Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Staged and monitored groundwater management in the Mekong Delta under deep uncertainty

Livija Jakaite
University of Bedfordshire
Vitaly Schetinin
vitaly.schetinin@beds.ac.uk
University of Bedfordshire

Article

Keywords:

Posted Date: November 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11584104/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: Competing interest reported. V.S. is an Editorial Board Member for Scientific Reports. The authors declare no other competing interests.

Staged and monitored groundwater management in the Mekong Delta under deep uncertainty

L. Jakaite* and V. Schetinin*

School of Computer Science, University of Bedfordshire,

Park Street, Luton, LU1 3JU, UK

*Corresponding authors:

livija.jakaite@lj-analytics.co.uk

vitaly.schetinin@beds.ac.uk

Abstract

Groundwater in Vietnam's Mekong Delta is being depleted and salinised, yet aquifer conditions are only partly observed, and the rates of degradation and the effects of management are deeply uncertain. We model annual management as a manager acting on a belief about the aquifer's condition and choosing among six measures from the region's four-pillar strategy. Degradation is anchored to published head-decline rates and initial salinity to an open salinity model; unobserved inputs are varied in 1000 scenarios. Business as usual leaves aquifers saline in 17 % of years; active management reduces this to about 4 %. The efficient strategy is staged: wait while the aquifer is in good condition, recharge when it is stressed, and apply the full package only when it is critical. Applying the full package every year forgoes 19 % of net benefit. A monitoring survey that reaches managers before the annual decision raises net benefit by about 6 %, more than a more accurate later survey (4 %). Caution pays in 18 % of scenarios, mostly where degradation exceeds 2.5 times the published rates. Rules that look ahead with a belief have the lowest regret, and perfect knowledge of the dynamics is worth a third of a timely survey.

Introduction

Groundwater depletion is a global concern[1]. The Vietnamese Mekong Delta depends on groundwater for domestic supply and much of its economy, and decades of abstraction have lowered hydraulic heads, accelerated land subsidence and allowed saline water to move into aquifers[2–4]. In the delta's channels, salt intrusion is also advancing, driven by sediment starvation.[5] Because the delta lies only about 0.8 m above sea level on average,[6] extraction-induced subsidence adds directly to its exposure to sea-level rise. A recent review of the delta's groundwater[7] summarises a four-pillar response: risk-zoned extraction limits, a shift to surface-water supply, managed aquifer recharge (MAR), which is increasingly used worldwide to secure stressed aquifers,[8] and digital metering with economic instruments. How these measures should be sequenced, and how much protection is worth buying when aquifer conditions are only partly observed, is less clear.

Two features make the decision hard. First, the state of an aquifer is observed through fragmented and noisy monitoring, so managers act on beliefs rather than on known states. Second, the dynamics themselves are deeply uncertain:[9,10] published head-decline rates are delta-wide averages, the effectiveness of measures is known mainly from pilots, and salinisation may be near-irreversible once established. Predictions of groundwater management models are themselves uncertain.[11] Previous studies have addressed such uncertainty with stochastic and robust optimisation of pumping from coastal aquifers[12] and with Bayesian decision networks under data scarcity.[13] Two simple responses are common. A manager may optimise expected benefit year by year under the best current estimate, or may apply strong measures whenever failure is plausible, in the spirit of satisficing[14] and adaptive management.[15] Decision making under deep uncertainty (DMDU) offers tools to compare such strategies across many plausible futures,[16,17] such as robust decision making and scenario discovery,[18,19] adaptive pathways,[20] which complement each other,[21] and robustness metrics.[22] Recent work frames water planning under uncertainty as a control problem, in which policies map future observations to actions,[23] and shows that learning from observations can make flexible, staged infrastructure plans worthwhile, including under groundwater uncertainty in data-scarce regions.[24,25] DMDU is, however, rarely combined with an explicit treatment of partial observability, in

