



Robust Decision Making and Info-Gap Decision Theory for water resource system planning



Evgenii S. Matrosov, Ashley M. Woods, Julien J. Harou*

Department of Civil, Environmental and Geomatic Engineering, University College London, Chadwick Building, Gower Street, London WC1E 6BT, UK

ARTICLE INFO

Article history:

Received 5 August 2011
 Received in revised form 5 March 2013
 Accepted 9 March 2013
 Available online 29 March 2013
 This manuscript was handled by
 Konstantine P. Georgakakos, Editor-in-Chief,
 with the assistance of Aris P. Georgakakos,
 Associate Editor

Keywords:

Water resources planning
 Robust Decision Making (RDM)
 Info-Gap Decision Theory (IGDT)
 Uncertainty
 Infrastructure planning

SUMMARY

Stationarity assumptions of linked human–water systems are frequently invalid given the difficult-to-predict changes affecting such systems. In this case water planning occurs under conditions of deep or severe uncertainty, where the statistical distributions of future conditions and events are poorly known. In such situations predictive system simulation models are typically run under different scenarios to evaluate the performance of future plans under different conditions. Given that there are many possible plans and many possible futures, which simulations will lead to the best designs? Robust Decision Making (RDM) and Info-Gap Decision Theory (IGDT) provide a structured approach to planning complex systems under such uncertainty. Both RDM and IGDT make repeated use of trusted simulation models to evaluate different plans under different future conditions. Both methods seek to identify robust rather than optimal decisions, where a robust decision works satisfactorily over a broad range of possible futures. IGDT efficiently charts system performance with robustness and opportuneness plots summarising system performance for different plans under the most dire and favourable sets of future conditions. RDM samples a wider range of dire, benign and opportune futures and offers a holistic assessment of the performance of different options. RDM also identifies through ‘scenario discovery’ which combinations of uncertain future stresses lead to system vulnerabilities. In our study we apply both frameworks to a water resource system planning problem: London’s water supply system expansion in the Thames basin, UK. The methods help identify which out of 20 proposed water supply infrastructure portfolios is the most robust given severely uncertain future hydrological inflows, water demands and energy prices. Multiple criteria of system performance are considered: service reliability, storage susceptibility, capital and operating cost, energy use and environmental flows. Initially the two decision frameworks lead to different recommendations. We show the methods are complementary and can be beneficially used together to better understand results and reveal how the particulars of each method can skew results towards particular future plans.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Water resource systems are sensitive to climate and population changes, making supply infrastructure planning difficult. Planning models grapple with the inherent uncertainty of future conditions when the statistical distributions of future conditions are unknown or not trusted. Under such ‘Knightian’ uncertainty (Knight, 1921) uncertainty is unquantifiable and the most likely realisation of the future is unknown. In such situations methods that rely on traditional Bayesian decision analysis to characterise uncertainty using probability theory may not be appropriate (Groves and Lempert, 2007). In such situations of ‘deep’ or ‘severe’ uncertainty, recent research has argued it is more appropriate to strive for robustness (Ben-Haim, 2001; Dessai and Hulme, 2007; Lempert

et al., 2006; Lempert and Collins, 2007) rather than optimality. A ‘robust’ system performs satisfactorily, or satisfices (Simon, 1959) performance criteria, over a wide range of uncertain futures rather than performing optimally over the historical period or a few scenarios.

Robust Decision Making (RDM) (Lempert and Popper, 2003) and Info-Gap Decision Theory (IGDT) (Ben-Haim, 2001) are two decision making frameworks that seek robustness. Both use trusted simulation models to consider a wide spectrum of plausible futures each with different input parameters to represent uncertainty.

Both approaches have been applied to water management. Groves and Lempert (2007) use RDM to identify vulnerabilities of the California Department of Water Resources’ California Water Plan (CWP). Lempert and Groves (2010) apply RDM to identify climate change vulnerabilities of the Inland Empire Utilities Agency’s 2005 Integrated Water Resource Plan and to develop a more robust plan including adaptive strategies. Hipel and Ben-Haim (1999) use

* Corresponding author. Tel.: +44 20 7679 0536.

E-mail address: j.harou@ucl.ac.uk (J.J. Harou).

IGDT to represent different sources of hydrological uncertainty. IGDT was also used within a water resources – timber production management problem in Australia to test the robustness of planting strategies to uncertainties in wildfire return periods and the knock-on impacts on municipal water availability (McCarthy and Lindenmayer, 2007). The sensitivity of UK flood management decisions to uncertainties in flood inundation models was investigated with IGDT (Hine and Hall, 2010). To our knowledge this paper constitutes the first IGDT application to water supply planning.

In a recent comparison of the two approaches, Hall et al. (2011) highlight the strengths of IGDT and RDM for robust system planning. In that application they find both tools come to similar conclusions to a climate change problem but provide different insights about the performance and vulnerabilities of the analysed strategies. They describe IGDT as a tool comparing candidate strategies' performance under a wide range of plausible futures (robustness) and their potential for rewards (opportuneness) under favourable future conditions. RDM identifies under which combination of future conditions a particular strategy becomes vulnerable to failure through 'scenario discovery' (Bryant and Lempert, 2010). Identifying different failure conditions provides scenarios to test plans and devise new strategies. IGDT by contrast, provides the facility to simultaneously compare the robustness and opportuneness of multiple strategies but does not quantify their vulnerabilities.

Since both methods are broadly similar (Hall et al., 2011), most analysts will use either one. Our study shows this can lead to the analysis missing important information that could have been uncovered if both frameworks had been used. Joint application reveals complementary information that would not be available if only one method were implemented.

This paper provides the first joint application of RDM and IGDT to a water resource system planning problem: identifying robust water supply portfolios for London and the Thames Basin for estimated future conditions in 2035. We critically assess how the two methods help formulate robust plans and how joint use increases understanding and insights.

The paper first reviews RDM and IGDT theory (Section 2). Section 3 describes the Thames basin and Section 4 introduces the system's capacity expansion problem and future uncertainties. Sections 5 and 6 describe the RDM and IGDT applications and Section 7 discusses the lessons learned from each application, and the limitations and the benefits of joint implementation. Section 8 concludes the paper.

2. Background on selected planning methods

2.1. Robust Decision Making (RDM)

RDM is a planning framework designed to help decision makers formulate robust plans for the future under conditions of Knightian, or 'deep' uncertainty (Lempert and Collins, 2007). RDM represents such uncertainty by considering system performance under a wide range of futures. RDM favours the concept of robustness over optimality and assumes that a strategy that is able to satisfy minimum performance criteria over a wide range of plausible futures is preferable.

RDM helps decision makers develop new strategies that are more robust than those initially considered (Hall et al., 2011). In RDM, the initial preferred plan is evaluated under a wide spectrum of plausible futures. RDM characterises the vulnerabilities of the initial strategy using a process known as *scenario discovery* (Bryant and Lempert, 2010; Groves and Lempert, 2007; Lempert et al., 2006) which rigorously identifies sets of future conditions, or 'scenarios', where the system under the preferred plan would be under most stress. Decision makers use this information to improve the

candidate strategy generating new strategies that hedge against those conditions which most frequently cause system failures. The new strategies can be resubmitted into the RDM framework and the process repeated iteratively until a suitably robust strategy is found.

An RDM analysis begins with the selection of one or more candidate strategies. Sometimes current policy can be selected as the candidate strategy or it can be one or more proposed future plans. Other times it can be chosen through a traditional utility or regret analysis (Lempert and Groves, 2010). The second step characterises the vulnerabilities of the candidate strategy by identifying under which combination of uncertain conditions it fails to meet performance criteria. To identify the vulnerabilities a trusted simulation model calculates system performance criteria for different combinations of input conditions representing a wide spectrum of plausible future states. Each run is evaluated as a success if it is able to satisfy minimum performance criteria or a failure if it is not. Performance criteria are either absolute thresholds where a strategy fails if its performance crosses a certain threshold such as a cost limit, or they are relative, where the performance of the candidate strategy is compared to the performance of an ideal strategy in the same state of the world; for example using regret (deviation from optimality) (Savage, 1954). Analysts then use cluster-finding algorithms to identify under which combinations of future conditions a particular strategy becomes vulnerable (failure clusters). Clusters are regions of the solution space bounded by one or more dimensions that have a relatively high density of failure points compared to the density in the whole of the solution space.

In the final step planners propose improvements to address the vulnerabilities uncovered in the previous step developing alternative strategies or discarding the strategy altogether if its vulnerabilities are unacceptable. The process returns to the second step to analyse an improved strategy or to the first step for a new strategy, and the process is repeated until planners agree on a 'robust' strategy. After this process, planners produce trade-off summaries which are used to compare alternative strategies.

2.2. Info-Gap Decision Theory (IGDT)

Info-Gap Decision Theory (IGDT) (Ben-Haim, 2001) is a non-probabilistic technique seeking to maximise robustness of a decision given minimum performance requirements. IGDT characterises the uncertainty of future system stresses as a group of nested sets defined by the parameter \tilde{u} , a best estimate of a future parameter, u . The deviation between u and \tilde{u} is scaled by h , which represents an increment or 'horizon' (Info-Gap parlance) of uncertainty. This can be parameterised such that, $h : h \geq 0$. The Info-Gap uncertainty model is constructed as a nested set based on \tilde{u} and h such that $U(h, \tilde{u})$. Fig. 1 shows a schematic of an Info-Gap uncertainty model showing the scaling of each interval (horizon) of uncertainty from a best estimate.

In the context of water resource planning, a number of management options q_i are available. Simulating each option quantifies minimum performance Π , or a performance windfall (reward) R , for different horizons of uncertainty α , such that $\Pi(q_i, h)$ or $R(q_i, h)$.

For management option q_i there will be a range of performances or rewards for each horizon of uncertainty specified by $U(h, \tilde{u})$. There will be minimum and maximum levels of performance for each h , which are referred to as *robustness* and *opportuneness* respectively in IGDT (Ben-Haim, 2001).

