Supplementary Information

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

L. Jakaite and V. Schetinin

School of Computer Science, University of Bedfordshire, Luton, UK

S1 Model parameters and regional data

Rewards, costs and the observation model are normalised and hand-set (Table S2). The base-case transition kernel (Table S3) rescales every worsening transition of the original hand-set kernel so that the status quo's yearly worsening probability from Stressed equals $r/w = 0.10$ (head decline $r = 0.5\,\mathrm{m\,yr^{-1}}$, drawdown band $w = 5\,\mathrm{m}$); the original kernel ($r/w = 0.5$) is the fast-degradation stress case. Monitoring returns the current state with probability 0.82 (0.85 under MAR or the full package) and any other state with equal probability. Table S1 gives the salinity of each aquifer by province.

Table S1: Expected share of each aquifer that is not fresh (TDS $\geq$ 1 g $\mathrm{L^{-1}}$) by province (Gunnink et al., 2021); area as the union of 2020 districts ($\mathrm{km^2}$, 1 km cells) and the share covered by the salinity model.

Province	Area	Coverage	qh	$\mathrm{qp_3}$	$\mathrm{qp_{2\text{-}3}}$	$\mathrm{qp_1}$	$\mathrm{n_2^2}$	$\mathrm{n_2^1}$
Long An	4,839	1.00	0.97	0.82	0.91	0.78	0.38	0.29
Tien Giang	1,913	0.99	0.85	0.77	0.83	0.92	0.53	0.29
Ben Tre	2,348	0.98	0.91	0.76	0.85	0.97	0.89	0.80
Tra Vinh	2,352	0.95	0.96	0.42	0.16	0.63	0.85	0.95
Vinh Long	1,523	1.00	0.91	0.69	0.59	0.79	0.86	0.58
Dong Thap	3,385	1.00	0.98	0.67	0.82	0.82	0.39	0.38
An Giang	3,533	0.91	0.96	0.45	0.72	0.83	0.51	0.51
Kien Giang	5,582	0.97	0.86	0.51	0.29	0.44	0.60	0.68
Can Tho	1,735	1.00	0.79	0.24	0.20	0.48	0.27	0.44
Hau Giang	1,320	1.00	0.97	0.25	0.10	0.41	0.76	0.57
Soc Trang	3,298	0.99	0.90	0.57	0.14	0.20	0.69	0.81
Bac Lieu	2,477	1.00	0.99	0.74	0.14	0.11	0.24	0.66
Ca Mau	5,264	0.98	1.00	0.76	0.26	0.15	0.23	0.59

Table S2: Reward $R(a, \theta)$ and cost $C(a)$ (normalised units).

Action	Good	Stressed	Critical	Saline	Recovering	$C(a)$
Status quo	50	20	-30	-80	10	0
Extraction quotas	45	35	10	-20	25	2
Managed recharge	40	45	30	5	70	8
Surface water	35	40	35	15	60	12
Metering and fees	45	35	15	-10	30	3
Full package	40	50	40	25	80	15

Table S3: Base-case transition probabilities $P(\theta' \mid \theta, a)$, $r/w = 0.10$ (rows: from; columns: to).

Action	From	Good	Stressed	Critical	Saline	Recovering
Status quo	Good	0.920	0.056	0.016	0.004	0.004
	Stressed	0.050	0.800	0.070	0.030	0.050
	Critical	0.010	0.080	0.784	0.076	0.050
	Saline	0.000	0.030	0.120	0.800	0.050
	Recovering	0.150	0.070	0.050	0.030	0.700
Extraction quotas	Good	0.936	0.048	0.010	0.002	0.004
	Stressed	0.080	0.784	0.056	0.020	0.060
	Critical	0.020	0.120	0.704	0.056	0.100
	Saline	0.000	0.050	0.150	0.720	0.080
	Recovering	0.200	0.080	0.036	0.020	0.664
Managed recharge	Good	0.930	0.044	0.012	0.002	0.012
	Stressed	0.100	0.660	0.044	0.016	0.180
	Critical	0.030	0.120	0.556	0.044	0.250
	Saline	0.010	0.050	0.120	0.620	0.200
	Recovering	0.220	0.060	0.030	0.016	0.674
Surface water	Good	0.930	0.044	0.012	0.002	0.012
	Stressed	0.120	0.648	0.040	0.012	0.180
	Critical	0.040	0.140	0.504	0.036	0.280
	Saline	0.020	0.080	0.140	0.560	0.200
	Recovering	0.250	0.060	0.024	0.012	0.654
Metering and fees	Good	0.928	0.052	0.012	0.004	0.004
	Stressed	0.070	0.780	0.056	0.024	0.070
	Critical	0.020	0.100	0.716	0.064	0.100
	Saline	0.000	0.040	0.140	0.740	0.080
	Recovering	0.180	0.076	0.040	0.024	0.680
Full package	Good	0.940	0.036	0.012	0.002	0.010
	Stressed	0.150	0.604	0.036	0.010	0.200
	Critical	0.050	0.150	0.440	0.030	0.330
	Saline	0.020	0.080	0.120	0.480	0.300
	Recovering	0.280	0.056	0.020	0.010	0.634

