland cotton receives no applied irrigation. At 400 lb/acre its margin before fixed costs is \$82.61/acre. Seed revenue and harvest costs scale proportionally with lint, while other variable costs remain constant per planted acre. The dryland budget's own \$100.25/acre fixed cost is recorded as a benchmark but is not charged again. The model retains \$255.25/acre on all eligible land, including fallow, to keep the comparison with the earlier irrigated-and-fallow model consistent. These committed costs do not affect the optimal controls because they are constant across land choices.

Groundwater dynamics and scenario solution

Groundwater stock S_t is measured at the end of period t . Net inflow R includes modeled recharge and net lateral and spring flows before other withdrawal B . Cotton groundwater withdrawal P_t equals the assumed groundwater share $g = 0.80$ of applied irrigation. The baseline constraints are:

$$S_t = S_{t-1} + R - B - P_t, \quad S_0 = 1,750,000$$

$$S_t \geq 0, \quad 0 \leq P_t \leq 159,240$$

$$S_{20} \geq \theta S_0, \quad \theta \in \{0.90, 0.95\}$$

The annual ceiling is a screening proxy based on eligible acreage and the posted HPWD rule benchmark, not an aggregation of verified property-level water rights. Actual well capacity, contiguous tract eligibility, groundwater source allocation and reserve-bank withdrawals are not established. The assumed stock is derived from a historical TWDB recovery benchmark and is not a measurement of current usable county storage [2–4].

In the terminal-only baseline, storage may fall below the year 20 target during the horizon, provided it remains nonnegative and returns to the required level at the endpoint. Summing the state equation gives a cotton pumping budget of $(1 - \theta)S_0 + 20(R - B)$: 844,940 AF under the 90% target and 757,440 AF under the 95% target.

Annual groundwater floor comparison

Six additional cases impose the same 90% or 95% threshold at every year-end, replacing the nonnegative intermediate-stock constraint while retaining the terminal requirement. Each is paired with its matching baseline; all production, price, cost and hydrologic assumptions remain unchanged. This restricts cumulative pumping through each preceding period but leaves the final cumulative budget unchanged. It protects modeled year-end stocks, not storage continuously within each year. The 36 sensitivity experiments remain terminal-only and are not rerun with annual floors.

Baseline and sensitivity scenarios

Six terminal-only baselines combine two stock targets with three deterministic production sequences: wet repeats 2010, dry repeats 2011, and alternating uses 2010 in odd periods and 2011 in even periods. These are experimental production environments, not a weather forecast. Dryland yield remains 400 lb/acre in both environments because its weather-specific response has not been calibrated.

Thirty-six terminal-only sensitivity cases apply 12 tests under the 90% target: price $\pm 20\%$; water cost, initial stock and net inflow $\pm 25\%$; dryland yields of 300 or 500 lb/acre; and dry-period dryland yields of 100 or zero. Each test is applied to all three sequences. Stock tests change the absolute terminal threshold; inflow tests change the aggregate flow benchmark rather than an independently estimated recharge process.

Calibration tiers, shadow values and uncertainty analysis

Calibration tiers repeat the 90% terminal-only baselines under alternative parameter sets. Tier A uses the experimental curves and drip-budget costs. Tier B multiplies irrigated yields by 0.667 and replaces other variable cost with \$426.35 per acre, derived from the 2024 pivot budget (total variable cost less stripping, ginning and energy) [9]. Tier C adds dryland yields from a regression of county non-irrigated yield on annual precipitation. A separate yield-multiplier test varies the multiplier from 1.0 to 0.6 with drip-budget costs unchanged.

The shadow value of the stock constraint is the change in NPV per additional acre-foot of required year 20 stock, computed by resolving each case with the target changed by 2,000 AF. The conservation-cost frontier resolves each terminal-only baseline for targets from 50% to 105% of initial stock. Because committed fixed costs are identical across land choices, these comparisons use NPV before committed fixed costs, meaning

discounted returns above variable costs; the gain over dryland-only is that value minus the dryland-only NPV on all eligible land (\$134.39 million at 400 lb/acre). A uniform-quota comparison sets annual pumping at the 90% budget divided by 20 and reoptimizes acreage each year.

The uncertainty analysis resolves the model for 2,000 random parameter sets. Eleven inputs are drawn independently from uniform ranges (listed in the notes to Table 9). In each year the wet or dry condition persists with probability 0.6 and is otherwise redrawn using the wet-year probability. Influence is summarized by standardized rank regression coefficients. The ranges are analyst-specified bounds, so the results describe scenario sensitivity and are not confidence intervals.