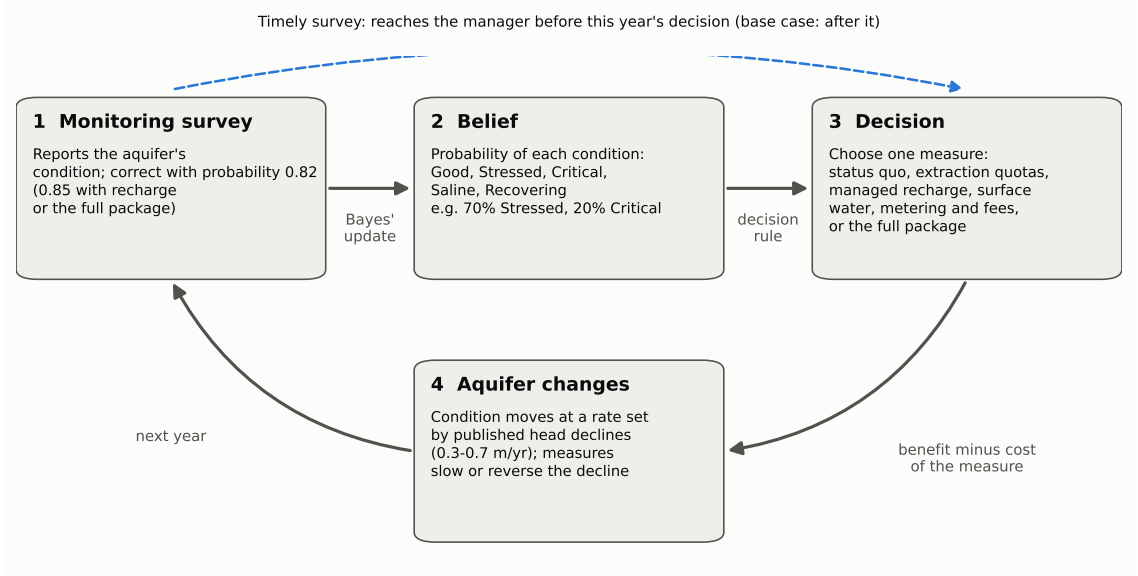


Figure 1. The yearly management cycle. A monitoring survey updates the manager's belief over five aquifer conditions; a decision rule chooses one of six measures; the aquifer's condition then changes, and the measure earns its benefit minus its cost. In the base case the survey of year t arrives after the decision of year t; in the "timely survey" variant (dashed) it informs that decision.

which the value of monitoring appears directly.

Partially observable Markov decision processes (POMDPs) provide that treatment.[26,27] In conservation they have shown when to stop managing and start surveying a species whose presence is uncertain,[28] and in fisheries that POMDP-based management avoids the overexploitation to which both precautionary and simplistic optimal-control rules are prone under measurement uncertainty.[29] In plain terms, the manager keeps a belief (a probability for each possible aquifer condition), updates it after each monitoring survey, and chooses measures both for their immediate effect and for how they change the aquifer's future condition (Fig. 1). This paper uses such a model of groundwater management in the Mekong Delta, built only on open data, to answer four questions: which management strategy is efficient when degradation proceeds at published rates; how much monitoring is worth, and whether its timing or its accuracy matters more; when caution pays, given that the planning model may be wrong; and how much better knowledge of the aquifer dynamics (calibration data) would be worth.

A stylised model is informative here for the following reasons. Local monitoring series for the delta are fragmented and not publicly available, and the effects and costs of the four pillars are known mainly from pilots. Rather than calibrate a

detailed model to uncertain inputs, we ask comparative questions: which decision rule, and which information, is worth more. We then test how sensitive the answers are to every input that is not observed (Table 1). Degradation rates are anchored to published head-decline rates and initial salinity to an open salinity model. The remaining hand-set inputs are varied over wide ranges in 1000 scenarios. Two inputs, the degradation rate and the cost of measures, decide whether caution pays, and we report the conditions under which it does. The others hardly change the conclusions. A value-of-information analysis shows that perfect knowledge of the dynamics would be worth about a third of a timely monitoring survey. The conclusions are therefore robust to the model's stylisation, and the analysis identifies which local data would matter most.

We find that the efficient strategy is staged and monitoring-led; that the timing of monitoring matters more than its accuracy; that caution pays only if aquifers degrade much faster than published rates suggest; and that rules which look ahead with a belief are robust across 1000 scenarios, so that timely monitoring is worth about three times as much as perfect calibration data.

Results

The yearly management cycle

Each year, a groundwater manager responsible for an aquifer receives a monitoring survey of uncertain reliability (Fig. 1). From the surveys so far, the manager maintains a *belief*: the probability that the aquifer is in good, stressed, critical, saline or recovering condition (for example, "70 % stressed, 20 % critical"). The manager then chooses a measure: do nothing new, set extraction quotas, recharge the aquifer, switch to surface water, introduce metering and fees, or apply all four pillars. The aquifer's condition then changes, more or less favourably depending on the measure. Decision rules differ in how they use the belief: a *year-by-year optimiser* chooses the measure with the best expected result for this year; a *cautious rule* chooses only measures that are very likely to avoid losses; a *look-ahead rule* also accounts for how today's measure changes the aquifer's future condition. We compare these rules over 16-year futures, with degradation anchored to published head-decline rates (Meth-

Table 1. Provenance of the model inputs and their influence on the conclusions. Influence is measured by the standardised rank regression coefficient (SRRC) of the welfare gain of the cautious rule over the best targeted rule across 1000 scenarios (scenario discovery), the partial expected value of perfect information (EVPPI) and the regional runs (see Results).

Input	Base value	Source	Uncertainty explored	Influence on conclusions
Head-decline rate r	$0.5\,\mathrm{m\,yr^{-1}}$	Minderhoud et al.[3]; government assessment in Chau et al.[7]	r 0.3–$1.0\,\mathrm{m\,yr^{-1}}$; r/w 0.03–0.5	High (SRRC +0.59; EVPPI 6.7): sets where caution pays
Drawdown band width w	5 m	Assumption	2.5–10 m	Acts only through r/w
Initial salinity	Per province and aquifer	Open data[30]	65 province–aquifer units; P(Saline) 0.02–0.97	Ranking unchanged in all 65 units (SRRC +0.06)
Cost of measures	Normalised	Hand-set	×0.5–3	High (SRRC −0.67): sets where caution pays
Effectiveness of measures	Hand-set relative effects	Measures from Chau et al.[7]	25–100 % of base	Low (SRRC −0.07; EVPPI 1.9)
Saline persistence	0.80	Hand-set	0.60–0.95	Low (SRRC +0.16; EVPPI 0.3)
Rewards and Saline penalty	Normalised; −80	Hand-set	Penalty −40 to −160; shadow price 0–160 per Saline year	Low ($\leq$ 3 units; SRRC +0.15)
Monitoring accuracy	0.82	Assumption (fragmented monitoring)	0.4–0.99; survey before or after the decision	Changes the value of monitoring, not the ranking
Horizon and discount factor	16 years; 0.95	Assumption	Not varied	–

ods). Effects are reported in normalised units of discounted net benefit and as a percentage of the look-ahead rule's base-case net benefit (462.2 units).