Table S4: Regret by rule: original sample (1000 scenarios) and independent replicate (500 scenarios); bootstrap 95 % CIs.

Rule	Mean (original)	Mean (replicate)	95th pct (original)	95th pct (replicate)
Near-optimal planner	4.0 [3.5, 4.5]	4.1 [3.4, 5.0]	23.4 [18.9, 26.0]	24.1 [20.8, 27.8]
Look-ahead rule	4.2 [3.7, 4.8]	4.4 [3.6, 5.2]	23.0 [19.2, 27.1]	25.9 [21.5, 28.7]
Blended rule	16.6 [15.5, 17.8]	18.8 [17.2, 20.5]	54.4 [51.3, 57.0]	59.9 [55.5, 63.8]
Cautious rule	22.8 [21.2, 24.5]	25.3 [22.9, 27.8]	80.0 [73.5, 82.5]	90.5 [75.3, 94.1]
Year-by-year optimiser	28.9 [27.0, 31.0]	29.8 [27.1, 32.7]	95.7 [84.8, 108.7]	97.4 [90.0, 108.6]
Always full package	86.8 [82.7, 91.0]	92.6 [86.7, 99.3]	215.3 [208.1, 227.8]	228.0 [216.5, 241.0]
Always MAR	95.1 [92.3, 97.9]	100.8 [96.8, 105.1]	180.7 [171.2, 188.5]	189.7 [167.1, 208.5]

Table S5: Convergence: paired differences at 5000 futures for three random seeds (95 % CIs).

Comparison	Seed 2025	Seed 7	Seed 11
Year-by-year optimiser − Cautious rule	9.0 [8.2, 9.9]	7.5 [6.6, 8.4]	7.1 [6.2, 8.0]
Blended rule − Cautious rule	2.7 [2.3, 3.1]	2.5 [2.1, 2.9]	2.7 [2.3, 3.1]
Look-ahead rule − Year-by-year optimiser	0.3 [-0.3, 1.0]	1.1 [0.4, 1.8]	0.7 [0.0, 1.4]
Near-optimal planner − Year-by-year optimiser	0.7 [-0.1, 1.4]	1.6 [0.8, 2.4]	1.4 [0.6, 2.3]
Near-optimal planner − Look-ahead rule	0.3 [-0.5, 1.2]	0.5 [-0.4, 1.4]	0.7 [-0.2, 1.6]
Oracle − Near-optimal planner	53.2 [51.9, 54.5]	53.5 [52.2, 54.9]	54.1 [52.7, 55.4]
Look-ahead rule: survey before − after	27.0 [25.9, 28.0]	27.9 [26.8, 29.0]	28.3 [27.3, 29.4]

S2 Decision rules: technical details

The main text uses plain names for the rules; the technical names used in the code and results files are given in brackets. Let b_t be the belief over the five states, $R(a, \theta)$ the reward and $C(a)$ the cost. All belief-based rules update b_t with Bayes' rule after each survey and propagate it with the transition model, $b_{t+1} = b_t P(\cdot \mid \cdot, a_t)$.