Numerical implementation

The model is solved by SciPy HiGHS linear programming under deterministic perfect foresight. All 48 distinct cases—42 terminal-only cases and six annual-floor cases—reach optimal status. Verification checks constraints, production interpolation, primal–dual objectives and accounting. The six terminal baselines are reproduced within \$0.01; all annual-floor stocks meet their thresholds and their NPVs do not exceed paired terminal-only NPVs. Sensitivity results are assumption tests, not confidence intervals.

An independent Python implementation of the same linear program, built from the workbook inputs, reproduces all 42 terminal-only NPVs within \$0.26 and the six annual-floor NPVs within \$0.17. The additional analyses described above use that implementation and the same deterministic perfect-foresight assumption; all additional cases reach optimal status.

Results

Baseline land allocation and discounted returns

With a 90% terminal target, average irrigated acreage is 78,136 acres in the wet sequence, 25,229 in the dry sequence and 61,170 in the alternating sequence (Table 2). Dryland occupies the remaining land at the assumed 400 lb/acre yield, leaving no fallow acreage. These are arithmetic averages of annual decisions, not constant annual acreage prescriptions.

Table 2. Baseline optimized land allocation and county returns

Condition and year 20 target	Irrigated acres	Dryland acres	Pumping thousand AF	NPV million dollars	Gain vs fallow million dollars
Wet, 90% retained	78,136	28,024	844.94	139.61	28.17
Dry, 90% retained	25,229	80,931	844.94	-247.12	95.26
Alternating, 90% retained	61,170	44,990	844.94	-0.00	53.53
Wet, 95% retained	70,045	36,115	757.44	104.07	37.15
Dry, 95% retained	22,616	83,544	757.44	-250.37	99.03
Alternating, 95% retained	58,558	47,602	757.44	-3.41	57.48

Notes: Acreage values are annual means. Pumping is the 20-year total for cotton only. Gain compares the dryland-enabled solution with the same model restricted to irrigated cotton and fallow. Negative returns include committed fixed costs on all eligible acres.

Wet-sequence net present value is \$139.61 million with the 90% target and \$104.07 million with the 95% target. Dry-sequence values remain negative at -\$247.12 million and -\$250.37 million. The alternating sequence is approximately at break-even with the 90% target (-\$2,854 in total) and has a -\$3.41 million value with the 95% target. Near-zero results should not be interpreted as precisely estimated break-even outcomes.

Stock targets and timing of withdrawals

All six baseline terminal constraints bind. Raising the target from 90% to 95% reduces permitted cumulative cotton pumping by 87,500 AF and lowers average irrigated acreage. Discounting favors earlier withdrawals when production conditions repeat, followed by reduced irrigation and stock recovery. Figure 1 shows why the terminal target does not guarantee a sustained stock floor.




Figure 1. Groundwater stock paths for the three baseline 90% cases. The dashed line identifies the required year 20 stock, not an annual minimum. Stocks include the assumed inflow and other pumping.

Sensitivity of returns and land choices

Table 3 reports reoptimized net present values under the 90% target. Lint-price changes produce large differences in returns. A 20% price reduction eliminates irrigation in the dry sequence, with all eligible land assigned to dryland and cotton groundwater pumping falling to zero. The wet sequence becomes unprofitable even though its average irrigated acreage is unchanged. A 20% price increase improves dry-sequence returns but does not eliminate the loss after committed fixed costs.

Table 3. Reoptimized county NPV under one-at-a-time sensitivity tests

Test	Wet	Dry	Alternating
Cotton price -20%	-145.65	-374.55	-227.60
Cotton price +20%	424.87	-91.06	236.86
Water cost -25%	165.68	-217.38	26.36
Water cost +25%	113.54	-276.87	-26.37
Initial stock -25%	122.09	-248.73	-1.71
Initial stock +25%	156.78	-245.52	1.70
Net inflow -25%	33.01	-256.61	-10.11
Net inflow +25%	233.30	-238.05	9.23
Dryland yield 300 lb/acre	115.63	-328.19	-45.56
Dryland yield 500 lb/acre	163.59	-166.05	45.56
Dry-period dryland yield 100	139.61	-342.38	-53.53
Dry-period dryland yield 0	139.61	-342.38	-53.53

Notes: Values are millions of dollars over 20 years. Water cost changes from \$120 to \$90 or \$150 per applied AF. Initial stock changes to 1,312,500 or 2,187,500 AF. Net inflow changes to 37,698.75 or 62,831.25 AF/year. In the last two tests, the wet-period dryland yield remains 400 lb/acre.