IGDT identifies management options that perform acceptably well under a wide range of conditions; seeking robustness rather than optimality in a process known as *robust-satisficing* (Ben-Haim, 2010). Acceptably well is defined as some minimum level of system performance that must be achieved, Π_c .

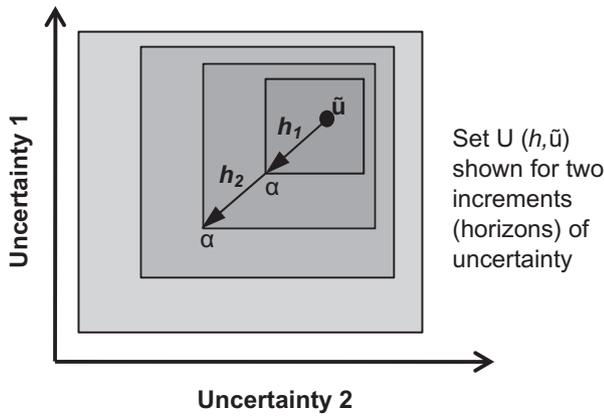


Fig. 1. Schematic of an Info-Gap uncertainty model showing the scaling (h) of each interval (horizon) of uncertainty (α) from a best estimate (\bar{u}). Note that the uncertainty model can be asymmetric to better represent the uncertainty surrounding the best estimate.

The resultant *robustness* \hat{h} is a measure of the maximum amount of uncertainty α (the maximum horizon of uncertainty) that can be tolerated whilst still ensuring a level of performance Π that is greater than a critical level Π_c .

$$\hat{h}(q_i, \Pi) = \max \left\{ h : \min_{u \in U(\alpha, \bar{u})} \Pi(q_i, u) \geq \Pi_c \right\} \quad (1)$$

As the requirement for performance increases; robustness decreases such that robustness is zero, $\hat{h}(q_i, u) = 0$, when the maximum possible level of performance $\Pi(q_i, u)$ is required. The robustness function (Eq. (1)) charts the performance of a management option as incrementally increasing uncertainty produces system inputs that result in lower system performance $\Pi(q_i, \bar{u}) \leq \Pi(q_i, u)$. Uncertain inputs may also result in a level of performance Π for a given management option that produces a

performance windfall Π_w , which gives rise to the IGDT ‘opportunity function’:

$$\hat{\beta}(q_i, \Pi) = \min \left\{ \alpha : \max_{u \in U(\alpha, \bar{u})} R(q_i, u) \geq \Pi_w \right\} \quad (2)$$

Info-Gap assumes users can identify a best estimate of the unknown parameter from which to start the uncertainty analysis; in our case the central estimate was used because we did not have evidence to suggest we should not use it. In cases of ‘severe’ uncertainty this estimate may in fact be a poor estimate of the real value. If so, because Info-Gap only considers a subset of the full uncertainty space around the best estimate, it could potentially ignore areas in which the real parameter value could fall. Info-Gap seeks local robustness around the choice of the best estimate instead of global robustness in the whole of the uncertainty space.

3. The Thames basin

The Thames river basin covers 16,133 km² and has 13 million inhabitants. Mean annual precipitation is 690 mm, with mean annual long-term effective precipitation at 280 mm (Merrett, 2007). Two private water companies manage the majority of municipal water supply: Thames Water Utilities Limited (TWUL) and Veolia Three Valleys Water (VTVW). The River Thames provides over 50% of TWUL’s supply and satisfies 70% of London’s demand (Jones, 1983).

Raw water from the Thames is stored in ten raised reservoirs between Slough and Teddington Weir. In the Lee valley thirteen storage reservoirs are fed from abstractions on the River Lee. The two reservoir groups are connected via a bulk transfer from the Thames Reservoirs along the Thames-Lee Tunnel (Halrow Ltd., 2010). The Thames basin has two existing surface water – groundwater conjunctive-use schemes in addition to normal groundwater abstraction. The North London Artificial Recharge Scheme (NLARS) recharges excess treated water to an aquifer in the Lee valley for later use during dry periods and the West Berkshire Ground Water Scheme (WBGW) augments supply during droughts. TWUL began

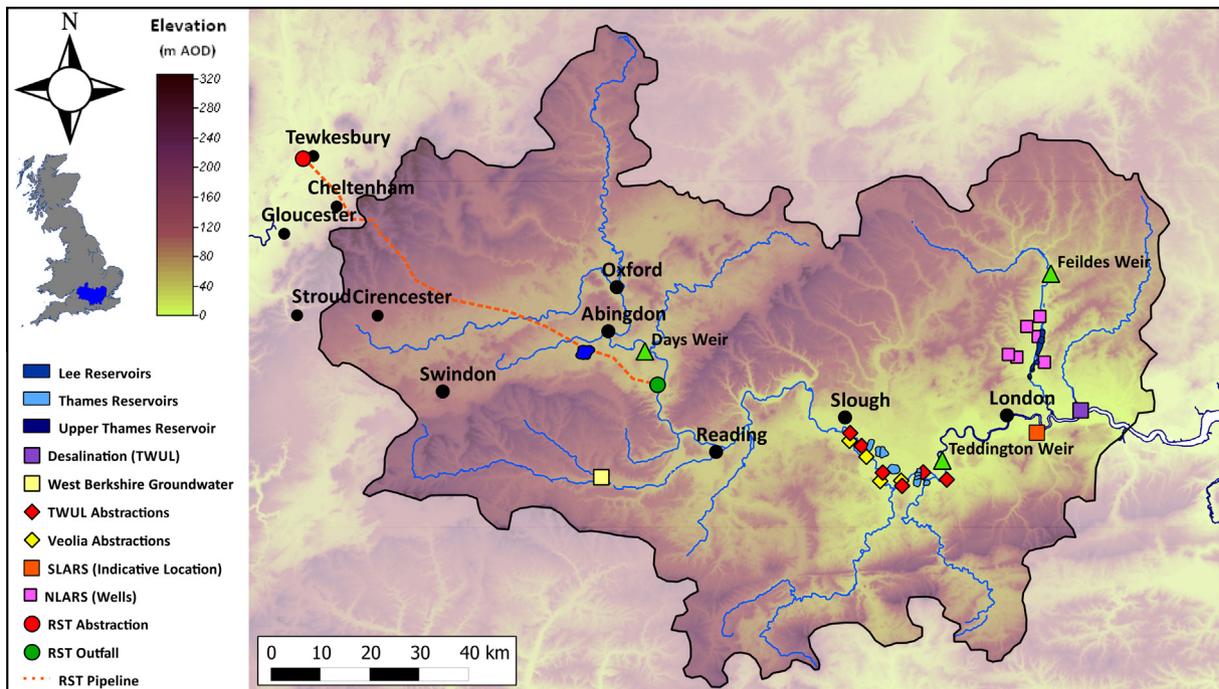


Fig. 2. The Thames basin showing major urban centres, and current and possible future water supply infrastructure. Adapted from Matrosov et al. (2011).

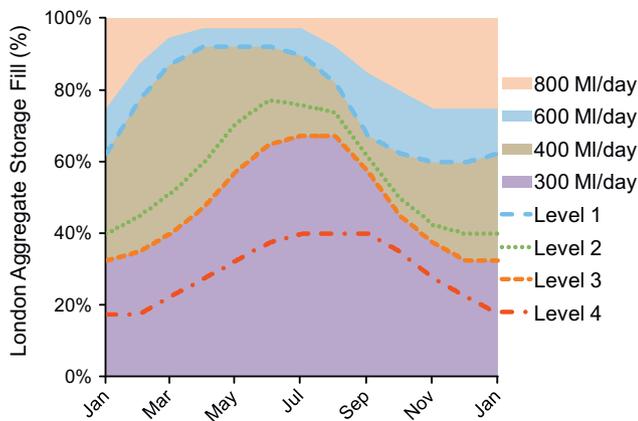


Fig. 3. The lower Thames control diagram showing the four levels of minimum environmental flows at Teddington and demand restriction thresholds as a function of the day in the year and storage in the LAS. Adapted from Matrosov et al. (2011).

operation of a desalination plant along a tidal stretch of the Thames in 2009. The Thames basin and water resource system is displayed in Fig. 2.

We use a Thames Basin water resource system model built with IRAS-2010 (Matrosov et al., 2011), an open-source rule-based water resource system simulator. The model includes 28 nodes (junctions, reservoirs, aquifers, treatment plants) and 33 links (rivers, canals, transfers, etc.) and runs on a weekly time-step. IRAS-2010 tracks flows and storages of every node and link in the system and registers system performance at every time-step. At the end of the simulation performance is summarised using performance metrics. Time series of naturalised gauged flows from the National River Flow Archive (NRFA) provide inflows into the system. One aggregated surface storage node, the London Aggregate Storage (LAS), represents the combined surface storage of all the reservoirs in the basin. Abstraction to surface storage on the Thames is controlled downstream by environmental minimum flows (EMF) at Teddington Weir. The EMF varies between 800 and 300 MI/day and is controlled by the Lower Thames Control Diagram (LTC) (Fig. 3). The LTC also determines when storage levels in LAS trigger the four levels of water use restrictions. For details on the model refer to Matrosov et al. (2011).

4. Problem formulation

The Thames basin has experienced six major droughts in the last 90 years (Marsh et al., 2007). London's population is expected to increase by 16% by 2028 (Office for National Statistics, 2008). Population growth and the likelihood that climate change will reduce summer precipitation in south-east England (Murphy et al., 2009) increases the threat of water scarcity. Water supply infrastructure expansions will form part of the solution to reduce the

region's vulnerability to drought. In this case study we apply the RDM and IGDT planning frameworks to an infrastructure planning problem for 2035.

4.1. Infrastructure portfolios (IPs)

Thames Water (2010) describes 37 plausible future water supply options including multiple variants of the same option. Of these we selected seven options (Table 1), leading to 19 possible infrastructure portfolios (IPs) (unique combinations of supply options). Selection criteria included the public availability of data and the expected impact of each option on system performance. Adding the current arrangement led to the 20 possible IPs, or 'options', considered in this case study (Table 2).