- *Year-by-year optimiser* (Prescriptive): $a_t = \arg\max_a \sum_\theta b_t(\theta)\,[R(a,\theta) - \mu C(a)]$ with cost aversion $\mu = 0.8$ and a narrow prior.
- *Cautious rule* (Philosophical): utility $U(a, \theta) = R(a, \theta) + \lambda O(a)$, where $O(a)$ is a static option value computed from the prior ($\lambda = 1$). Only satisficing actions with $P_{b_t}(U > 0) > \beta = 0.9$ are considered (all actions if none qualifies); among them the rule maximises $b_t^\top U(a, \cdot) + \nu\, I(a; b_t)$, where I is the expected reduction in belief entropy from the next survey ($\nu = 1$). The rule uses a broad prior and a model-averaged likelihood.
- *Blended rule* (Integrated): combines the two utilities with an entropy-driven weight on the year-by-year utility,

  $$\alpha_t = \alpha_{\min} + (\alpha_{\max} - \alpha_{\min})\,[1 - \tanh(\kappa \bar{H}_t / 2)],$$

  where $\bar{H}_t$ is the belief entropy normalised to $[0, 1]$ ($\kappa = 5$, $\alpha_{\min} = 0.05$, $\alpha_{\max} = 0.95$). A diffuse belief gives a small α_t, so the rule leans towards caution; a concentrated belief gives a large α_t. The information term is weighted by $1 - \alpha_t$ and the likelihood is blended in the same proportions (Fig. S4).
- *Look-ahead rule* (QMDP; Littman et al., 1995): $a_t = \arg\max_a \sum_\theta b_t(\theta)\, Q_t(\theta, a)$, where Q_t are the finite-horizon action values of the fully observed problem with net reward $R - C$ and discount factor 0.95.
- *Near-optimal planner* (PBVI; Pineau et al., 2003): finite-horizon point-based value iteration on beliefs reachable under the planning model (collected from simulated episodes, optionally with a regular grid on the belief simplex). Its value at the initial belief is a lower bound on the optimum, and the QMDP value an upper bound; the bracket and its convergence are shown in Fig. S6 and Table S5.

S3 Corrections to the earlier version of the model

Figure S1 shows the effect of each correction on the kernel of an earlier, exploratory version of the model: letting actions change the aquifer, counting costs, and propagating beliefs.

S4 Supplementary figures

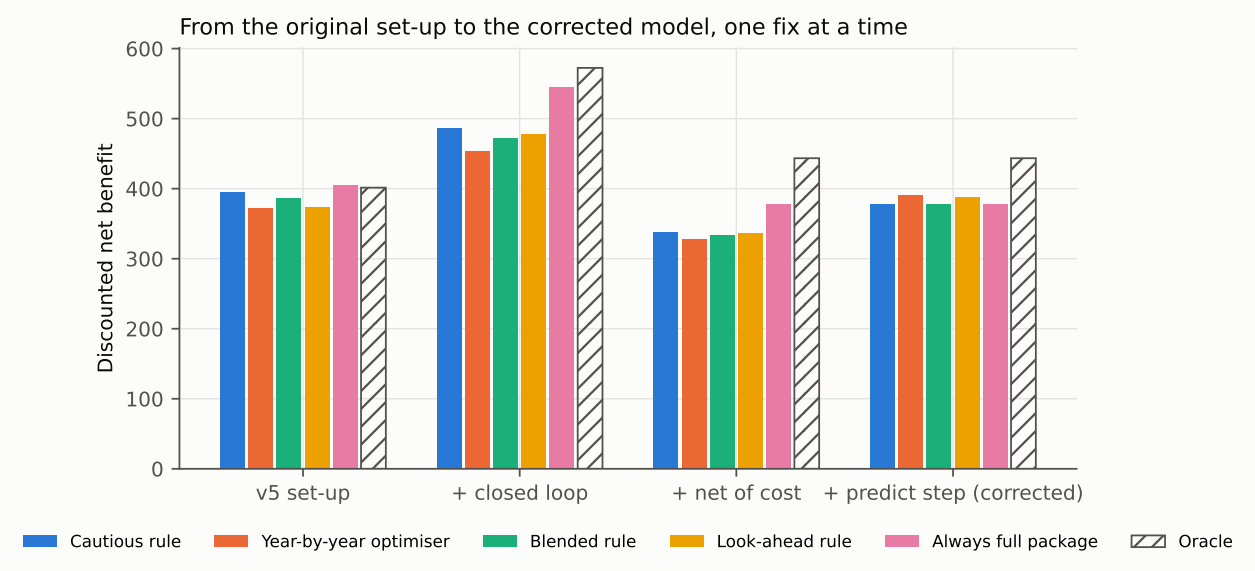


Figure S1: Effect of each correction to the earlier version of the model (earlier kernel).