The efficient strategy is staged

Business as usual leaves the aquifer saline in 17 % of years; every active rule reduces this to 4–6 % (Table 2). The best rules (near-optimal planner, look-ahead rule, year-by-year optimiser) are statistically close (planner minus year-by-year optimiser +0.7 to +1.6 units across three seeds) and share a staged pattern: they wait in about 40 % of years, use MAR, quotas or metering when the belief indicates stress, and

Table 2. Base case. 5000 simulated 16-year futures (seed 2025; ± one standard error). Net benefit is discounted and in normalised units; the percentage is relative to the look-ahead rule. The action mix is the share of decisions.

Rule	Net benefit	% of look-ahead	Saline years (%)	Saline at $t = 15$ (%)	Action mix
Near-optimal planner	462.5 ± 0.9	100.1	4.4	2.3	wait 38, MAR 28, quotas 4, full 30
Look-ahead rule	462.2 ± 1.0	100.0	4.2	2.0	wait 41, MAR 16, quotas 2, full 41
Year-by-year optimiser	461.9 ± 1.0	99.9	4.2	1.9	wait 41, MAR 12, metering 4, full 43
Blended rule	455.6 ± 0.9	98.6	4.0	1.4	wait 40, full 59
Cautious rule	452.8 ± 0.9	98.0	3.9	1.3	wait 39, full 61
Always MAR	399.6 ± 1.2	86.4	6.0	2.3	MAR 100
Always full package	374.6 ± 1.1	81.0	3.6	0.8	full 100
Always status quo	−76.0 ± 4.9	−16.4	17.2	15.0	wait 100
Oracle (state known before acting)	515.8 ± 0.8	111.6	3.8	1.2	wait 54, full 45

apply the full package in 30–43 % of years (Supplementary Fig. S10). Applying the full package every year buys 0.6 percentage points less saline exposure at the cost of 88 units, 19 % of the net benefit. The cautious rule earns 7.1–9.0 units (1.5–2.0 %) less than the year-by-year optimiser across seeds (all $p < 0.001$); the blended rule recovers 2.5–2.7 of these units. The price of caution rises with costs, from about 2.5 units at half the base costs to 94–102 at three times the base costs (Supplementary Fig. S5). The near-optimal planner's value (462.0 at the initial belief) is stable to within 0.12 units across belief sets from 100 to more than 11,000 points per year, including at the extremes of the degradation range, so the gap to the Oracle (53–54 units, about 12 %) reflects information timing, not planning error (Supplementary Fig. S6 and Table S5).

Monitoring: timing matters more than accuracy

With the survey arriving after the decision, raising accuracy from 0.82 to 0.99 adds 18 units (3.9 %) for the look-ahead rule (Supplementary Fig. S3), less than the discounted cost of a metering programme (34 units). A survey of unchanged accuracy that arrives *before* the annual decision adds 27–28 units across seeds (5.8–6.1 %), about one and a half times as much. With timely surveys, accuracy also becomes worth more ($0.6 \to 0.99$: +57 against +41 units), and at accuracy 0.99 the belief-based rules come within 2 units of the Oracle (Fig. 2).

When the model is wrong, caution becomes insurance

When all rules plan with the base model but the aquifer follows different dynamics (Supplementary Fig. S11), the year-by-year optimiser keeps a significant advantage

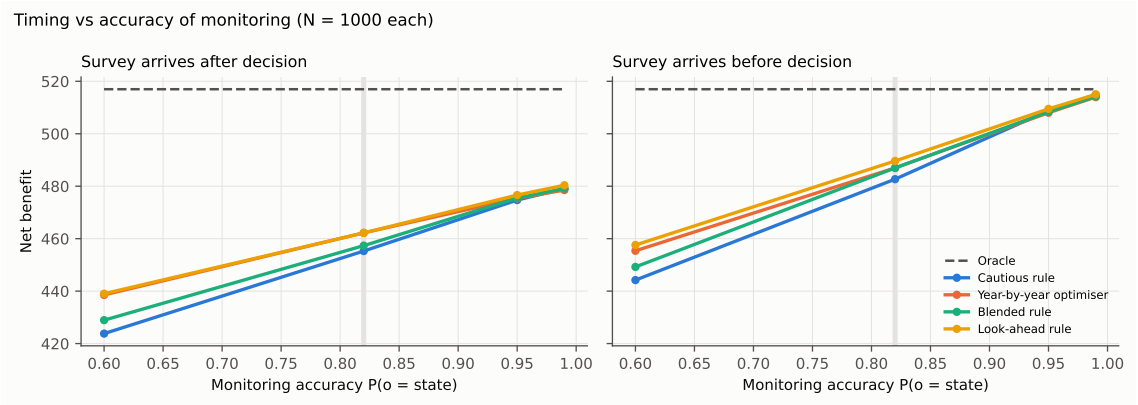


Figure 2. Value of monitoring accuracy and timing. Net benefit against monitoring accuracy when the survey arrives after (left) or before (right) the annual decision (1000 futures per point).

over the cautious rule in eight of twelve scenarios, including more or less persistent salinity and hidden tipping points. If measures are only 25–50 % as effective as assumed, or in a combined pessimistic case, the two rules are statistically tied. If the aquifer degrades five times faster than published rates suggest, the cautious rule wins by 16.8 units (95 % CI 14.2–19.4) and keeps saline exposure at 12.3 % of years against 15.0 %.