We investigate the Upper Thames Reservoir (UTR) with capacities of 75, 100 or 150 million m^3 ($M m^3$) (UTR75, UTR100, UTR150), the River Severn Transfer (RST) and the South London Artificial Recharge Scheme (SLARS). We also investigate the operation of the desalination plant at 80 MI/day (DESAL80) and 140 MI/day (DESAL140).

The proposed UTR is filled from River Thames abstractions during periods of high flows. The UTR would provide continual supply to the Swindon and Oxfordshire Water Resource Zone (SWOX WRZ), with additional releases into the River Thames during drought conditions for re-abstraction downstream for London. The River Severn Transfer (RST) would import water from the River Severn to the west of the Thames basin (Fig. 2). When activated during drought conditions the RST would import an amount equivalent to the release of UTR150 (Thames Water, 2010).

SLARS pumps from the confined chalk aquifer in South London (Jones et al., 2005) and functions analogously to NLARS. For details on the function and release capacities of the UTR, RST and SLARS please refer to Thames Water (2010).

All infrastructure additions incur capital costs except for the desalination plant which is already built to operate to a maximum capacity of 150 MI/day. Variable and fixed operating costs are considered for each proposed option.

Variable operating costs are calculated by multiplying flow from supply nodes by the unit energy consumption, and unit energy cost. Energy consumption results from desalination, pumping groundwater, to overcome elevation gain and friction losses in the RST, and pumping from the River Thames to the UTR and from the Thames Estuary to the desalination plant. We assume a pumping efficiency of 75% (OFWAT, 2010) and a consumption of 2 kW h/ m^3 for desalination of brackish water (Raluy et al., 2005a, 2005b; Stokes and Horvath, 2006) (Table 3).

4.2. Uncertainties

Uncertainties considered in this study include natural hydrological variation, climate change perturbation of natural hydrology,

Table 1
The seven possible future options (see in Fig. 2).

ID	Name	Description	Capacity
UTR75	Upper Thames Reservoir	High capacity reservoir to supply Swindon and Oxford (SWOX) and London. Daily maximum supply set to 157.5MI/d ^a .	75 M m^3
UTR100	Upper Thames Reservoir	High capacity reservoir to supply Swindon and Oxford (SWOX) and London. Daily maximum supply set to 202MI/d ^a .	100 M m^3
UTR150	Upper Thames Reservoir	High capacity reservoir to supply Swindon and Oxford (SWOX) and London. Daily maximum supply set to 315MI/d ^a .	150 M m^3
RST	River Severn Transfer	High capacity bulk transfer of freshwater from the River Severn to the River Thames ^a .	315 MI/d
DESAL80	London Desalination	Reverse osmosis desalination plant	80 MI/d
DESAL140	London Desalination	Reverse osmosis desalination plant.	140 MI/d
SLARS	South London Artificial Recharge Scheme	Artificial groundwater storage and recovery scheme (South London) ^a (Fig. 2).	22 MI/d

^a (Thames Water, 2010).

Table 2

The 20 water supply 'infrastructure portfolios' (IPs), alternatively called 'options' or 'plans' evaluated in the case study.

Infrastructure portfolio	Future supply options
1	DESAL80
2	DESAL80, SLARS
3	DESAL140
4	DESAL140, SLARS
5	UTR75, DESAL80
6	UTR75, DESAL80, SLARS
7	UTR75, DESAL140
8	UTR75, DESAL140, SLARS
9	UTR100, DESAL80
10	UTR100, DESAL80, SLARS
11	UTR100, DESAL140
12	UTR100, DESAL140, SLARS
13	UTR150, DESAL80
14	UTR150, DESAL80, SLARS
15	UTR150, DESAL140
16	UTR150, DESAL140, SLARS
17	RST, DESAL80
18	RST, DESAL80, SLARS
19	RST, DESAL140
20	RST, DESAL140, SLARS

projected London water demand and future energy prices up to a 2035 planning horizon.

4.2.1. Natural hydrological variability

Significant flow variability has been observed on the River Thames since records began in the 1880s, including several periods of prolonged drought (Cole and Marsh, 2006; Marsh et al., 2007). To consider natural hydrological variability each IRAS-2010 Thames simulation is run over the 85-year historical period (1920–2005) assuming any one of these hydrological years could occur in 2035.

4.2.2. Climate change perturbation

Climate change uncertainty is represented using monthly climate change perturbation factors that are multiplied by historical river flow time series to estimate future flows. Perturbation factors based on UKCP09 probabilistic climate change projections (Murphy et al., 2009) were used to create sets of twelve monthly flow factors (Christierson et al., 2012) valid from 2020–2035. Each set of monthly flow factors represents one hydrological future.

The UKCP09 climate projections were developed by the UK Met office and are based on "perturbed physics" and multimodel ensembles produced by a variant of the HadCM3 climate model (Johns et al., 2003) and the results of 13 other Global Climate Models (GCMs). Murphy et al. (2007) provide a description of multi-

model and perturbed physics ensembles. The projections are based on the B1, A1B, A1F1 emissions found in the IPCC's (Intergovernmental Panel on Climate Change) Special Report on Emissions Scenarios (SRES) (Nakicenovic and Swart, 2000). Please consult Murphy et al. (2009) for a detailed description of the UKCP09 projections.

The Thames catchment was simulated using the Probability Distributed Model (PDM) (Moore, 1985) and the Catchmod (Wilby et al., 1994) rainfall run-off models using Latin Hypercube Sampling of the UKCP09 projections for the A1B scenario. A large number of modelled river flow times-series was generated representing the future and historical time periods to take into account uncertainty in hydrological parameters using the Generalised Likelihood Uncertainty Estimation (GLUE) methodology (Beven, 2006; Beven and Binley, 1992). Models that closely reproduced the historical time series were used to generate a distribution of monthly flow factors. Please see Christierson et al. (2012) for a detailed description of the flow perturbation factors.

An ensemble of 100 equiprobable flow factor sets was used. The flow factor sets were classified by their percentage perturbation effect on winter (November–April) and summer (March–October) historical flows.

For the IGDT application, the central estimate (median) flow factor set is used as the best estimate. The IGDT uncertainty model creates sets of monthly perturbation factors that diverge from the median flow factor set at structured intervals defined by the IGDT uncertainty model.

4.2.3. London water demand

London mean annual water demand estimates for the time horizon were obtained from TWUL's stochastic water demand forecasting tool (Thames Water, 2010). The demand projections for 2035 were fitted to a gamma distribution with shape 311 and scale 6.9.

For the RDM application, nine levels of London water demand were obtained from deciles of the gamma distribution. Including the demand values at the 0.01 and 0.99 percentile of the gamma distribution raised the total demand values considered to 11.

For IGDT an uncertainty model is built around this best estimate of demand $Y|x = 0.5$ (2144.5 Ml/day) with the increments of uncertainty derived conditionally from the underlying gamma distribution such that: $x = 0.525, 0.550, 0.575, \dots, 0.975$. The same sample intervals were used to sample the left-side of the gamma distribution.

4.2.4. Energy prices

A uniform distribution of price estimates (£0.09–£0.22/kWh) was assumed based on current energy price (£0.09) (Eurostat,

Table 3

Costs and release rates of the management options.

Management option	Release rate (ML/day)	Capital costs (M£)	Energy use (kW h/m ³)	Fixed operating costs (k£/yr)
UTR75	133.5 to Thames 24 to SWOX ^a	536.9 ^a	0.11 ^b	360 ^a
UTR100	178 to Thames 24 to SWOX ^a	725.5 ^a	0.11 ^b	370 ^a
UTR150	267 to Thames 48 to SWOX ^a	821.6 ^a	0.11 ^b	380 ^a
RST	267 to Thames 48 to SWOX ^a	579.6 ^b	1.08 ^{b,c}	7,000 ^c
SLARS	22 ^a	2.3 ^b	0.23 ^b	53 ^a
DESAL80	80	–	2.15 ^{b,d}	830 ^a
DESAL140	140	–	2.15 ^{b,d}	930 ^a

^a (Thames Water, 2010).

^b Potential energy equation with 75% pump efficiency.

^c (OFWAT, 2010).

^d (Raluy et al., 2005a, 2005b).

Table 4
Uncertain system parameters for strategic water resource forecasting in the Thames Basin for the 2020–2035 planning horizon.

Parameter	RDM	IGDT
Hydrological variability	Historical flows ^a	Same as for RDM
Climate change perturbation	100 monthly flow perturbation sets valid for 2020–2035 ^a (Monte-Carlo sampling)	A single central estimate climate change flow factor set ^b with 40 intervals of uncertainty either side of the central estimate
Water demand	11 water demand levels obtained using LHC sampling of deciles of a gamma distribution + 2 extreme values from distribution forecasts ^b	A single central estimate of demand with 40 conditional intervals of uncertainty scaled by the gamma distribution either side of the central estimate ^b
Energy prices	14 energy prices ^c obtained using LHC sampling of 13-quantiles of a uniform distribution	A single central estimate with 40 horizons of uncertainty either side of the central estimate ^b

^a National River Flow Archive (NRFA).

^b (Thames Water, 2010).

^c (Eurostat, 2011).

2011) assuming prices will increase (DECC, 2010). For the RDM application the uniform distribution was divided into 13 equiprobable intervals whose boundaries were used to produce 14 energy costs. For IGDT a best estimate of 13 p/kWh was used to construct the uncertainty model bounded by the upper and lower estimates described above.

Uncertainties for both RDM and IGDT are summarised in Table 4.

4.3. Performance criteria

IGDT and RDM require performance criteria to classify each simulation as a failure or success. Success or failure is based on five satisficing metrics: reliability of water supply service, reservoir storage susceptibility, environmental performance, energy consumption and total costs (capital and operating).