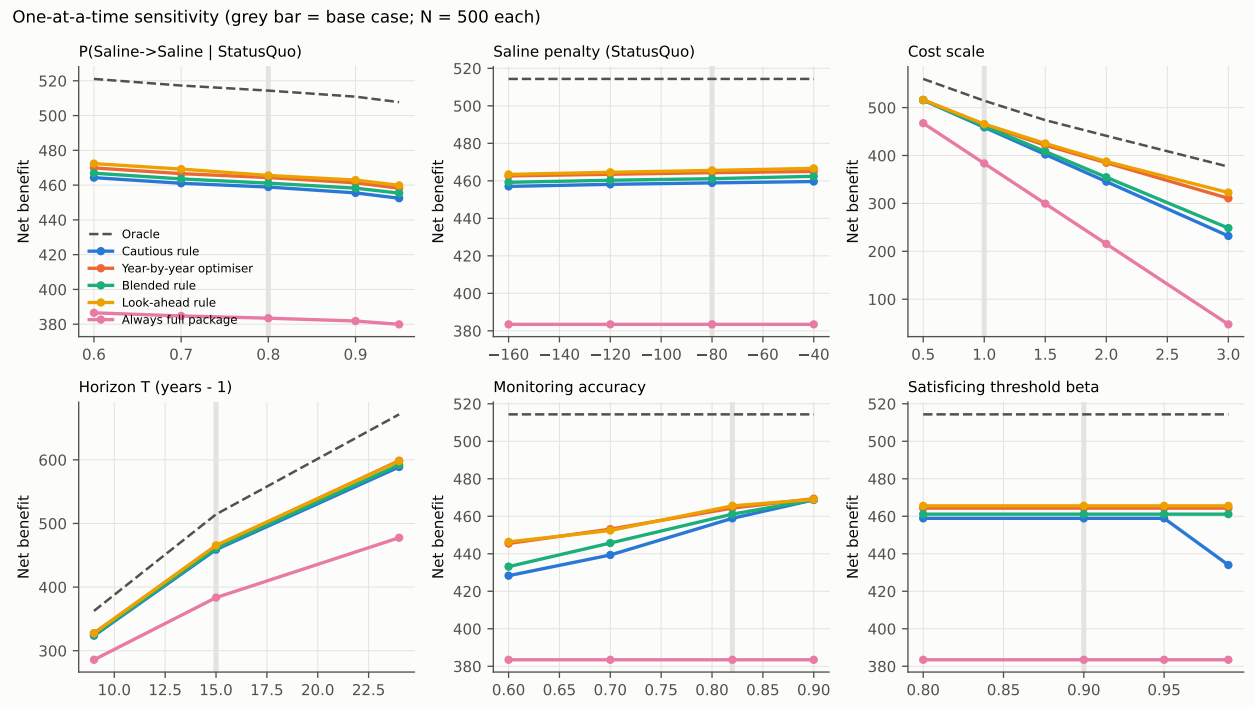


Figure S2: One-at-a-time sensitivity of the base case.