The $\pm 25\%$ water-cost tests change returns without changing mean irrigated acreage in these solved cases. In the alternating sequence, lower water cost raises NPV to \$26.36 million and higher water cost reduces it to $-\$26.37$ million. Thus, a similar aggregate acreage allocation can coexist with substantially different profitability.

Net inflow strongly affects the water budget. Reducing the inflow benchmark by 25% lowers wet-sequence average irrigated acreage from 78,136 to 54,895 acres; increasing it by 25% raises acreage to 101,378. Changing initial stock by $\pm 25\%$ produces smaller acreage changes because only 10% of the initial-stock change enters the cumulative withdrawal budget under the proportional terminal target.

Dryland yield assumptions determine the value of replacing fallow. At 300 lb/acre, dryland remains planted in the baseline acreage pattern but the dry-sequence NPV falls to $-\$328.19$ million. At 100 lb/acre in dry periods, the dryland margin is negative, and the optimizer instead leaves an average of 80,931 acres fallow in the dry sequence. NPV then returns to $-\$342.38$ million, the corresponding irrigated-and-fallow result. An anticipated zero dryland yield gives the same land choice. This is avoidance of planting under known unfavorable conditions, not an estimate of losses from unexpected crop failure.

Annual stock protection and the timing of pumping

The annual-floor comparison separates protection of intermediate stocks from protection only at the terminal date (Table 4). All six annual-floor cases meet the specified year-end threshold. At the 90% target, minimum stock rises from 80.4% to 90.0% in the wet sequence, 63.2% to 90.0% in the dry sequence and 84.2% to 90.0% in the alternating sequence. The paired policies end with the same stock and cumulative cotton pumping.

Table 4. Discounted returns and minimum stock under paired policies

Condition and target	Terminal NPV million \$	Annual NPV million \$	NPV cost million \$	Terminal minimum %	Annual minimum %
Wet, 90%	139.61	126.16	13.45	80.4	90.0
Dry, 90%	-247.12	-251.64	4.52	63.2	90.0
Alternating, 90%	-0.00	-1.03	1.03	84.2	90.0
Wet, 95%	104.07	80.47	23.60	82.2	95.0
Dry, 95%	-250.37	-255.30	4.93	66.3	95.0
Alternating, 95%	-3.41	-4.64	1.23	89.2	95.0

Notes: NPV cost is terminal-only NPV minus annual-floor NPV; it measures foregone discounted modeled returns, not a fee or a monetized environmental benefit. Minimum stock is the lowest modeled year-end stock as a percentage of the initial stock. The near-zero terminal NPV for alternating 90% is -\$2,854 before rounding.

Under dry conditions with a 90% target, imposing an annual floor reduces first-five-year pumping from 796,200 to 342,485 AF, deferring 453,715 AF beyond that interval. The 95% dry comparison defers 502,455 AF. The annual-floor paths reach the threshold early and subsequently balance cotton pumping with net inflow after other withdrawals (Figure 2).




Figure 2. Dry-condition groundwater paths under terminal-only targets and annual year-end floors. The common initial stock and flow assumptions isolate the effect of the policy timing constraint.

Mean irrigated acreage, mean dryland acreage and total 20-year cotton pumping are unchanged in all six policy pairs. The annual floor nevertheless lowers NPV by \$1.03–\$13.45 million under the 90% target and \$1.23–\$23.60 million under the 95% target. These differences arise from reallocated production timing and discounting within the specified model, rather than lower cumulative water use. The unchanged acreage averages are a case-specific result, not a general property of annual floors.

Calibration evidence, shadow values and uncertainty

County yield data and plot-to-county transfer

County data do not identify a diminishing-returns response. In the linear specification the coefficient on irrigation depth is 81 lb/acre per AF/acre with an HC3 standard error of 238 ($p = 0.74$); the Newey-West standard error is 185 ($p = 0.67$), and the no-trend, first-difference and logarithmic specifications give p -values of 0.77 to 0.91 (Table 5). A pairs bootstrap gives a 95% interval of -348 to 517 lb/AF and leave-one-out estimates range from -27 to 247 , while the experimental treatment steps imply marginal products of 735 to $1,857$ lb/AF. The interval spans both zero and the experimental values, and the minimum detectable effect is roughly 666 lb/AF [8].