The service reliability criterion is based on the frequency of water use restrictions imposed in the Thames basin over the full simulation. The restrictions considered are levels 2 and 3 from the LTCD (Fig. 3) which correspond to sprinkler and non-essential use bans respectively. Average service reliability is calculated by assessing the number of weeks ($W_{fail,i}$) each restriction level (i) was imposed in the simulation over the whole 85-year simulation (W_{sim}):

$$R_i = \left(1 - \frac{W_{fail,i}}{W_{sim}}\right) \quad (3)$$

Thresholds for L2 and L3 failures were calculated from the current maximum frequency of supply restrictions TWUL will impose; 1 in 10 years for sprinkler bans (level 2–L2) and 1 in 20 years for non-essential use bans (level 3–L3). For L2 restrictions we use the 9 month (40 week) hosepipe ban imposed by TWUL in 2007 (BBC, 2007) as a representative L2 restriction duration. An allowable frequency of L2 restrictions of 1 in 10 years translates to 332 weeks of allowed L2 restrictions over the 85-year simulation. L3 restrictions can be imposed no more frequently than 1 in 20 years (half that of level 2) and we use a representative failure duration of 4.5 months in line with historical durations of non-essential use bans in south-east England (WaterBriefing, 2006). Following Eq. (3), a reliability of 0.925 for restriction L2 (sprinkler ban) is used for this study and a reliability of 0.973 is used for restrictions at L3.

A storage susceptibility metric is defined as the lowest storage level reached by LAS. Its emergency storage of 22.5% of capacity (45 M m³) is the threshold below which failure occurs due to pressure-related distribution problems in the network (Cookson and Weston, 2008).

Environmental performance is calculated downstream of the last abstraction on the Thames at Teddington Weir. A failure occurs whenever the flow at this node drops below 800 MI/day, the minimum environmental flow during normal conditions (Fig. 3). The

duration and severity of these failures are quantified using the Shortage Index (SI) adapted from Fredrich (1975):

$$SI = \frac{100}{N_T} \sum_{t=1}^{N_T} \left(\frac{WS_t}{WD_t}\right)^2 \quad (4)$$

where WS_t is the flow shortage at week t , WD_t is the weekly flow minimum (800 MI/day \times 7 days) and N_T is the total weeks in the simulation. Higher SI values signal worse performance.

The energy consumption criterion is the total cumulative energy consumed by infrastructure during each simulation run. The total cost criterion combines capital and cumulative operating costs of each simulation. Capital costs, fixed operating costs and energy consumption figures for infrastructure options are provided in Table 3.

Because of a lack of explicit performance requirements for the environmental, energy and cost criteria we consider the worst 15% of performances to be failures.

In our RDM and IGDT implementations, if a simulation fails in at least one of the performance criteria outlined above, that particular simulation is considered a failure.

5. RDM implementation

We perform one iteration of the RDM framework. From among the twenty possible infrastructure portfolios (IPs) we select one as the candidate strategy. We then identify its vulnerability scenarios and determine which IP would perform better inside these scenarios to consider possible tradeoffs between IPs. Ways to address weaknesses in the IPs are considered.

We use the enumeration method to sample 20 infrastructure portfolios (Table 2), 101 climate scenarios (100 equiprobable hydrological realisations and the historical record), 11 levels of estimated London water demand and 14 values of energy costs described in Section 4.2 leading to 311,080 simulations.

5.1. Choosing a candidate strategy using regret analysis

For RDM one of the twenty IPs must be selected as the initial candidate strategy. We perform a regret-based ranking (Savage, 1954) of the proposed IPs to identify which performs best over the simulated set of plausible futures.

Regret, or deviation from an ideal performance of criteria c , is the difference in the performance P of the best performing strategy (s') and that of the strategy in question, s , for the same input parameters, j .

$$R_c(s,j) = |P_c(s',j) - P_c(s,j)| \quad (5)$$

For cost, energy and environmental performance; where low values are better, an absolute value is required. In addition to considering

regret for each criteria, we also consider multi-criteria regret, where the regrets are standardised over a 0 to 1 interval and weighted equally.

Fig. 4 shows box and whisker plots of normalised regret for the five performance criteria and the multi-criteria aggregate regret. Relative performance is assessed on three ordering statistics: the lower quartile (lower end of the box), median (red line inside the box) and the upper quartile (upper end of the box) value of regret for each IP. We use median regret to select the candidate IP.

Fig. 4f (Energy consumption regret), the easiest to interpret, shows that all IPs that include DESAL140 have median energy regret close to 1 and IPs with DESAL80 have near-zero regret. Desalination is twice as energy intensive as RST and eclipses the pumping energy requirements (Table 3).

IP 17, which includes the RST, has the lowest median regret (median regret is 0) because the RST perennially provides out-of-Thames-basin flow (unlike the UTR which frequently empties during severe droughts) and thus relies least on desalination.

Fig. 4b, c and e show regret for service reliability, storage susceptibility and environmental performance. IPs with more and larger infrastructure perform better with these criteria as they can

supply more water during droughts. IP 1, the baseline IP, has the highest median regret followed by IP 2 which includes SLARS. SLARS has little effect on median regret of IPs in the service and environmental performance criteria. Increasing desalination output to 140 Ml/day strongly reduces median regret in all IPs as does the UTR. Larger UTR capacities improve regret. In all three criteria the increase in regret from running desalination at 80 Ml/day rather than 140 Ml/day outweighs the amelioration in median regret that the UTR100 provides over the UTR75. In all three criteria, IP 20, which includes the RST, DESAL140 and SLARS, has the lowest median regret. IP 19 performs nearly as well followed by IPs 16 and 15 which replace the RST with the UTR150. IPs with less and smaller infrastructure (e.g. IPs 1 and 2) perform better in the cost criterion (Fig. 4d). SLARS results in a small increase in cost regret as does the higher desalination output due to high energy use. The RST incurs large capital and operating costs leading to the highest regret values. IP 20, the most costly portfolio in every modelled future, has regret values approaching 1.

Fig. 4a shows equally weighted multi-criteria regret. All options with DESAL140 have higher regret than those with DESAL80 while regret decreases with higher UTR volumes. IPs with RST and UTR

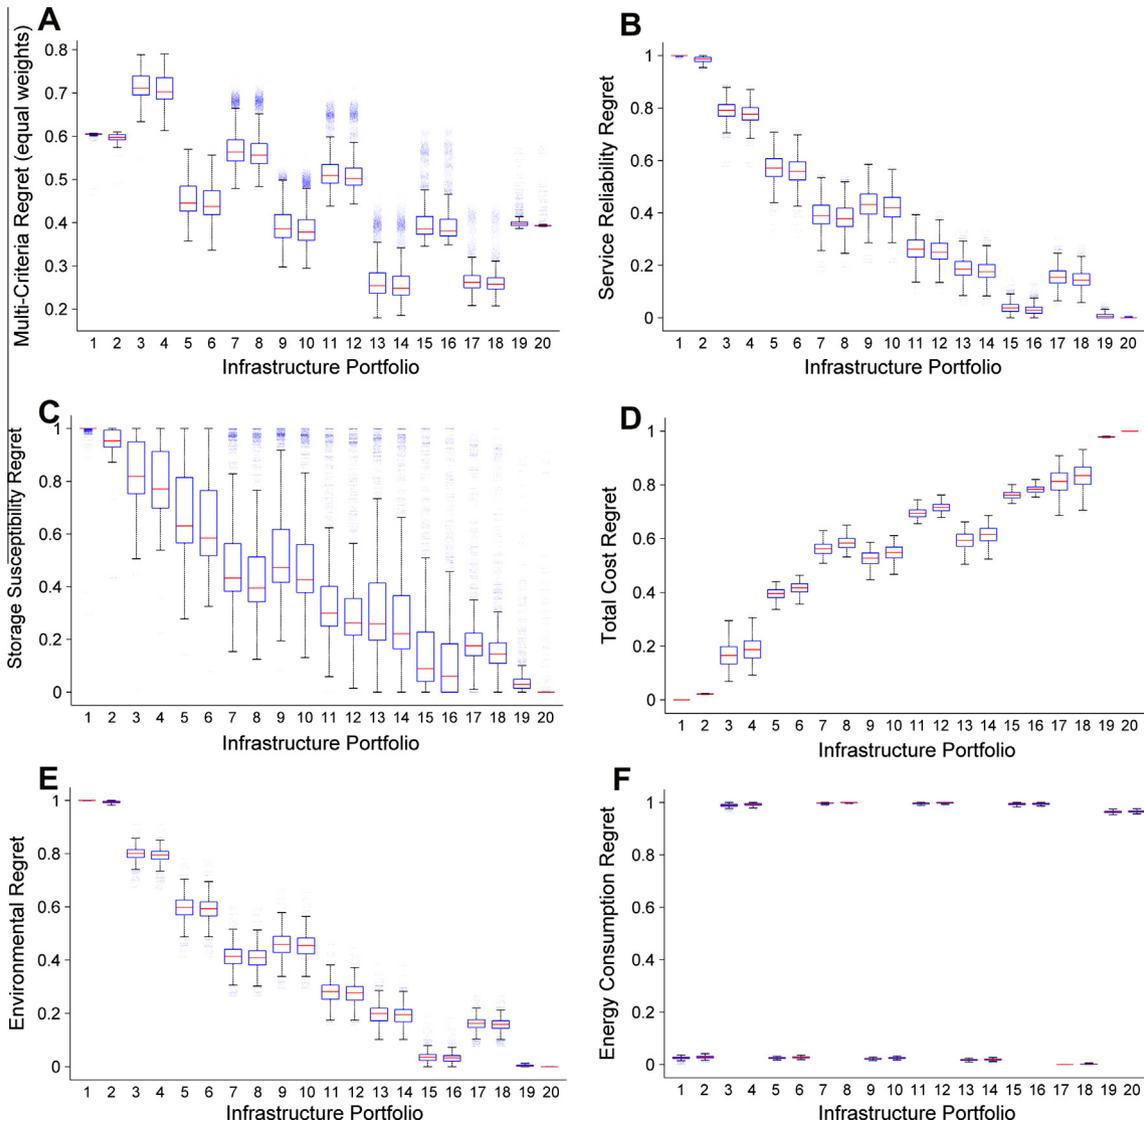


Fig. 4. (a–f) Box and whisker plots of regret in each of the five performance criteria and an aggregated regret calculated using equal weighting of all criteria. The median fractional regret is represented by the red line in each box. The top and bottom of the boxes are the upper and lower quartiles respectively while the whiskers represent the regret value within 1.5 times the interquartile range from the upper and lower quartiles. Points represent outliers.