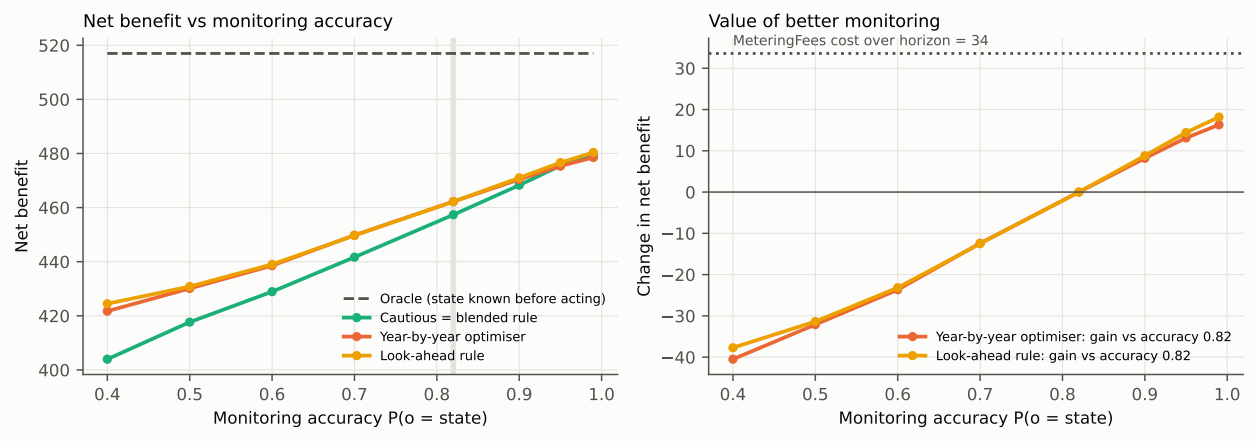


Figure S3: Value of more accurate monitoring when the survey arrives after the decision.

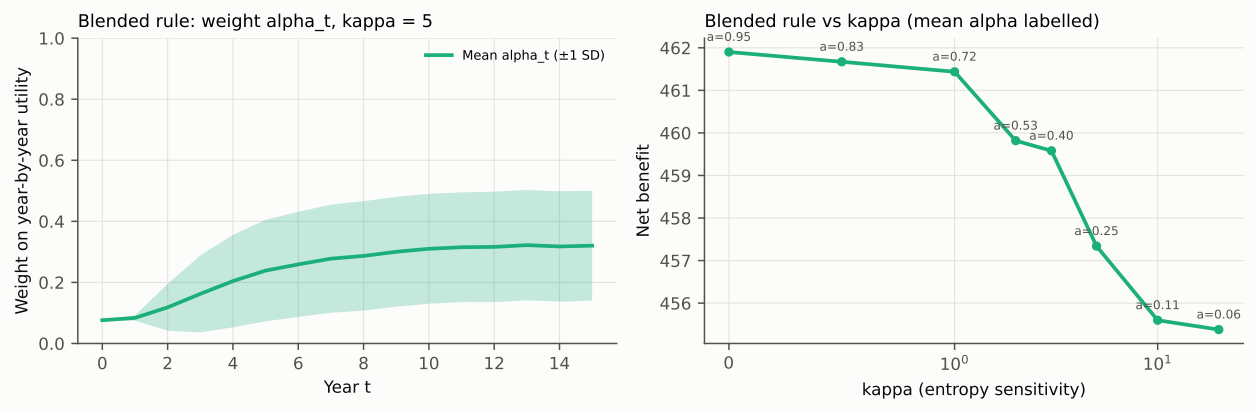


Figure S4: Blended rule: weight α_t on the year-by-year utility over time, and sensitivity to κ.

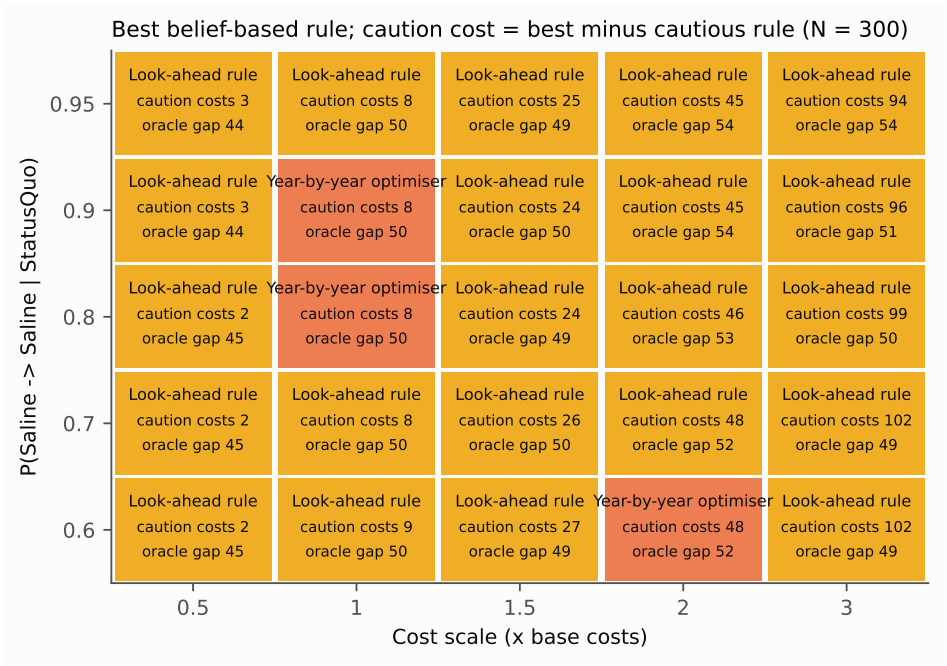


Figure S5: Best belief-based rule and price of caution over saline persistence and cost.