Caution pays only under fast degradation

Across 1000 scenarios, caution beats the best targeted rule in 18.0 % of scenarios, ties in 19.8 % and loses in 62.2 % (an independent sample of 500 scenarios gives 18.8 %, 14.6 % and 66.6 %). Two factors decide the outcome: cost (standardised rank coefficient -0.67) and degradation rate ($+0.59$; $R^2 = 0.82$). Caution wins in 94 % of the scenarios with $r/w > 0.245$ (2.5 times the base rate) and costs at most 1.6 times the base, and loses in 86 % of scenarios at or below the base rate (Fig. 3). The effectiveness of measures and initial salinity hardly matter.

The look-ahead rules are the robust choice (Table 3): the near-optimal planner and the simple look-ahead rule have the lowest mean, 95th-percentile and maximum regret (mean under 1 % of net benefit), even though they plan with the base model. The year-by-year optimiser, efficient in the base case, has the largest worst-case regret of the belief-based rules, because it does not weigh the future cost of saline years and loses most when costs are high and degradation is fast. The blended rule roughly halves the year-by-year optimiser's worst-case regret and reduces the

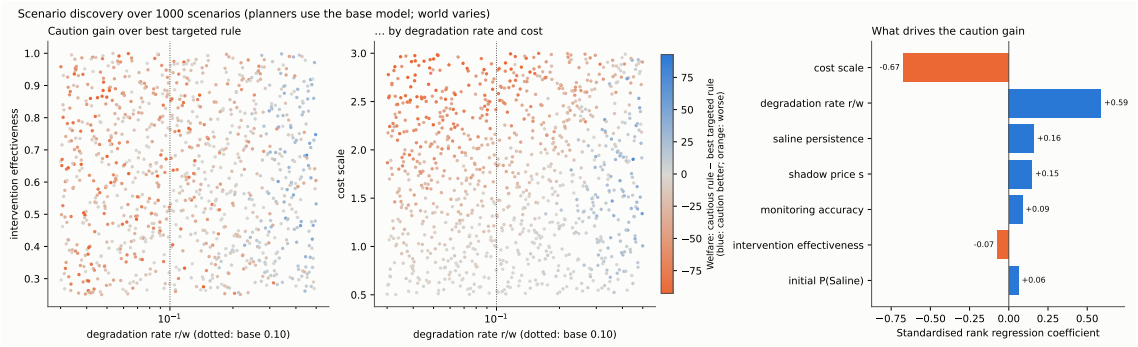


Figure 3. Scenario discovery over 1000 scenarios. Welfare of the cautious rule minus the best targeted rule against degradation rate, effectiveness of measures and cost, and the standardised rank regression coefficients of its drivers.

Table 3. Regret over 1000 scenarios. Regret is the best welfare among the seven rules minus the rule's welfare, in units and, for the mean, as a percentage of the look-ahead rule's base-case net benefit.

Rule	Mean	Mean (%)	Median	95th percentile	Maximum	Best in (%)
Near-optimal planner	4.0	0.9	0.4	23.4	76.5	36.7
Look-ahead rule	4.2	0.9	0.8	23.0	81.6	26.3
Blended rule	16.6	3.6	8.8	54.4	108.2	10.5
Cautious rule	22.8	4.9	12.8	80.0	151.0	16.3
Year-by-year optimiser	28.9	6.3	18.7	95.7	200.7	6.2
Always full package	86.8	18.8	70.0	215.3	294.4	3.7
Always MAR	95.1	20.6	90.8	180.7	308.4	0.3

cautious rule's by 28 %. The independent replicate gives the same ordering of mean and 95th-percentile regret for all seven rules (Supplementary Table S4).

Value of calibration data

Perfect knowledge of the dynamics is worth 8.0 welfare units on average (1.7 %; median 3.7; 90th percentile 23.5). Knowing the degradation rate alone gives 84 % of this (6.7); the effectiveness of measures (1.9) and saline persistence (0.3) add little. A pre-decision survey is worth 26.7 on average (5.8 %) and is positive in every scenario, about three times the expected value of perfect calibration (Fig. 4). Calibration data become valuable only when the true degradation rate is far above the published rates, and better monitoring accuracy reduces their value (rank coefficient −0.23). In the independent replicate, EVPI is 7.2 on average (median 3.1), knowing the degradation rate is worth 6.8, and a pre-decision survey 26.2.