Table 5
Upper quartile, median and lower quartiles of regret for IPs 13, 14, 17 and 18.

	IP 13	IP 14	IP 17	IP 18
Upper quartile	0.284	0.276	0.278	0.272
Median	0.254	0.248	0.262	0.258
Lower quartile	0.236	0.232	0.249	0.246

150 have similar regret values. IPs 13, 14, 17 and 18 have the lowest regret values. Table 5 summarises the lower quartile, median and upper quartile regret of the four best performing IPs.

In this case study we chose the option with the lowest median value of regret, IP 14, to take forward to the next RDM phase.

5.2. Characterising vulnerabilities of the candidate strategy

Having identified infrastructure portfolio 14 (UTR150, DESAL80, SLARS) as our candidate strategy, RDM's scenario discovery stage identifies under which combinations of uncertain parameters IP 14 is vulnerable to failure. Following Bryant and Lempert (2010) we use a modified version of the Patient Rule Induction Method (PRIM) (Friedman and Fisher, 1999) to characterise the vulnerabilities of IP 14. Vulnerabilities are defined as combinations of input conditions that produce high levels of failures.

PRIM is an interactive statistical cluster-finding algorithm that finds one or more low-dimensional boxes, or 'scenarios' in a hyper-dimensional space where the density of interesting points inside each box is higher than the space outside the box. PRIM seeks to maximise the density, coverage and interpretability of boxes. Density is defined as the total number of interesting points (in this case failure points) over the total points inside the box whilst coverage is the total number of interesting points inside the box compared to the total in the entire space. A trade-off or balance between coverage and density must be subjectively made by the analyst. Interpretability is how easily decision makers can interpret each box as a coherent 'scenario'. Interpretability is approximated by the number of dimensions that define each scenario.

We apply a modified version of PRIM using the "Scenario Discovery Toolkit"¹ in the publically available statistical computing environment R which implements the PRIM code with a variety of useful features and visualisations (Bryant and Lempert, 2010). The Scenario Discovery Toolkit produces density vs. coverage trade-off plots. Each point on the plot represents a unique scenario described by one or more dimensions, and each dimension is represented by one uncertain system parameter. The analyst selects the box with an appropriate density to coverage ratio. Once a box or 'failure scenario' has been identified, the failure points inside the box are removed from the solution space and a new plot can be generated in order to find additional boxes. This process is repeated until the decision maker is satisfied by the boxes identified.

Each member of the solution space for IP 14 is classified as success or failure using only four of the five performance criteria: service reliability, storage susceptibility, environmental performance and total energy consumed. Total cost was not used because all IP 14 runs have the same capital and fixed operating costs and only variable operating costs differ and are directly proportional to the energy consumption. This resulted in a failure density of 54% with the majority of failed simulations resulting from failure in the service reliability criterion.

Two scenarios were identified using PRIM covering 78% of all failure points with a density of 93% inside the two scenarios. In total the scenarios include 44% of all simulations. Fig. 5a and b shows

the density vs. coverage trade-off plots generated by PRIM (called 'peeling trajectories' in PRIM jargon). The colours of the points represent each dimension being constricted; a change in colour represents the constriction of a new uncertainty dimension. The circled points represent the boxes chosen as having the appropriate coverage and density to describe scenarios. Choosing an alternative point in Fig. 5a would result in PRIM generating an alternate subsequent trade-off curve (Fig. 5b) as choosing an alternative scenario would result in different points remaining in the solution set.

The dimensions and boundaries of our chosen vulnerability clusters are outlined in Table 6. Table 7 provides the density and coverage of each box.

Scenario 1 describes futures where summer river flows are more than 8% lower than historical (1920–2005) and London's mean annual demand is greater than 2098 Ml/day. In this scenario dry summers cause IP 14 to fail, even with moderate water demands. The high coverage (63%) of this scenario shows IP 14 is vulnerable to dry summers combined with moderate to high demand levels. Change in winter hydrology is not included in this scenario because changes in winter hydrology do not substantially affect this IP. In futures where summer flows are low, the system suffers from a lack of storage and cannot buffer the summer–winter disparity in flow, even with the desalination plant providing 80 Ml/day of additional supply during droughts.

Scenario 2 describes a scenario where London's mean annual demand ranges from low (in the interval between the first and second deciles) to high and includes dry summers and extremely dry winters (6% and 13% drier than historical respectively). This failure scenario shows that IP 14 performs poorly under dry climates no matter how low demand is.

The price of energy was not found to contribute to the vulnerability of IP 14. The density of both scenarios is high (94% and 97% for scenario 1 and scenario 2 respectively) suggesting that these scenarios are exceptionally dangerous for IP 14.

5.3. Initial findings and recalculating regret

Overall, moderate to high London demands combined with moderately dry summers cause IP 14 to fail. Low future demands when confronted with dry summers and extremely dry winters also lead to frequent system failure.

To check if any other IPs perform better under these vulnerability scenarios, regret is recalculated for all IPs for futures that fall within the range of conditions dictated by scenarios 1 and 2. We also examine the regret for all futures outside of the two scenarios. Regret plots for both scenarios are provided in the [Supplementary material](#). IP 18 has the lowest median regret inside the scenarios whilst IP 14 has the lowest median regret outside each scenario. Therefore, there is a trade-off between IP 14 and IP 18. Because of the similarity between the scenarios and this single trade-off, we combine scenarios 1 and 2 into a single vulnerability scenario and recalculate multi-criteria regret for all futures inside and outside of the conditions posed by scenarios 1 and 2 (Fig. 6a and b).

IPs 13, 14, 17 and 18 are still the best performing IPs for futures contained within the two vulnerability scenarios. Tables 8 and 9 summarise the lower quartile, median and upper quartile regret values of the four best performing IPs. For futures inside the two scenarios, regret is higher for all four IPs. However, the increase in the range of regret for values not considered outliers is considerably greater for IP 13 and 14. IP 18 has the lowest upper quartile and median value of regret in scenarios 1–2. Compared to the increase in range of regret for IP 14, the range in regret increases minimally (median regret remains the same) indicating that IP 18's relative performance remains stable inside and outside the vulnerability scenarios. For futures outside the vulnerability

¹ <http://cran.r-project.org/web/packages/sdtoolkit/index.html>.

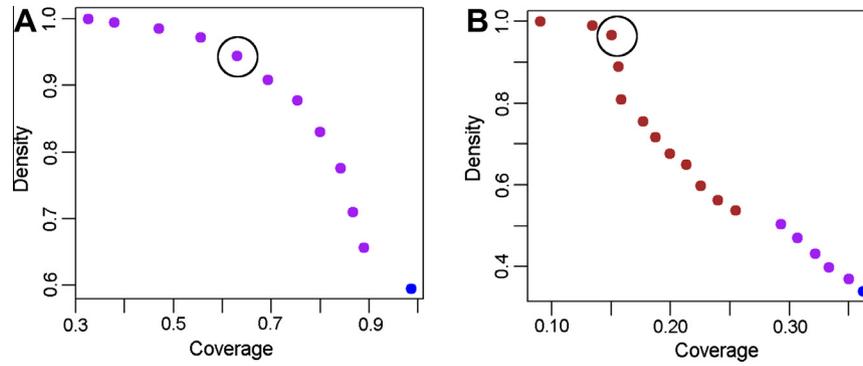


Fig. 5. Density vs. coverage trade-off plots produced by the scenario discovery toolkit. Colours represent different dimensions. Circled points were chosen as scenarios. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 6
Dimensions and boundaries of the two vulnerability scenarios identified using PRIM for infrastructure portfolio 14.

Scenario/dimension	1	2
Fractional change in summer hydrology	<0.92	<0.94
Fractional change in winter hydrology	–	<0.87
London demand (Ml/day)	>2098	>1933

Table 7
Density and coverage of the two vulnerability scenarios for infrastructure portfolio 14.

Scenario	Density	Coverage
1	0.94	0.63
2	0.97	0.15
Total	0.95	0.78

scenarios, IP 14 has the lowest regret values for each of the ordering statistics.

Fig. 7 shows the median regret of each of the twenty IPs. IP 18 has the lowest median regret inside scenarios 1–2 while IP 14 has the lowest median regret in all futures outside the scenarios. Relative to the regret of every other IP however, the difference in median regret between the IP 14 and 18 is small.

IP 14 and IP 18 are the same except IP 14 includes the UTR150 and IP 18 includes the RST. Both the UTR150 and the RST provide the same amount of water to the Thames at the beginning of droughts. The UTR however has maximum storage and is refilled from the Thames which cannot occur during droughts whilst the RST in our model continuously provides water from a neighbouring wetter region. The uncertainty in the reliability of the RST supply

(e.g. possible correlation of droughts in the Thames and Severn regions) was not considered in this study. We used a supply estimate for the transfer based on estimates of what it could provide under historical drought conditions (Thames Water, 2010).

The RST is costlier and more energy intensive, but has better performance in the storage susceptibility, service and environmental performance criteria. IP 14 however is vulnerable because it lacks the storage necessary to maintain service during extended droughts. RDM suggests that these two strategies are likely to be the most robust into the future. Section 7.1 discusses the merits and usefulness of our RDM implementation and its recommendations.

6. Info-Gap analysis

The formulation of a fractional-error Info-Gap uncertainty model for climate change perturbation, demand and energy cost is used to build robustness and opportuneness functions based on our system performance requirements.