Table 5. County irrigated-yield regressions on irrigation depth, 2000–2024 ($n = 18$)

Specification	Irrigation coefficient	Standard error	p-value
Linear, HC3 errors	81.13	237.96	0.738
Linear, Newey-West (2 lags)	81.13	185.32	0.668
Linear, no trend, HC3	-33.33	161.54	0.839
First differences, HC3	48.06	161.68	0.771
Log-log elasticity, HC3	0.035	0.311	0.911

Notes: The dependent variable is county irrigated cotton lint yield (lb/acre); controls are annual precipitation and, except where stated, a linear trend. Coefficients are lb/acre per AF/acre, except the log-log row, which is an elasticity. Years 2011 and 2019–2024 have no irrigated yield. The Durbin-Watson statistic for the linear model is 1.25. The pairs-bootstrap 95% interval (5,000 draws) is -347.9 to 516.9 .

Irrigation depth in the county series is itself coarse. Twenty-one of 25 annual values lie within 0.02 acre-inch of a whole number of acre-inches and the standard deviation is 0.23 AF/acre, which is consistent with an agency estimate rather than a measured application and biases a regression slope toward zero. Applied water is also chosen in response to weather, so the slope would remain confounded even if depth were measured exactly. The county series is therefore used to calibrate the level of yields rather than their response to water.

Three comparisons place county irrigated yields at about two-thirds of experimental plot yields (Figure 3). County irrigated yield in 2010 was 970 lb/acre, 0.68 of the 1,435 lb/acre fully irrigated plot yield at nearly the same depth (0.66 versus 0.68 AF/acre). County irrigated yields averaged 958 lb/acre in 2012–2018, 0.64 of the 1,500 lb/acre drip-budget yield, and the pivot budget expects 1,000 lb/acre, 0.67 of the drip budget [5, 9].




Figure 3. County irrigated cotton yield and irrigation depth for years with data (points, shaded by year), with the 2010 and 2011 experimental treatment means (solid) and the same curves multiplied by 0.667 (dashed). County depth is TWDB cotton water divided by cotton acreage.

Yield transfer and calibration tiers

Table 6 rescales the experimental curves while holding drip-budget costs and the 400 lb/acre dryland yield fixed. Results are reported as NPV before committed fixed costs, which excludes the \$255.25 per acre charge on all eligible land (present value \$415.2 million) because that charge is identical across land choices. In the wet sequence NPV falls by about \$131 million for each 0.1 reduction in the multiplier. Irrigation disappears from the dry sequence at multipliers below about 0.92 and from the wet and alternating sequences below about 0.68; each such case returns the dryland-only value of \$134.39 million.

Table 6. NPV before committed fixed costs under rescaled irrigated yields, 90% terminal target

Irrigated-yield multiplier	Wet	Dry	Alternating
1.000 (baseline)	554.85	168.11	415.23
0.900	423.54	134.39	320.01
0.800	292.23	134.39	235.72
0.700	160.92	134.39	151.42
0.667	134.39	134.39	134.39
0.600	134.39	134.39	134.39

Notes: Values are millions of dollars over 20 years. A value of 134.39 is the dryland-only NPV on all eligible land; it indicates that no land is irrigated. Other variable cost remains at the drip-budget value of \$530.54 per acre.

Table 7 adds the cost and dryland-yield calibrations. Tier B multiplies irrigated yields by 0.667 and replaces other variable cost with \$426.35 per acre derived from the pivot budget; Tier C adds dryland yields from a regression of county non-irrigated yield on annual precipitation (573 lb/acre in the wet sequence and 194 lb/acre in the dry sequence; $R^2 = 0.27$, 18 years). Under Tier B the stock constraint binds only in the wet sequence, with a shadow value of \$111 per acre-foot; the dry sequence irrigates no land, and the alternating sequence pumps 573,988 AF, below the 844,940 AF budget, so profitability rather than stock limits irrigation. Under Tier C no land is irrigated in any sequence, which does not match observed irrigated acreage. The baseline (Tier A) is

therefore best read as an upper bound on the value of irrigation, and Tiers B and C as bracketing alternatives. The Tier C dryland yields are per harvested acre and may overstate returns per planted acre.