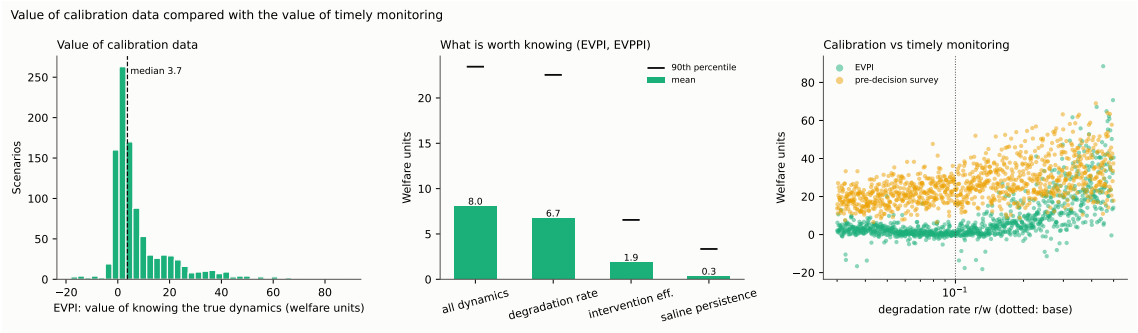


Figure 4. Value of calibration data. Distribution of EVPI, partial EVPI by factor, and EVPI and the value of a pre-decision survey against the degradation rate (1000 scenarios).

Regional robustness

Starting the simulation from the salinity of each of the 65 province–aquifer units (Fig. 5) does not change the ranking: the near-optimal planner is the best rule in all units, and the year-by-year optimiser beats the cautious rule by 4.0–7.5 units in every unit (all $p \leq 0.0001$). Initial salinity changes how much is at stake (net benefit 432–466; saline years 3–13 %), not which rule to use. Across published degradation rates and band widths (r/w 0.03–0.40) the ranking is unchanged, a pre-decision survey is worth 20–39 units and the price of caution is 5–32 units (Supplementary Fig. S9).

Discussion

Where degradation proceeds at published rates, the efficient strategy is to stage the measures: wait while aquifers are in good condition, recharge when they are stressed, and apply the full package only when they are critical. This is nearly optimal, whereas applying the full package everywhere and every year buys little extra protection and forgoes about a fifth of the net benefit.

Staging depends on knowing when an aquifer moves from one condition to another, so the timing of information matters. The largest gain available, about 6 % of net benefit, comes from surveys that reach managers before the annual decision. This supports the digital-metering and integrated-forecasting pillar,[7] with emphasis on the timing of information rather than its precision. Of what is unknown about the dynamics, the degradation rate is the parameter most worth measuring, and

routine head-trend monitoring is the calibration data with the highest value.

How the information is used matters as much. Rules that look ahead with a belief, including the simple look-ahead rule, hedge automatically: they become more protective as evidence of degradation accumulates. They are the robust choice across 1000 scenarios, whereas both the year-by-year optimiser and the fixed cautious rule have large worst-case regret. Caution remains a reasonable insurance where degradation may be much faster than published delta-wide averages; scenario discovery locates this region (degradation more than about 2.5 times the base rate, at moderate costs).

The results connect DMDU practice with partial observability. Scenario discovery[19] and regret-based robustness[22] identify where a cautious strategy is justified, while the belief-based model makes the value of monitoring explicit. Adaptive pathways[20] prescribe signposts and triggers; here the belief plays the role of a continuous signpost, and the value of a timely survey quantifies how much an earlier signpost is worth. As in planning studies that value learning,[23,25] the benefit of a staged strategy here depends on observations arriving in time to act on them. The result also parallels fisheries management, where neither fixed precautionary margins nor simple optimal control was sufficient under measurement uncertainty.[29] The finding that the robust rules look ahead with a belief rather than satisfice cautiously suggests that robustness can come from learning and timely information as much as from protective margins.

The study has limitations. The degradation rates are delta-wide averages mapped to states through an assumed band width; the relative effects of the measures, their costs and the rewards are normalised and hand-set; the five-state model abstracts a spatially heterogeneous, multi-aquifer system; the salinity model represents conditions up to about 2015; and no local monitoring series or stakeholder elicitation were used. Table 1 and scenario discovery treat these uncertainties explicitly, but the conclusions are conditional on the ranges sampled. Calibration to provincial monitoring data, starting with head trends, and a spatially explicit, multi-aquifer version are natural extensions.

In summary, for groundwater management in the Mekong Delta under deep uncertainty, a staged, monitoring-led strategy planned with look-ahead belief updating

is efficient at published degradation rates and robust across a wide range of plausible dynamics. Caution pays only if aquifers degrade much faster than published rates suggest, and the most valuable information is timely monitoring, followed by head-trend data that pin down the degradation rate.

Methods

Study area and data

The study covers the 13 provinces of the Vietnamese Mekong Delta (boundaries as of 2020). All inputs are open:

- *Groundwater salinity.* The three-dimensional model of Gunnink et al.[30,31] (CC-BY 4.0) gives, for $1\,\mathrm{km} \times 1\,\mathrm{km} \times 5\,\mathrm{m}$ voxels, the expected total dissolved solids (TDS) and the probability that TDS is below $1\,\mathrm{g\,L^{-1}}$, for the aquifers from the Holocene (qh) to the Early Pliocene (n_2^1), based on data collected from the 1980s to 2015. We use the expected share of each aquifer that is not fresh ($\mathrm{TDS} \geq 1\,\mathrm{g\,L^{-1}}$) per province (Fig. 5).

- *Degradation rates.* Hydraulic heads in the delta's aquifers declined by about 0.3–$0.7\,\mathrm{m\,yr^{-1}}$ on average over 1991–2016;[3] an earlier government assessment, summarised by Chau et al.,[7] reported 0.3–$0.5\,\mathrm{m\,yr^{-1}}$.

- *Measures and context.* The action set follows the four-pillar strategy and the regulatory framework summarised by Chau et al.[7]

- *Boundaries.* Province and district polygons from geoBoundaries[32] (OCHA ROAP and Government of Viet Nam, CC BY 3.0 IGO).