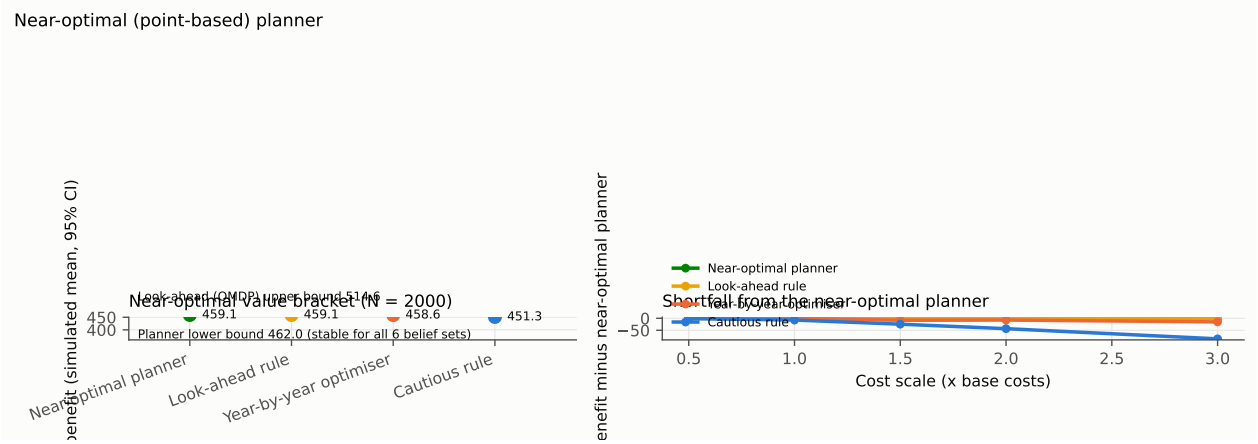


Figure S6: Near-optimal (point-based) planner: value bracket and shortfall of the other rules across costs.

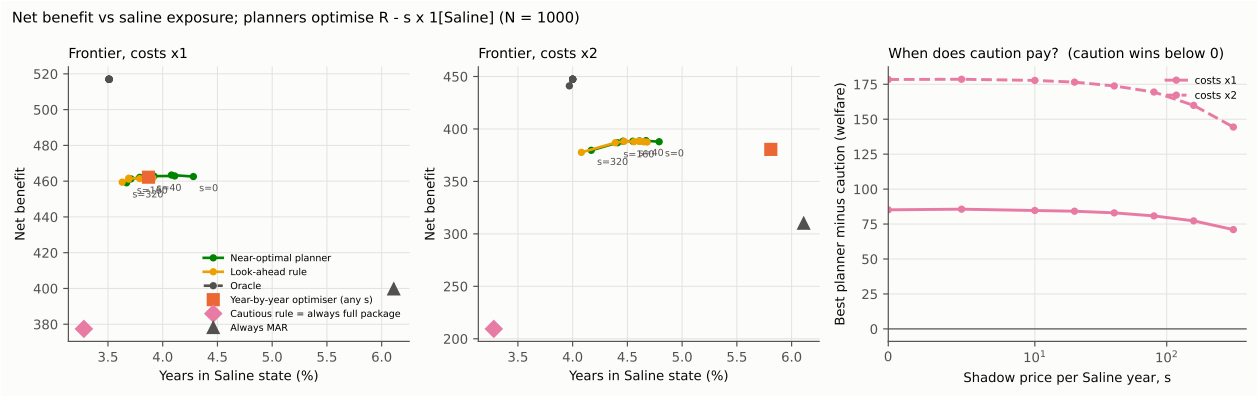


Figure S7: Net benefit against saline exposure when the near-optimal planner and the look-ahead rule price Saline years at s.