Table 7. Calibration tiers under the 90% terminal-only target

Tier and condition	NPV before fixed costs	Irrigated acres	Dryland acres	Shadow value \$/AF
A, wet	554.85	78,136	28,024	399.9
A, dry	168.11	25,229	80,931	36.7
A, alternating	415.23	61,170	44,990	38.9
B, wet	251.10	78,136	28,024	111.0
B, dry	134.39	0	106,160	0.0
B, alternating	209.30	53,080	53,080	0.0
C, wet	332.00	0	106,160	0.0
C, dry	0.00	0	0	0.0
C, alternating	168.50	0	53,080	0.0

Notes: NPV before fixed costs is in millions of dollars. Tier A uses the experimental curves and drip-budget costs; Tier B uses a yield multiplier of 0.667 and pivot-budget variable cost; Tier C adds regression-based dryland yields. Acreage values are annual means. Shadow value is the change in NPV per additional acre-foot of required year 20 stock.

Figure 4 maps the alternating sequence over multipliers from 0.5 to 1.0 and dryland yields from 200 to 600 lb/acre, with other variable cost interpolated between the pivot and drip budget values. The stock constraint binds in 24 of 99 cells, all with multipliers of 0.80 or higher and dryland yields of 450 lb/acre or less. Irrigation is zero in 15 cells, all with multipliers of 0.70 or lower and dryland yields of 400 lb/acre or more. The county-calibrated multiplier of 0.667 lies in the region where profitability, not the stock, limits irrigation.




Figure 4. Mean irrigated acreage in the alternating sequence under the 90% target across irrigated-yield multipliers (horizontal axis) and dryland yields (vertical axis). The black line separates cases in which the stock constraint binds (lower right) from cases in which it does not; the dashed line marks 0.667.

Shadow value of groundwater and the cost of conservation

At the 90% target an additional acre-foot of required terminal stock costs \$400 of discounted returns in the wet sequence and \$37 to \$39 in the dry and alternating sequences (Table 8). These are scarcity values on top of the \$120 per acre-foot water cost. Raising the target from 90% to 95% costs \$406, \$37 and \$39 per acre-foot, close to the shadow values. In the wet sequence the constraint does not bind for targets up to about 73% of initial stock: all eligible land is irrigated each year, cotton pumping totals 1.148 million AF, and the stock still ends at 1.27 million AF because inflow offsets pumping. In the dry sequence, pumping limited only by nonnegative stock totals 2.42 million AF and exhausts the stock in year 14 (Figure 5).

Table 8. Shadow value of the stock constraint and cost of conservation, terminal-only baselines

Condition	Shadow value 90% target \$/AF	Cost 90% to 95% \$/AF	Unconstrained pumping thousand AF	NPV before fixed costs, 90%	Gain over dryland-only
Wet	399.9	406.2	1,148.0	554.85	420.46
Dry	36.7	37.2	2,419.9	168.11	33.72
Alternating	38.9	38.9	2,166.4	415.23	280.84

Notes: Shadow value is computed by resolving each case with the year 20 target raised by 2,000 AF.

Unconstrained pumping is the 20-year cotton pumping with only nonnegative stock and the annual ceiling. NPV and gains are millions of dollars before committed fixed costs; the dryland-only NPV is \$134.39 million.




Figure 5. NPV before committed fixed costs as a function of the required year 20 stock, terminal-only baselines. The dotted line is the dryland-only value for all eligible land.

Of the 844,940 AF cotton pumping budget at the 90% target, 669,940 AF (79%) is twenty years of net inflow after other withdrawals; only 175,000 AF comes from drawing down the initial stock. This explains why the results respond more to the inflow benchmark than to the stock benchmark. The observed record points the same way. Observed Ogallala pumping averages 166,306 AF per year over the 23 years with data and was 101,696 AF in 2024, against a modeled inflow of 50,265 AF per year, and cotton groundwater at the assumed share (84,928 AF in 2024) is twice the constant 42,247 AF per year implied by the 90% target. Taken at face value, observed pumping and the inflow benchmark imply cumulative net withdrawals of about 2.67 million AF over those years, more than the 1.75 million AF initial benchmark, and TWDB cotton water from all sources in 2024 (106,160 AF) exceeds Ogallala irrigation for all crops (95,854 AF). These series come from different data products and are not additive, so they bound rather than test the stock and groundwater-share assumptions [8].