The pumped Pleistocene aquifers ($\mathrm{qp}_{2\text{-}3}$, qp_1) are mostly fresh in a southern band (Hau Giang, Soc Trang, Bac Lieu, Ca Mau: 10–26 % not fresh for $\mathrm{qp}_{2\text{-}3}$) and mostly brackish or saline upstream and in the north-east (Long An, Tien Giang, Ben Tre, Dong Thap: 82–91 %) (Fig. 5; Supplementary Table S1 and Fig. S8). No local monitoring time series were used.

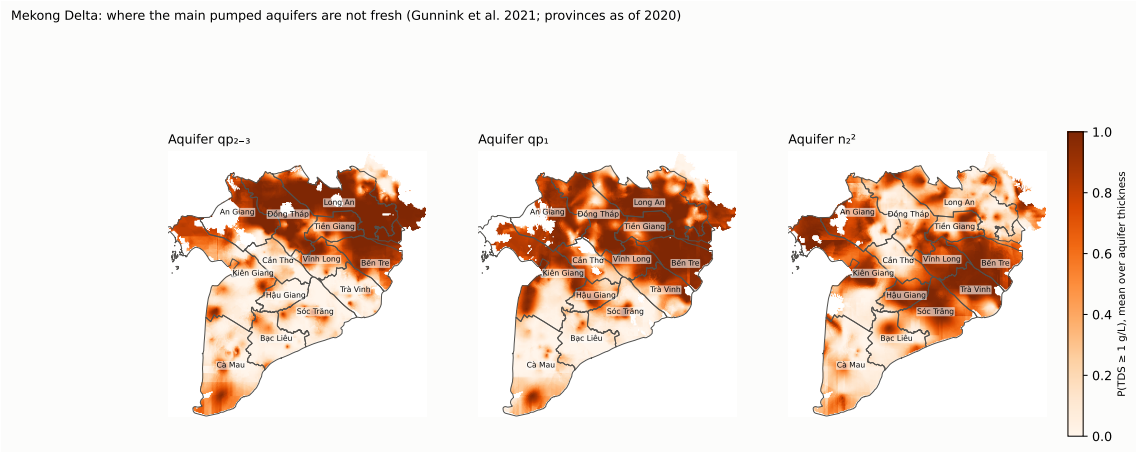


Figure 5. Groundwater salinity in the Vietnamese Mekong Delta. Probability that groundwater is not fresh ($\mathrm{TDS} \geq 1\,\mathrm{g\,L^{-1}}$), averaged over aquifer thickness, for the three main pumped aquifers, with province boundaries (2020). Data: Gunnink et al.[31]

Decision model

Each year $t = 0, \ldots, 15$ the manager chooses one of six measures a: status quo, extraction quotas, MAR, surface water (inter-provincial conveyance and well sealing), metering and fees, or the full package of all four pillars. The aquifer is in one of five states θ: Good, Stressed, Critical, Saline or Recovering. The measure earns a reward $R(a, \theta)$ and incurs a cost $C(a)$, both in normalised units; outcomes are evaluated as the discounted net benefit $\sum_t 0.95^t \, [R(a_t, \theta_t) - C(a_t)]$. To make effects comparable, we also report them as a percentage of the look-ahead rule's base-case net benefit (462.2 units). After acting, the manager observes a noisy signal of the current state, correct with probability 0.82 (0.85 when MAR or full-package infrastructure is in place), and the aquifer moves according to $P(\theta' \mid \theta, a)$.

Anchoring degradation to published rates. Treating Good, Stressed and Critical as drawdown bands of width w, a head decline of $r\,\mathrm{m\,yr^{-1}}$ moves an aquifer one band down with yearly probability r/w under the status quo. Starting from a hand-set kernel in which the measures' relative effects were specified, every transition into a more severe state is rescaled so that the status quo's yearly worsening probability from Stressed equals r/w. The base case is $r = 0.5\,\mathrm{m\,yr^{-1}}$ and $w = 5\,\mathrm{m}$ ($r/w = 0.10$). Saline is near-absorbing under the status quo ($P(\text{Saline} \to \text{Saline}) = 0.80$). The original hand-set kernel, which implies $r/w = 0.5$, is retained as a fast-degradation

stress case. The full parameter set is given in Supplementary Section S1 (Tables S2 and S3).

Provenance of the inputs. Table 1 lists, for every input, whether it comes from data or is assumed, how its uncertainty is explored, and how much it changes the conclusions.

Decision rules

All belief-based rules update the belief with Bayes' rule after each survey and project it forward with the transition model. Technical details are given in Supplementary Information, Section S2.

- *Year-by-year optimiser* ("prescriptive"): chooses the measure with the highest expected benefit net of cost this year, under a narrow prior.
- *Cautious rule* ("philosophical"): satisficing; considers only measures whose benefit is positive with probability above 0.9, gives credit to keeping options open and to learning, and uses a broad prior and a model-averaged likelihood.
- *Blended rule*: weights the two rules above according to how uncertain the belief is, leaning towards caution when the belief is diffuse (Supplementary Fig. S4).
- *Look-ahead rule* (QMDP[33]): values each measure by its long-run effect, averaging over the belief.
- *Near-optimal planner* (point-based value iteration, PBVI[34]): solves the full belief-based planning problem approximately, and serves as a benchmark.
- Reference rules: always status quo, always MAR, always the full package, and an Oracle that knows the true state before acting (an upper bound).