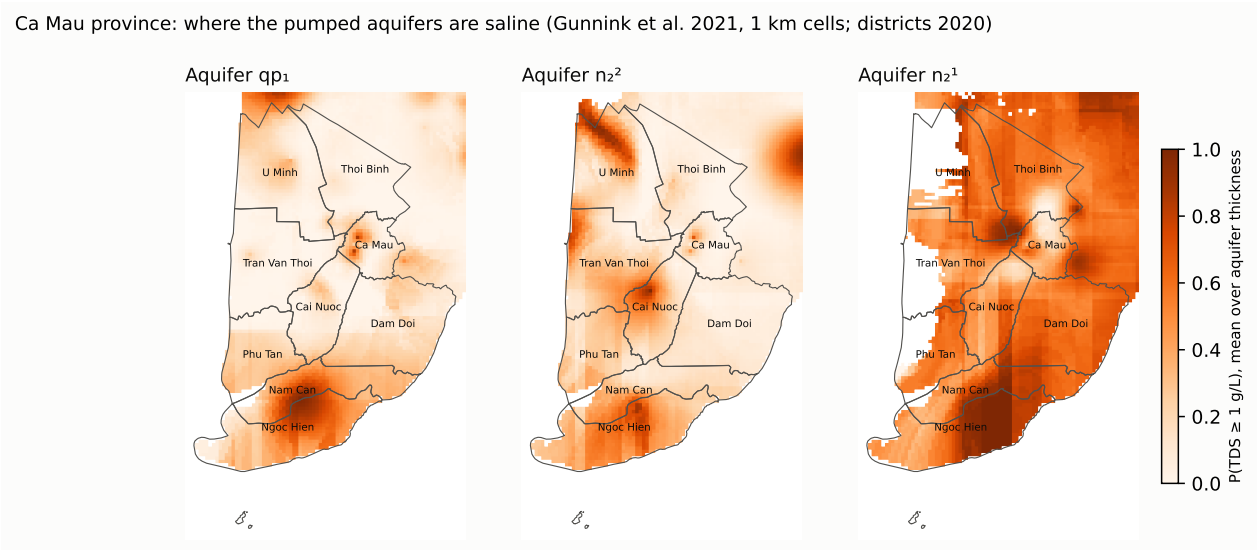


Figure S8: Ca Mau province: probability that the pumped aquifers are not fresh, by district.

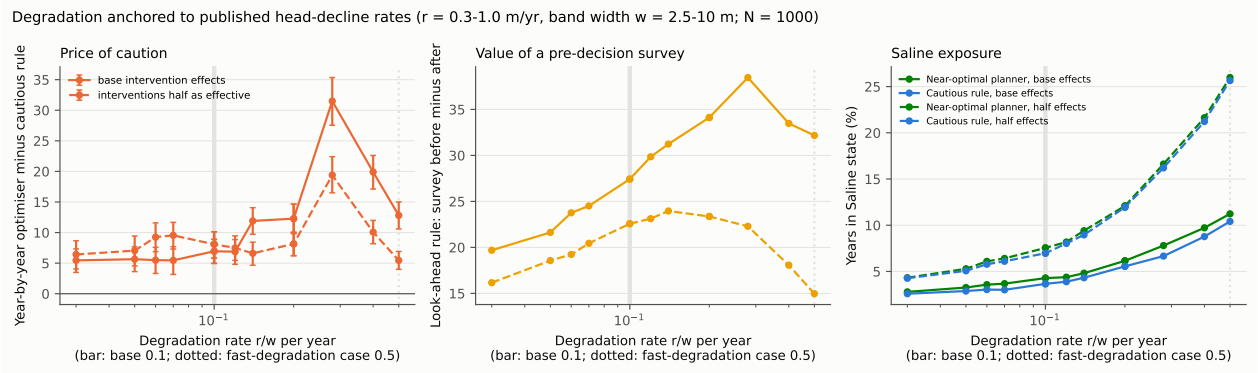


Figure S9: Price of caution, value of a pre-decision survey and saline exposure across published head-decline rates (solid: base intervention effects; dashed: half as effective).