Taking the convention u_1 = climate change perturbation, u_2 = water demand and u_3 = energy cost; κ_l and κ_r scale and, σ_r and σ_l bound the right (upper) and left (lower) intervals of the Info-Gap model respectively. The 3-dimensional Info-Gap Uncertainty model becomes:

$$U(h, \tilde{u}) = \{u : \max[\sigma_l, (1 - \kappa_l h)\tilde{u}_{ij}] \leq u_{ij} \leq \min[\sigma_r, (1 + \kappa_r h)\tilde{u}_{ij}]\} h \geq 0, \quad i = 1, 2, 3, \quad j = 1, 1 - 12 \quad (6)$$

where $\kappa_l = [0.005, 0.0125, 0.1]$ is the scaling factors for the left hand side of the Info-Gap model for u_i ; $\kappa_r = [0.005, 0.0125, 0.25]$ is the scaling factors for the right hand side of the Info-Gap model for

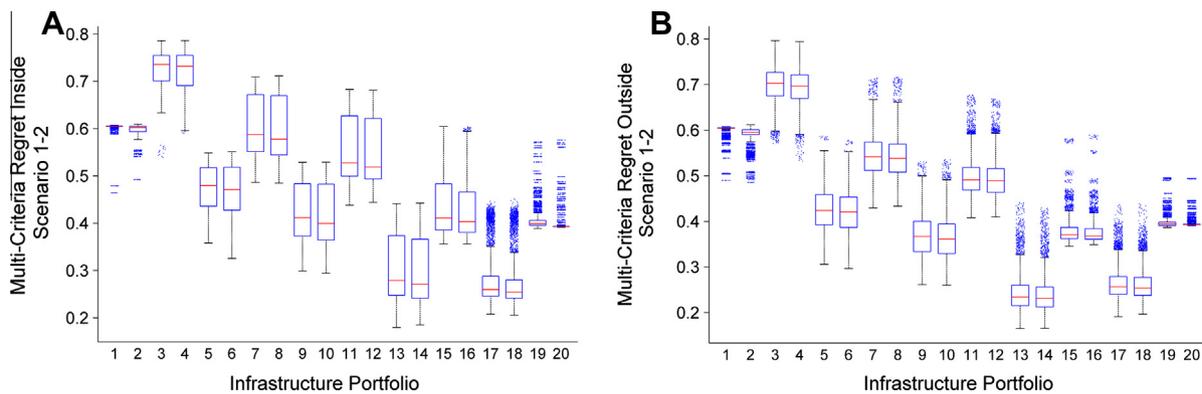


Fig. 6. (a) Multi-criteria regret of all modelled futures in both vulnerability scenarios and (b) regret of all modelled futures outside the two vulnerability scenarios.

Table 8

Upper quartile, median and lower quartiles of regret for IPs 13, 14 and 17 for futures comprised in vulnerability scenarios 1–2.

	IP 13	IP 14	IP 17	IP 18
Upper quartile	0.374	0.366	0.288	0.280
Median	0.279	0.271	0.260	0.254
Lower quartile	0.248	0.241	0.246	0.241

Table 9

Upper quartile, median and lower quartiles of regret for IPs 13, 14, and 17 for futures outside vulnerability scenarios 1–2.

	IP 13	IP 14	IP 17	IP 18
Upper quartile	0.260	0.256	0.279	0.277
Median	0.234	0.231	0.256	0.254
Lower quartile	0.215	0.213	0.240	0.238

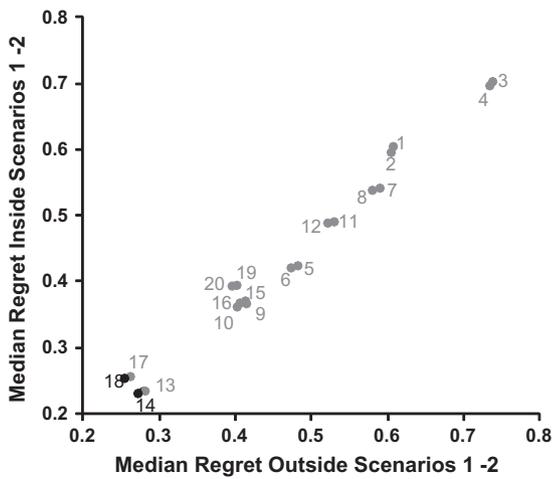


Fig. 7. Median multi criteria regret of each of the twenty policies inside and outside scenarios 1–2. IP 18 has the lowest median regret inside the vulnerability scenarios while IP 14 has the lowest regret outside the scenarios.

u_i ; $\sigma_l = [1.2, 1914.81, 9]$ is the lower boundaries of the Info-Gap model for u_i ; and $\sigma_r = [0.8, 2391.88, 22]$ is the upper boundaries of the Info-Gap model for u_i .

Best estimates of 2144.5 Ml/day and 13 p/kW h were used for u_2 and u_3 respectively. The best estimate for u_1 is defined by the median flow factor set from the 100 sets generated from UKCP09 climate models.

The use of scaling factors for the best estimate of the flow factors results in a compression or extension of the seasonal range of flows that is prescribed by the best estimate. When the scaling factor is less than 1 as in the case of the robustness function the seasonal range is compressed and it is extended when the scaling factor is greater than one as is the case of the opportuness function.

6.1. Robustness function

The robustness function (Eq. (7)) from the 3D Info-Gap uncertainty model (Eq. (6)) uses the critical system performance criteria $\Pi_c(Res_o, Cost_t, Energy_p, Env_p, Rel_{L2,L3})$ to assess the performance of each management option q as a success or a failure:

$$\begin{aligned} & \hat{h}(\Pi_c(Res_o, Cost_t, Energy_p, Env_p, Rel_{L2,L3}), q) \\ &= \max\{h(\min_{u_j \in u(h)} \Pi_{Res_o, Cost_t, Energy_p, Env_p, Rel_{L2,L3}}(u_{1,2,3}, q)) \\ & \geq \Pi_c(Res_o, Cost_t, Energy_p, Env_p, Rel_{L2,L3})\} \end{aligned} \quad (7)$$

Table 10

Robustness to uncertainty for IPs 1–20. Robustness is the maximum number of increments of uncertainty away from the central estimates of inflow, demand and cost for which the system maintained minimum performance requirements in the categories of reservoir susceptibility, service reliability, environmental performance, total cost and energy consumption.

Option	Robustness	Failure criterion
1	FAIL	Reservoir susceptibility, service reliability L2 and L3
2	FAIL	Reservoir susceptibility, service reliability L2 and L3
3	FAIL	Service reliability L2 and L3
4	FAIL	Service reliability L2
5	0	Environmental performance
6	0	Environmental performance
7	3	Environmental performance
8	3	Environmental performance
9	2	Environmental performance
10	2	Environmental performance
11	5	Environmental performance
12	6	Environmental performance
13	7	Environmental performance
14	8	Environmental performance
15	7	Energy consumption
16	7	Energy consumption
17	9	Environmental performance
18	9	Environmental performance
19	5	Total cost
20	2	Total cost

where Res_o is the storage susceptibility; $Cost_t$ is the total cost; $Energy_p$ is the total energy consumption; Env_p is the environmental performance; and $Rel_{L2,L3}$ is the service reliability. The relative and absolute values for each performance criteria are outlined in Section 4.3 and are the same as those implemented in the RDM analysis.

Robustness analysis of the 20 infrastructure portfolios (IPs) required 800 simulations. Our 5 performance indicators: storage susceptibility, service reliability, environmental performance, total cost and total energy consumption are used to assess a simulation as a success or a failure and the interval of uncertainty at which they fail (Table 10). Simulation results showed that portfolio IPs 1–4 failed to meet our minimum performance requirements for the best estimate scenario and IPs 5 and 6 could only cope with the best estimate scenario. Furthermore, IP 7 and 8 only maintained minimum levels of system performance for 3 increments of uncertainty, whilst IP 9 and 10 failed after 2 increments. As a result these 10 IPs were excluded from further analysis due to poor robustness. The remaining 10 IPs included 90 successful simulations. Robustness curves showing performance for IPs 11–20 at each interval of uncertainty are plotted for each performance criteria (Fig. 8).

Results show that IP 17 and 18 are the most robust, tolerating 9 increments of uncertainty from our best estimate. This translates to a future that is 5% drier than our best estimate projection, with a total potable water demand of up to 2179 Ml/day and energy cost of up to 15.25 p/kW h.

IPs 14 is the next most robust option, able to maintain acceptable levels of system performance for 8 increments of uncertainty; a future 4.5% drier than our best estimate with a total demand of up to 2175 Ml/day and energy costs of up to 15 p/kW h). IPs 13, 15 and 16 are able to cope with 7 increments of uncertainty (4% dryer, demand up to 2172 Ml/day, energy cost 14.5 p/kW h). IPs 13, 14, 17 and 18 all fail as a result of environmental performance, whereas IP 15 and 16 fail due to high energy consumption.

IPs with greater investment in supply infrastructures are shown to not be necessarily more robust to uncertainty. This is particularly true for IPs that include DESAL140, which has high energy consumption (resulting in the failures of IPs 15 and 16) and therefore also high operating costs. Combining DESAL140 with the large

capital investment in RST, causes early cost failures in IPs 19 and 20 even though they are able to supply large amounts of water. Although IPs 17 and 18 also include an RST, they only include DE-SAL80 which significantly reduces energy consumption and costs.

Robustness curves for environmental performance, total costs and energy consumption are almost linear reflecting the relatively linear relationship these performance measures have with water availability and water demand. Steep robustness curve gradients indicate minimal performance loss as uncertainty increases, and

the crossing of robustness curves marks the location where one IP becomes relatively more robust than another option.