Experiments and statistics

1. *Base case*: all rules, 5000 simulated futures, three random seeds.

2. *Monitoring*: accuracy 0.4–0.99, with the survey arriving after or before the annual decision.

3. *Model misspecification*: rules plan with the base model while the aquifer follows one of 12 alternative dynamics (weaker measures, different saline persistence, hidden tipping points, degradation five times faster).

4. *Robustness*: degradation rates $r = 0.3$–$1.0\,\mathrm{m\,yr^{-1}}$ and $w = 2.5$–$10\,\mathrm{m}$; initial salinity from each of 65 province–aquifer units; costs; Saline shadow prices (Supplementary Figs. S2 and S7).

5. *Scenario discovery*: a Latin hypercube of 1000 scenarios varying degradation rate (r/w 0.03–0.5, log-uniform), effectiveness of measures (0.25–1), saline persistence (0.60–0.95) and initial salinity (0.02–0.97), which planners do not know, and cost scale (0.5–3), monitoring accuracy (0.60–0.95) and a shadow price s (0–160) per Saline year, which they do. Welfare is net benefit minus s times discounted Saline years. Regions where caution wins are found with classification trees;[19] drivers with standardised rank regression; robustness with regret.[22]

6. *Value of information*: for every scenario, the near-optimal planner is given the true dynamics (expected value of perfect information, EVPI) or one true factor (partial EVPI, EVPPI),[35] the Bayesian data-worth logic long used in groundwater management,[36] and compared with the same planner using the base model.

Rules are compared on common random numbers. Paired differences are reported with bootstrap 95 % confidence intervals and Holm-corrected paired t-tests. An earlier, exploratory version of the model had three faults, now corrected: actions did not affect the aquifer, costs were not counted, and beliefs were not propagated. Their effects are documented in Supplementary Section S3 and Fig. S1.

Data availability

All data used are openly available: groundwater salinity[30] (CC-BY 4.0); province and district boundaries from geoBoundaries[32] (CC BY 3.0 IGO); head-decline rates

from Minderhoud et al.[3] and Chau et al.[7]. Simulation outputs are provided with the code.

Code availability

Python code (simulator, decision rules, PBVI solver, experiments, figures) and simulation outputs are openly available on Zenodo under the MIT licence, https://doi.org/10.5281/zenodo.23849418 (version 1.0.0),[37] and on GitHub at https://github.com/ljakaite/mekong-groundwater-decisions.

References

1. Famiglietti, J. S. The global groundwater crisis. *Nature Climate Change* **4**, 945–948 (2014).

2. Erban, L. E., Gorelick, S. M. & Zebker, H. A. Groundwater extraction, land subsidence, and sea-level rise in the Mekong Delta, Vietnam. *Environmental Research Letters* **9**, 084010 (2014).

3. Minderhoud, P. S. J. *et al.* Impacts of 25 years of groundwater extraction on subsidence in the Mekong delta, Vietnam. *Environmental Research Letters* **12**, 064006 (2017).

4. Minderhoud, P. S. J., Middelkoop, H., Erkens, G. & Stouthamer, E. Groundwater extraction may drown mega-delta: projections of extraction-induced subsidence and elevation of the Mekong delta for the 21st century. *Environmental Research Communications* **2**, 011005 (2020).

5. Eslami, S. *et al.* Tidal amplification and salt intrusion in the Mekong Delta driven by anthropogenic sediment starvation. *Scientific Reports* **9**, 18746 (2019).

6. Minderhoud, P. S. J., Coumou, L., Erkens, G., Middelkoop, H. & Stouthamer, E. Mekong delta much lower than previously assumed in sea-level rise impact assessments. *Nature Communications* **10**, 3847 (2019).

7. Chau, T. T. N., Nam, N. D. G., Minh, H. V. T., Y, L. N. & Ngan, N. V. C. Groundwater resources in the Mekong Delta: Current status, challenges, and sustainable approaches under dual impacts. *International Journal of Research and Scientific Innovation* **XII**, 699–712 (2025).

8. Dillon, P. *et al.* Sixty years of global progress in managed aquifer recharge. *Hydrogeology Journal* **27**, 1–30 (2019).

9. Walker, W. E., Lempert, R. J. & Kwakkel, J. H. Deep uncertainty. In *Encyclopedia of Operations Research and Management Science*, 395–402 (Springer, 2013).

10. Kwakkel, J. H., Walker, W. E. & Haasnoot, M. Coping with the wickedness of public policy problems: Approaches for decision making under deep uncertainty. *Journal of Water Resources Planning and Management* **142**, 01816001 (2016).

11. Delottier, H., Pryet, A. & Dupuy, A. Why should practitioners be concerned about predictive uncertainty of groundwater management models? *Water Resources Management* **31**, 61–73 (2017).

12. Sreekanth, J. & Datta, B. Stochastic and robust multi-objective optimal management of pumping from coastal aquifers under parameter uncertainty. *Water Resources Management* **28**, 2005–2019 (2014).

13. Mohajerani, H. *et al.* Application of Bayesian decision networks for groundwater resources management under the conditions of high uncertainty and data scarcity. *Water Resources Management* **31**, 1859–1879 (2017).

14. Simon, H. A. A behavioral model of rational choice. *The Quarterly Journal of Economics* **69**, 99–118 (1955).

15. Williams, B. K. Adaptive management of natural resources—framework and issues. *Journal of Environmental Management* **92**, 1346–1353 (2011).