A uniform annual quota equal to the 90% budget divided by 20 years, with acreage reoptimized each year, lowers NPV before fixed costs by 5.2% in the wet sequence, 3.7% in the dry sequence and 16.5% in the alternating sequence relative to the perfect-foresight solution. In this deterministic setting that difference indicates the value of flexible multi-year water accounting.

Uncertainty analysis

Across 2,000 random parameter sets, irrigation exceeds the dryland-only value by more than \$1 million in 93% of draws and mean irrigated acreage is zero in 7% (Table 9). The stock constraint binds in 44% of draws (shadow value above \$1 per acre-foot), with a median shadow value of \$77 per acre-foot among those. Standardized rank regression ranks the irrigated-yield multiplier (0.55), dryland yield (−0.44) and lint price (0.44) as the largest influences on the gain from irrigation, followed by the wet-year probability (0.20); initial stock (0.01), other withdrawals (−0.02) and the discount rate (−0.07) have little influence (Figure 6). For mean irrigated acreage the wet-year probability (0.33), net inflow (0.24) and the multiplier (0.23) rank highest.

Table 9. Distribution of outcomes across 2,000 random parameter sets

Output	5th percentile	25th	Median	75th	95th percentile
NPV before fixed costs, \$ million	50.0	142.2	233.2	344.3	532.4
Mean irrigated acres	0	31,848	47,839	63,696	83,692
Cotton pumping, thousand AF	0.0	360.4	577.4	767.7	1,010.7
Shadow value, \$/AF	0.0	0.0	0.0	65.4	288.0
Gain over dryland-only, \$ million	0.0	46.1	104.8	188.9	361.5

Notes: Eleven inputs are drawn independently from uniform ranges: lint price \$0.55 to \$0.90 per pound; water cost \$90 to \$160 per AF; net inflow 33,900 to 62,831 AF per year; initial stock 1,312,500 to 2,187,500 AF; terminal target 80% to 100%; groundwater share 0.6 to 1.0; dryland yield 250 to 550 lb/acre; irrigated-yield multiplier 0.6 to 1.0; discount rate 2% to 7%; wet-year probability 0.3 to 0.7; and other withdrawals 10,000 to 25,000 AF per year. The ranges are analyst-specified bounds, not estimated distributions.




Figure 6. Standardized rank regression coefficients for the gain from irrigation over dryland-only (left) and its distribution across 2,000 draws (right). The dashed line marks the median.

Discussion

The paired policy comparison shows that the timing of groundwater protection matters even when the endpoint requirement and cumulative pumping budget are the same. A terminal-only target permits intermediate depletion followed by recovery, whereas an annual floor restricts depletion at each modeled year-end. In the dry 90% case, the minimum stock differs by 26.8 percentage points and the annual floor reduces discounted returns by \$4.52 million. This illustrates a tradeoff between early production income and protection of intermediate stocks. It does not establish continuous protection between yearly observations or the physical feasibility of a county-scale pumping schedule.

Identical mean acreage and cumulative pumping conceal materially different annual decisions. An acreage average alone cannot show whether irrigation is concentrated early or distributed across later periods. With deterministic production curves, a fixed water share and constant inflow, the solved baseline cases retain the same cumulative water allocation under both policies. The annual floor redistributes that allocation over time. More variable weather, changing lift costs, binding hydraulic limits or alternative crops could change both total water use and average acreage; the present results do not establish invariance under those conditions.

The NPV differences in Table 4 measure the modeled opportunity cost of the stricter timing constraint. They are not estimates of the total social value of groundwater conservation. The objective includes cotton returns but excludes environmental services, household water security, future well reliability and other nonmarket benefits. It also excludes public implementation and enforcement costs. A policy evaluation would need those benefits and costs before comparing the annual floor with a terminal-only target on welfare grounds.

Dryland cotton improves returns relative to fallow when its margin before committed costs is positive, yet the dry sequences remain unprofitable after fixed costs under both stock policies. The advantage depends on the assumed 400 lb/acre dryland yield; lower dry-period yields can make fallow preferable. Committed fixed costs affect reported profitability but not optimal controls because they are constant across choices. Allowing equipment, rental or other obligations to be released after land-use changes would require a different cost model and could change the interpretation of financial losses.