Storage susceptibility (Fig. 8a) and service reliability (Fig. 8e) produce nonlinear robustness curves. This is because demand and level of service are a function of the time of year and the amount of storage in London reservoirs. As a result there are stages in the Info-Gap analysis where conditions get harsher and the resulting lower reservoir levels activate demand management measures sooner. Service reliability goes down as more weeks

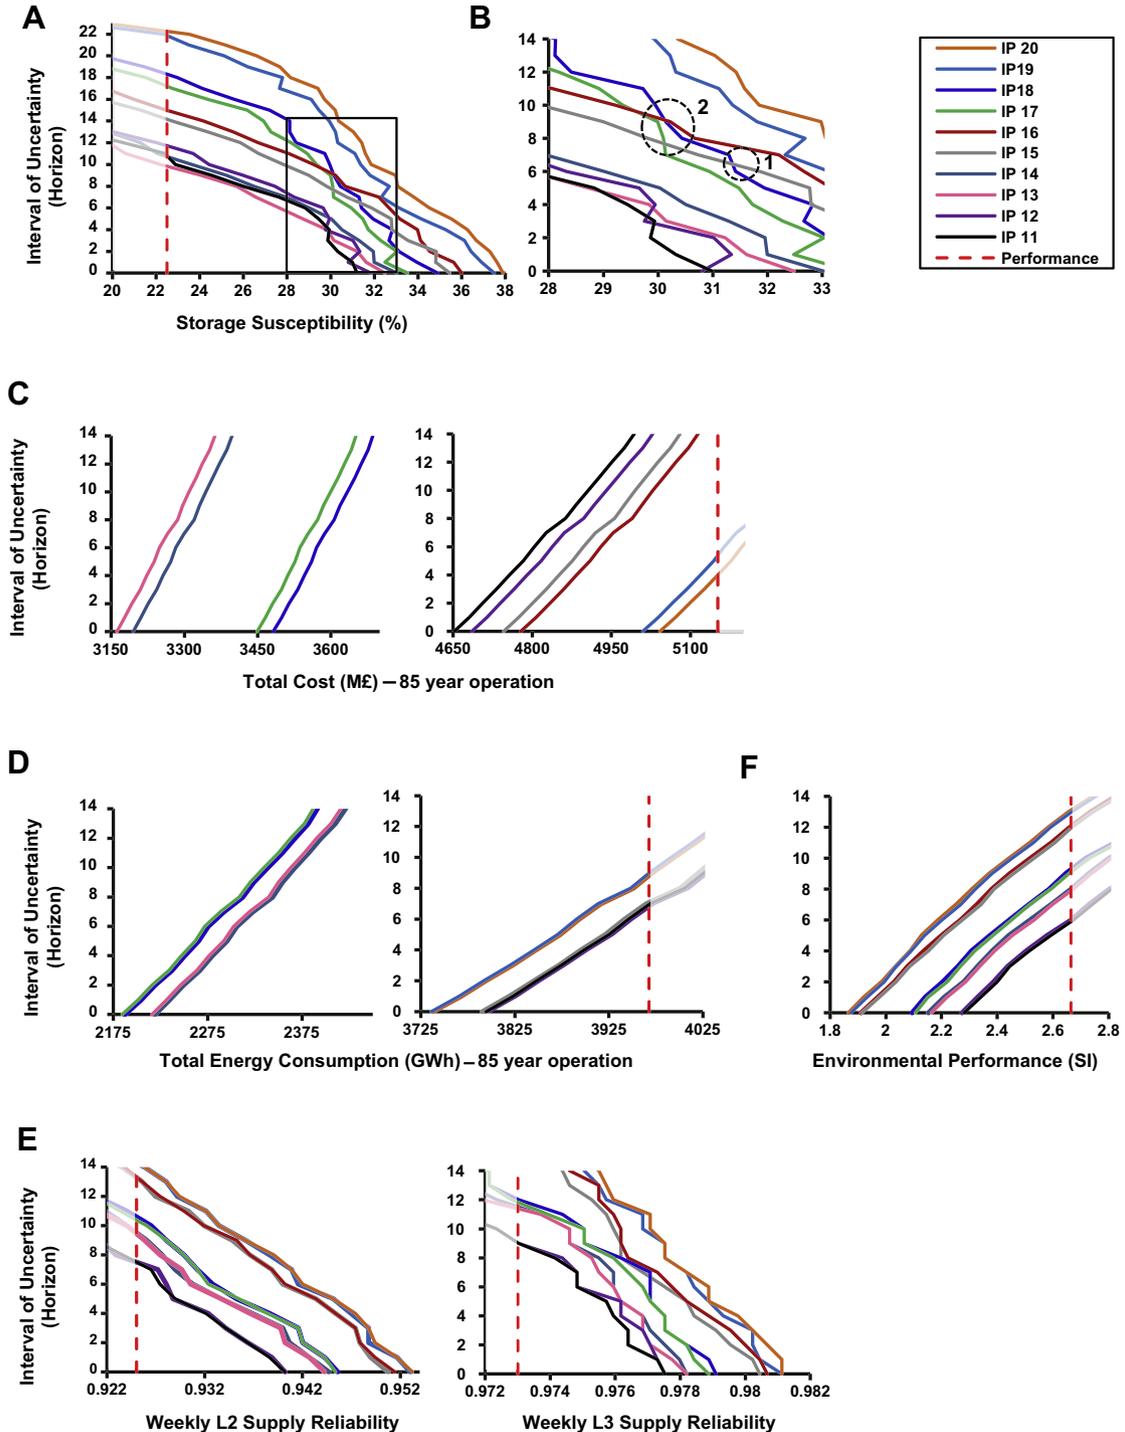


Fig. 8. Robustness curves for IPs 11–20 for storage susceptibility (a and b), cost (c), energy consumption (d), service reliability (e) and environmental performance (f). Results show that for cost, energy and environmental performance there is a linear decrease in performance with increasing demand and decreasing water availability. Storage susceptibility has crossing robustness curves (1 and 2 in Fig. 8A). Where curves cross at 1, IP 18 becomes more robust to uncertainty than IP 15. At 2 the same occurs for IP 17 and IP 18; they both become more robust than IP 16.

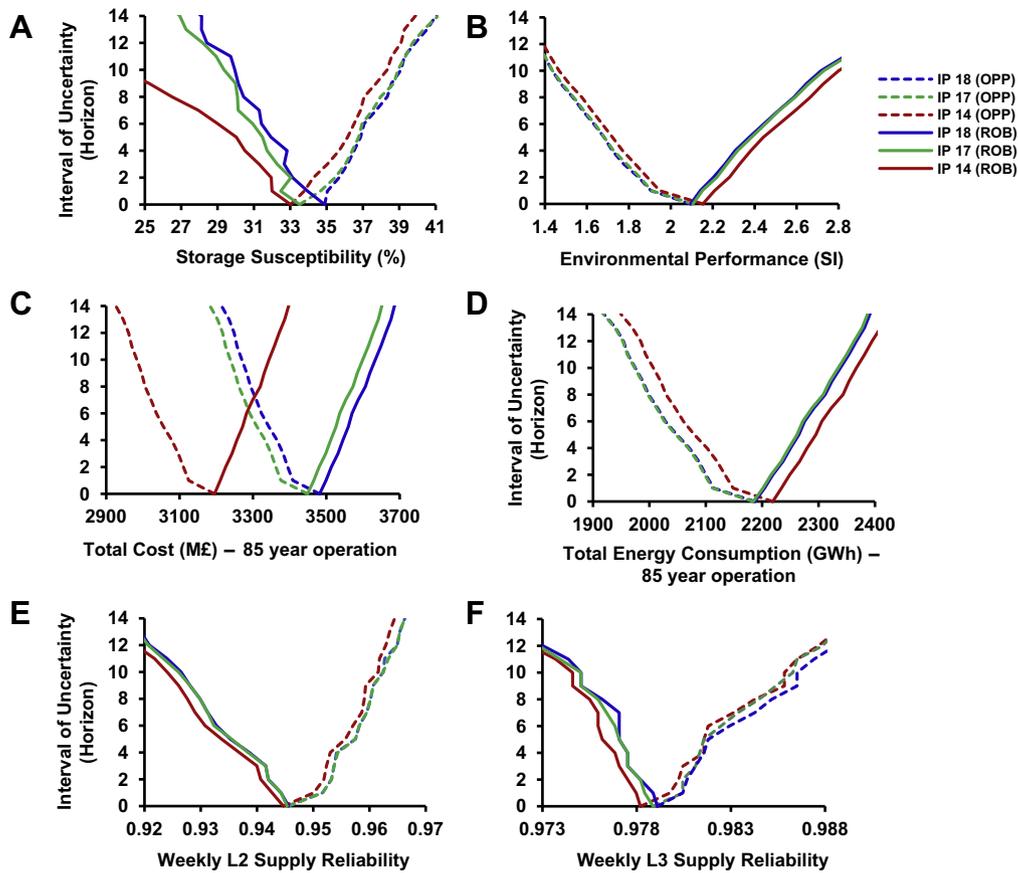


Fig. 9. Opportuneness curves (dotted lines) for IPs 15, 16, 17 and 18 for storage susceptibility (a), environmental performance (b), cost performance (c) and energy consumption (d) and service reliability at L2 (e) and L3 (f). Curves show the improvement in performance for each option as we move incrementally towards increasingly more benign futures. The result show relative linearity between improvements in each performance measure and each successively more benign future. Robustness are also included to aid analysis of possible trade-offs (solid lines). Results show that performance rewards are comparable between the options except for total costs where IP 14 shows significant cost rewards (reductions) compared with IPs 17 and 18. These cost rewards are combined with lower rewards in the remaining performance metrics.

are spent in restrictions, but storage susceptibility improves when compared to the previous more benign interval of uncertainty. This is demonstrated when comparing Fig. 8b and e for IP 19, where a dramatic improvement in storage susceptibility at uncertainty interval 8 compared to interval 7 is matched by a sudden drop in L2 and L3 service reliability over the same 2 increments.

IPs 19 and 20 outperform IPs 13–18 in terms of storage susceptibility at each increment of uncertainty. This is because RST and DESAL140 combine to provide large amounts of water during periods of drier conditions and higher demand. The high energy consumption and total costs causes these options however, to fail after only 5 and 2 increments of uncertainty respectively.

Crossing of robustness curves provides valuable information. IPs 15 and 16 perform better than 17 and 18 in terms of storage susceptibility for the first 6 intervals of uncertainty. However, the steeper gradient of the robustness curves for IPs 17 and 18 shows that their storage susceptibility in this criterion is less affected by the increasingly harsh conditions at each successive interval of uncertainty. At the 7th interval of uncertainty IP 18 crosses the robustness curve for IP 15: beyond this interval of uncertainty IP 18's storage susceptibility is more robust to uncertainty than option 15's. This occurs again at the 8th interval of uncertainty where robustness curves for IP 15 and 17 cross. As conditions get harsher, IPs 17 and 18 perform comparatively better than IPs 15 and 16. Crossing robustness curves are also a feature of the service reliability performance criteria however, the crossings are oscillatory and do not denote a permanent shift.