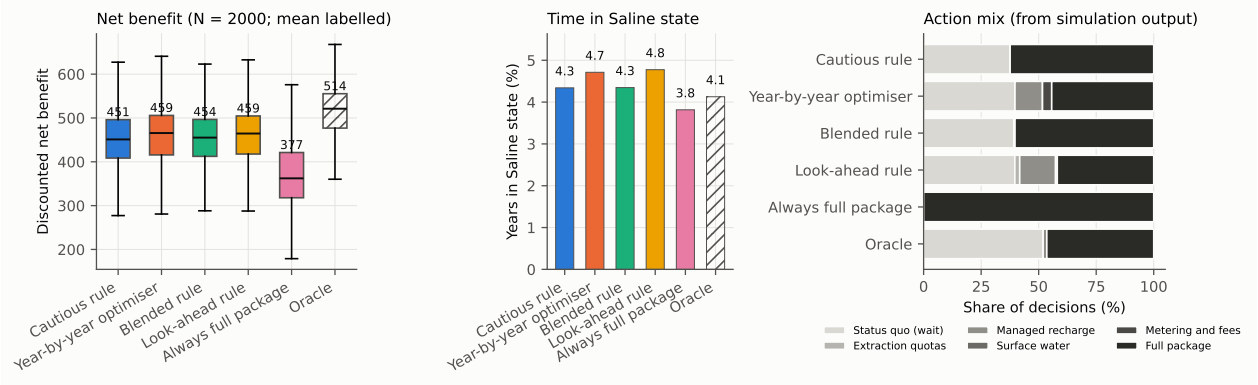


Figure S10: Base case (2000 futures): net benefit, time in the Saline state and action mix by rule.

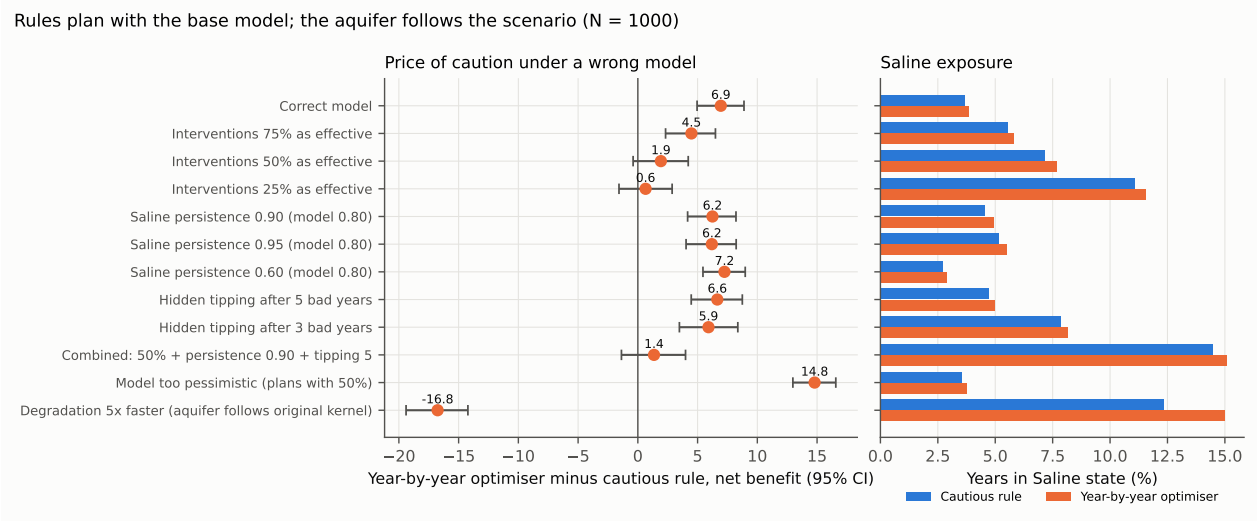


Figure S11: Price of caution (year-by-year optimiser minus cautious rule) and saline exposure when the planning model is wrong (1000 futures per scenario).

References

Jan L. Gunnink, Hung Van Pham, Gualbert H. P. Oude Essink, and Marc F. P. Bierkens. 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. doi: 10.5194/essd-13-3297-2021.

Michael L. Littman, Anthony R. Cassandra, and Leslie Pack Kaelbling. Learning policies for partially observable environments: Scaling up. In *Machine Learning Proceedings 1995*, pages 362–370, 1995. doi: 10.1016/B978-1-55860-377-6.50052-9.

Joelle Pineau, Geoff Gordon, and Sebastian Thrun. Point-based value iteration: An anytime algorithm for POMDPs. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, pages 1025–1032, 2003.