The calibration analysis qualifies the level of the baseline returns more than the structure of the constraint. The experimental curves come from subsurface-drip research plots, and county irrigated yields are about two-thirds of those plot yields. Rescaling the curves lowers NPV before committed fixed costs and removes irrigation from the dry sequence at multipliers below roughly 0.92 and from the wet and alternating sequences below roughly 0.68, so the baseline irrigated returns are better read as an upper bound than as a county estimate. What survives rescaling is the logic of the constraint: it carries a positive shadow value only where irrigated margins are high enough for water to be scarce, as in the wet sequence under the county-calibrated Tier B, and elsewhere profitability rather than the stock limits irrigation.

County data cannot substitute for the experiments in identifying the response curve. The county depth series is a rounded agency estimate, applied water is chosen in response to weather, and a regression with 18 observations cannot distinguish the experimental slopes from zero. The experiments in turn offer two years, three treatments per year and cultivars that share main plots; the concave interpolation imposes diminishing returns that three points per year cannot test, although the steps between treatments exceed the reported least-significant-difference thresholds. A two-source design is therefore appropriate: experiments for the shape of the response, and county and budget data for the level of yields and as an external check. The Tier C result, in which no land is irrigated, shows that the deterministic margin model does not reproduce observed irrigated acreage once county dryland yields are used. Production risk, crop insurance, rotation and abandonment are omitted and could explain part of that gap.

The shadow values measure the cost of an exogenous stock requirement, not the gain from coordinating pumping, so they are not comparable with the literature that contrasts optimal and competitive extraction [7]. A

competitive benchmark with a stock-dependent pumping cost is not computed here. Lift-cost feedback is probably small relative to the cost uncertainty already tested: with an assumed specific yield of 0.15 and roughly 575,000 acres of county area, a 10% drawdown of the 1.75 million AF benchmark corresponds to about 2 feet of water-level decline and drawing down the full benchmark to about 20 feet, which would raise energy cost by about 8% if cost scales with a total dynamic head of 250 feet. Those aquifer parameters are assumptions, not values from the supplied data, and should be verified before the cost feedback is set aside.

The findings remain provisional scenario evidence. Initial usable storage and inflow are historical benchmarks; groundwater share, alternative supply, physical well capacity and actual water-rights acreage are unverified. Transferring two experimental curves to county acres does not quantify farm heterogeneity or yield uncertainty. Perfect foresight omits surprise drought and crop failure, and the stock model omits spatial drawdown and lift-cost feedback. The sensitivity tests apply only to terminal targets, so robustness of the annual-floor comparison across those tests has not been established.

Further calibration should verify usable storage, source availability, pumping capacity and property-level constraints, estimate dryland responses by production environment, and distinguish avoidable from committed costs. Those steps would support an empirical policy interpretation. Until then, the comparison demonstrates how annual year-end protection changes the feasible timing and discounted value of irrigation under explicit assumptions; it should not be used as a county pumping recommendation.

Conclusion

This study asked how the form of a groundwater stock rule changes cotton irrigation in Lubbock County. A linear program over 20 years compared a year 20 stock target with an annual year-end floor under wet, dry and alternating production sequences. Four results stand out. First, the two rules permit the same cumulative pumping and mean acreage, but the annual floor moves pumping later and raises the lowest modeled stock from 63–84% to 90% of the initial stock, at a cost of \$1.0 to \$13.5 million in net present value, or 0.2% to 2.7% of discounted returns before committed fixed costs. Second, the shadow value of the stock constraint is about \$400 per acre-foot in the wet sequence and \$37 to \$39 in the dry and alternating sequences, so water is scarce only where irrigated margins are high. Third, county irrigated yields are about two-thirds of the experimental plot yields, and county data cannot identify the yield response to water; rescaling the curves removes irrigation from the dry sequence and, below a multiplier of about 0.68, from the others. Fourth, across 2,000 random parameter sets the gain from irrigation depends mainly on the yield multiplier, dryland yield and cotton price, and hardly at all on initial storage.

These results are conditional scenario evidence. The model uses a single-tank aquifer, perfect foresight, historical storage and inflow benchmarks, and an assumed groundwater share, and it does not reproduce observed irrigated acreage once county dryland yields are used. They should not be read as a county forecast or a pumping recommendation. The most valuable extensions are well water-level data for a historical backcast, a stochastic dynamic program with weather and price shocks, and data on crop and irrigation-system mix. Until then, the paper's main contribution is to show that the timing of protection, and not only its endpoint, determines how much irrigated production a stock rule forgoes.