16. Marchau, V. A. W. J., Walker, W. E., Bloemen, P. J. T. M. & Popper, S. W. (eds.) *Decision Making under Deep Uncertainty: From Theory to Practice* (Springer, Cham, 2019).

17. Cunha, M. d. C. Water and environmental systems management under uncertainty: From scenario construction to robust solutions and adaptation. *Water Resources Management* **37**, 2271–2285 (2023).

18. Lempert, R. J., Groves, D. G., Popper, S. W. & Bankes, S. C. A general, analytic method for generating robust strategies and narrative scenarios. *Management Science* **52**, 514–528 (2006).

19. Bryant, B. P. & Lempert, R. J. Thinking inside the box: A participatory, computer-assisted approach to scenario discovery. *Technological Forecasting and Social Change* **77**, 34–49 (2010).

20. Haasnoot, M., Kwakkel, J. H., Walker, W. E. & ter Maat, J. Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. *Global Environmental Change* **23**, 485–498 (2013).

21. Kwakkel, J. H., Haasnoot, M. & Walker, W. E. Comparing robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty. *Environmental Modelling & Software* **86**, 168–183 (2016).

22. Herman, J. D., Reed, P. M., Zeff, H. B. & Characklis, G. W. How should robustness be defined for water systems planning under change? *Journal of Water Resources Planning and Management* **141**, 04015012 (2015).

23. Herman, J. D., Quinn, J. D., Steinschneider, S., Giuliani, M. & Fletcher, S. Climate adaptation as a control problem: Review and perspectives on dynamic water resources planning under uncertainty. *Water Resources Research* **56**, e24389 (2020).

24. Fletcher, S., Lickley, M. & Strzepek, K. Learning about climate change uncertainty enables flexible water infrastructure planning. *Nature Communications* **10**, 1782 (2019).

25. Fletcher, S., Strzepek, K., Alsaati, A. & de Weck, O. Learning and flexibility for water supply infrastructure planning under groundwater resource uncertainty. *Environmental Research Letters* **14**, 114022 (2019).

26. Kaelbling, L. P., Littman, M. L. & Cassandra, A. R. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* **101**, 99–134 (1998).

27. Chadès, I., Pascal, L. V., Nicol, S., Fletcher, C. S. & Ferrer-Mestres, J. A primer on partially observable Markov decision processes (POMDPs). *Methods in Ecology and Evolution* **12**, 2058–2072 (2021).

28. Chadès, I. *et al.* When to stop managing or surveying cryptic threatened species. *Proceedings of the National Academy of Sciences* **105**, 13936–13940 (2008).

29. Memarzadeh, M. & Boettiger, C. Resolving the measurement uncertainty paradox in ecological management. *The American Naturalist* **193**, 645–660 (2019).

30. Gunnink, J., Van Pham, H., Oude Essink, G. & Bierkens, M. Three-dimensional hydrogeological and salinity model of the Mekong delta, Vietnam. Zenodo dataset, CC-BY 4.0, https://doi.org/10.5281/zenodo.4441776 (2021).

31. Gunnink, J. L., Pham, H. V., Oude Essink, G. H. P. & Bierkens, M. F. P. The three-dimensional groundwater salinity distribution and fresh groundwater volumes in the Mekong Delta, Vietnam, inferred from geostatistical analyses. *Earth System Science Data* **13**, 3297–3319 (2021).

32. Runfola, D. *et al.* geoBoundaries: A global database of political administrative boundaries. *PLOS ONE* **15**, e0231866 (2020).

33. Littman, M. L., Cassandra, A. R. & Kaelbling, L. P. Learning policies for partially observable environments: Scaling up. In *Machine Learning Proceedings 1995*, 362–370 (1995).

34. Pineau, J., Gordon, G. & Thrun, S. Point-based value iteration: An anytime algorithm for POMDPs. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, 1025–1032 (2003).

35. Eidsvik, J., Mukerji, T. & Bhattacharjya, D. *Value of Information in the Earth Sciences: Integrating Spatial Modeling and Decision Analysis* (Cambridge University Press, 2015).

36. Feyen, L. & Gorelick, S. M. Framework to evaluate the worth of hydraulic conductivity data for optimal groundwater resources management in ecologically sensitive areas. *Water Resources Research* **41** (2005).

37. Jakaite, L. & Schetinin, V. Mekong groundwater decisions: code and simulation outputs for 'Staged and monitored groundwater management in the Mekong Delta under deep uncertainty'. Zenodo, version 1.0.0, MIT licence, https://doi.org/10.5281/zenodo.23849418 (2026).

Acknowledgements

The authors thank Assoc. Prof. Dr. Van Pham Dang Tri, Director of the Mekong Institute at Can Tho University, Viet Nam, for his valuable comments and insights on this research. This research received no specific funding.

Author contributions

L.J.: conceptualisation, methodology, software, formal analysis, writing (original draft). V.S.: conceptualisation, methodology, validation, writing (review and editing), supervision. All authors reviewed the manuscript.

Competing interests

V.S. is an Editorial Board Member for *Scientific Reports*. The authors declare no other competing interests.

Additional information

Supplementary information accompanies this paper. **Correspondence** and requests for materials should be addressed to L.J. or V.S.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SREPMekongSupplementaryInformation.pdf