References

- [1] Snowden, C., Ritchie, G., and Thompson, T. (2013). Water Use Efficiency and Irrigation Response of Cotton Cultivars on Subsurface Drip in West Texas. *Journal of Cotton Science*, 17, 1–9. <https://www.cotton.org/journal/2013-17/1/upload/jcs17-1.pdf>
- [2] Texas Water Development Board. (2013). Task13-042 groundwater modeling report, Table 3. <https://www.twdb.texas.gov/groundwater/docs/GAMruns/Task13-042.pdf>
- [3] Texas Water Development Board. GR03-22 groundwater modeling report. Historical modeled flow benchmark. <https://www.twdb.texas.gov/groundwater/docs/GAMruns/GR03-22.pdf>
- [4] High Plains Underground Water Conservation District. Posted rules, amended rules effective December 1, 2023, Rule 5.2. <https://www.hpwd.org/rules>
- [5] Texas A&M AgriLife Extension. (2024). Estimated Costs and Returns per Acre, Irrigated Cotton, Drip, South Plains Extension District 2. User-supplied file 2024D2CottonDrip.pdf.
- [6] Texas A&M AgriLife Extension. (2024). Estimated Costs and Returns per Acre, Dryland Cotton, Herbicide Tolerant, South Plains Extension District 2. User-supplied file 2024D2CottonDryland.pdf.
- [7] Gisser, M., and Sánchez, D. A. (1980). Competition versus optimal control in groundwater pumping. *Water Resources Research*, 16(4), 638–642.
- [8] County series in Lubbock_Water_Cotton_Master_Dataset_2000_2024.xlsx, compiled from TWDB Report 152 (cotton acreage and water use) and Report 20 (Ogallala pumpage), USDA NASS (county cotton yields and Texas upland cotton prices) and NOAA (annual precipitation and temperature).
- [9] Texas A&M AgriLife Extension. (2024). Estimated Costs and Returns per Acre, South Plains Extension District 2, pivot-irrigated cotton. User-supplied file 2024D2CottonPivot.pdf.
- [10] Burt, O. R. (1964). Optimal resource use over time with an application to ground water. *Management Science*, 11(1), 80–93.
- [11] Feinerman, E., and Knapp, K. C. (1983). Benefits from groundwater management: magnitude, sensitivity, and distribution. *American Journal of Agricultural Economics*, 65(4), 703–710.
- [12] Koundouri, P. (2004). Current issues in the economics of groundwater resource management. *Journal of Economic Surveys*, 18(5), 703–740.
- [13] Tsur, Y., and Graham-Tomasi, T. (1991). The buffer value of groundwater with stochastic surface water supplies. *Journal of Environmental Economics and Management*, 21(3), 201–224.
- [14] Scanlon, B. R., et al. (2012). Groundwater depletion and sustainability of irrigation in the US High Plains and Central Valley. *Proceedings of the National Academy of Sciences*, 109(24), 9320–9325.
- [15] Hornbeck, R., and Keskin, P. (2014). The historically evolving impact of the Ogallala Aquifer: agricultural adaptation to groundwater and drought. *American Economic Journal: Applied Economics*, 6(1), 190–219.
- [16] Terrell, B. L., Johnson, P. N., and Segarra, E. (2002). Ogallala aquifer depletion: economic impact on the Texas High Plains. *Water Policy*, 4(1), 33–46.
- [17] Provencher, B., and Burt, O. (1993). The externalities associated with the common property exploitation of groundwater. *Journal of Environmental Economics and Management*, 24(2), 139–158.
- [18] Pfeiffer, L., and Lin, C.-Y. C. (2012). Groundwater pumping and spatial externalities in agriculture. *Journal of Environmental Economics and Management*, 64(1), 16–30.
- [19] Hendricks, N. P., and Peterson, J. M. (2012). Fixed effects estimation of the intensive and extensive margins of irrigation water demand. *Journal of Agricultural and Resource Economics*, 37(1), 1–19.
- [20] Mieno, T., and Brozović, N. (2017). Price elasticity of groundwater demand: attenuation and amplification bias due to incomplete information. *American Journal of Agricultural Economics*, 99(2), 401–426.
- [21] Lobell, D. B., Cassman, K. G., and Field, C. B. (2009). Crop yield gaps: their importance, magnitudes, and causes. *Annual Review of Environment and Resources*, 34, 179–204.