IP 14 is only one increment of uncertainty less robust when compared to IPs 17 and 18. This IP includes UTR150, DESAL80 and SLARS. It is still an infrastructure heavy option, but it does not include the almost limitless capacity of the RST to maintain supplies. This slightly lower robustness comes with certain benefits: Fig. 9c and d show IP 14 is cheaper and less energy intensive than IPs 17 and 18. IPs 17 and 18 however outperform IP 14 in environmental performance, service reliability (L2 and L3) and storage susceptibility. The presence of SLARS in IP 18 makes it outperform IP 17 for storage susceptibility, environmental performance and service reliability (L2 and L3) at each increment of uncertainty at the expense of greater costs and energy consumption.

Results from the robustness analysis show IPs 14, 17 and 18 to be robust management options. Trade-offs between robustness in service reliability, environmental and storage susceptibility are at the expense of higher total costs and energy consumption, making it difficult to distinguish between these 3 IPs.

6.2. Opportuneness function

The opportuneness function is derived from Eq. (6) and is the inverse of the robustness function. Opportuneness analysis quantifies the performance reward in each of our performance criteria $\Pi_t(Res_t, Cost_t, Energy_p, Env_p, Rel_{L2,L3})$ achieved in more benign futures derived from the right side of the Info-Gap uncertainty model. These futures become increasingly wetter, with lower

demand and energy costs at each successive interval of uncertainty away from the best estimate:

$$\hat{\beta}(\Pi_r(\text{Res}_o, \text{Cost}_t, \text{Energy}_p, \text{Env}_p, \text{Rel}_{L2,L3}), q) = \min \left\{ h \left(\max_{u_i \in u(h)} \Pi_{\text{Res}_o, \text{Cost}_t, \text{Energy}_p, \text{Env}_p, \text{Rel}_{L2,L3}}(u_{1,2,3}, q) \right) \geq \Pi_r(\text{Res}_o, \text{Cost}_t, \text{Energy}_p, \text{Env}_p, \text{Rel}_{L2,L3}) \right\} \quad (8)$$

Opportuneness analysis included 40 simulations each for IPs 14, 17 and 18. As with robustness analysis, curves are plotted for the performance of each option for each increment of uncertainty away from the central estimate. Opportuneness analysis explores system performance in futures that are less harsh. Whilst we are planning for system robustness, opportuneness can help choose between IPs with similar robustness (e.g. IPs 14, 17 and 18). Crossing curves denote when one IP gives greater reward than another. However, in opportuneness analysis we are looking for the lowest curves with shallow gradients, indicating a high performance reward for small increments of uncertainty.

Opportuneness curves for IPs 14, 17 and 18 are plotted for each of our performance criteria (Fig. 9a–f). The corresponding robustness curves are also included to aid interpretation of possible trade-offs between robustness and opportuneness. The curves are relatively linear for environmental performance, total cost and energy consumption. For service reliability at L2 and L3 there are step changes in calculated reward but the responses are comparable for all 3 IPs and do not result in crossing curves. Storage susceptibility opportuneness curves are quite linear, although Fig. 9a shows IP 17's opportuneness curve to have a shallower gradients than IPs 14 and 18 for the first two increments of uncertainty.

Without any crossing opportuneness curves the ranking of the IPs is unambiguous. IPs 17 and 18 consistently give similar levels of reward and IP 14 gives less reward for all criteria compared to IP 17 and IP 18 except total costs. IP 14 is significantly cheaper at each interval of uncertainty; on average it is approximately £3.3M cheaper to operate each year than IPs 17 and 18 (based on it being £280M cheaper over the 85 years of simulation). This greater cost reward must be traded-off against a lower overall robustness.

Based on this analysis IGDT promotes IP 18 (current infrastructure, RST and DESAL80) as the preferred water supply infrastructure portfolio. Compared to IP 14, IP 18 is more robust overall and performs better at each increment of uncertainty for all performance criteria except cost. IP 14's cost opportuneness is counter-balanced by its lower robustness. IPs 17 and 18 are similar; IP 17 is more robust to cost uncertainty but IP 18 outperforms IP 17 for all other criteria. IGDT leads us to select IP 18 as the preferred plan for a conservative planner. If less deviation from the future best estimate were to occur IP 14 would be more appealing for its robustness to cost increases in harsher futures and greater cost rewards in more benign futures.

7. Discussion

7.1. RDM discussion

An RDM analysis begins by the selection of a candidate strategy. In this case study, the RDM process began with a multi-criteria regret analysis. We generated a solution space of 311,080 simulations using a computationally efficient system simulator. The solution space included dire, mild and opportune plausible futures allowing us to compare IPs under a wide range of conditions. Each IP was compared to the best performing IP for a given set of input parameters. We used median relative regret to select IP 14 as the

candidate strategy. IPs 13, 18 and 17 had the next lowest median regret.

The initial regret analysis also provided information on the relative performance of each IP. SLARS has minimal effect on the levels of regret and IPs with DESAL140 have higher regret in the cost and energy criteria than those with DESAL80. For storage susceptibility, service reliability and environmental performance however, DESAL140 improved regret in similar IPs.

The second step of the RDM framework involved a scenario discovery process using PRIM to isolate scenarios in which the candidate strategy is likely to fail. Two scenarios were identified: one which includes futures with dry summers and moderate-to-high demand and a second with dry summers, very dry winters and very low demand. In total 78% of failure points were covered in the three scenarios albeit with a high density of 95% (compared to the 55% failure density in the whole of the solution space). Given the spread of failure points around the solution space, PRIM was unable to produce scenarios with a better coverage to density ratio. This is a result of having multiple performance criteria that conflict; i.e. better environmental, storage and reliability performance often results in high energy consumption and high costs generating failures that are spread throughout the solution space.

A relative regret analysis of the simulations inside and outside vulnerability scenarios 1 and 2 showed that while IP 14 has the lowest median regret outside the vulnerability scenarios, IP 18 has the lowest median regret inside the vulnerability scenarios. The UTR150 in IP 14 and the RST in IP 18 provide the same amount of water to London when first activated during times of drought. However, because of how it is modelled the RST provides this release continuously and is unaffected by Thames area droughts (unlike the UTR).

From scenario 1 (which includes dry summers) and the regret analysis on scenarios 1 and 2, we identified IP 14 as vulnerable to failure because of its inadequate storage under droughts making it unable to compensate for the summer/winter hydrology disparity in climates that have dry summers. We identified that the RST is a useful, but an expensive and energy intensive alternative to address this disparity. We assumed the Severn basin will not experience drought at the same time as the Thames basin. A spatial drought correlation study would be required to establish the reliability of the RST and consider it in this study. This implies we cannot make the 'IP 14 or 18' decision for planners and that the relative strength of IPs that include the RST could be further lowered in regards to reliability. The price of energy was not found to affect the vulnerability of IP 14 because the majority of system failures come from the service reliability criteria.

In a full RDM analysis, decision makers could improve IP 14 and possibly IP 18 devising new IPs which could be resubmitted into the RDM process. The process would be repeated until a suitably robust IP is identified.

7.2. IGDT discussion

The Info-Gap analysis was implemented using the fractional-error Info-Gap uncertainty model (Eq. (6)). This model characterised uncertain system input parameters and provided an interval-based approach to progressively sample harsher and more benign futures as compared to a best estimate.

The IGDT approach uses the vector h to incrementally sample the uncertain space. Because the three uncertain system parameters (inflows, demands, energy costs) are all scaled by h for each 'horizon' or increment of uncertainty, all parameters change proportionately; becoming increasingly harsher or more benign in robustness and opportuneness analysis respectively. Inflows lower whilst demands and energy costs increase at each successive interval of uncertainty. Although these parameters can be scaled and

constrained within the Info-Gap uncertainty model to better represent natural relationships (Hall et al., 2011), conditions do change concurrently and proportionately. This has the advantage of requiring fewer simulations: only 40 simulations for each IP rather than a full grid search of 64,000 simulations for each IP (40 inflow, 40 demand and 40 energy intervals). One potential draw-back of the fractional-error IGDT approach is that it does not sample conditions where for instance; demand decreases but energy costs increase relative to our best estimate, or where inflow and demand are both lower than expected (vulnerable conditions as described by RDM's Scenario 2).

Results from our Info-Gap analysis show IP 18 is the most robust future infrastructure portfolio for maintaining adequate supply to London under harsher futures that are up to 5% dryer, with up to 2179 Ml/d potable water demand and energy costs up to 15.25 p/kW h. IP 18 maintains minimum system performance requirements for the greatest number of uncertainty intervals from our best estimate whilst maintaining water supplies to London (i.e. IP 18 fails as a result of environmental performance rather than service reliability). Plotting of robustness curves allows comparing the performance of different IPs for each metric at structured intervals of uncertainty around a best estimate. Robustness curves show IP 18 is the most robust to uncertainty in terms of storage susceptibility, service reliability and environmental performance. The steeper gradient of IP 18's robustness curves for storage susceptibility demonstrate there is less loss in performance at each successive interval of uncertainty when compared to IPs 15 and 16. The crossing of robustness curves at points 1 and 2 in Fig. 8b, marks the location where IP 18 becomes more robust than IP 15 and 16. The crossing of robustness curves results from the complex interactions between the supply infrastructure, environmental flow and demand reductions as dictated by the LTCD.

IGDT analysis suggests that IP 14 is also a good candidate strategy. This result is subject to the same caveats on RST as in RDM. Results from robustness and opportuneness analysis show IP 14 could be preferable to 18 if a decision maker could be more certain that the future will diverge less from the best estimate, or that the future will be more benign. Under both these conditions the improved cost performance of IP 14 compared to IP 18 for small deviations from the best estimate would potentially promote this as the most robust IP.

7.3. Thames study limitations

Several limitations affect the relevance of the results for real Thames region planning. The England and Wales water sector follows a 'twin-track' approach including both supply augmentation and demand management (e.g. water conservation, increased efficiency, metering, leakage reduction, communication campaigns, etc.). Our study only considers supply augmentation schemes to balance supply and demand; further research will bring in demand management options. The flow perturbation factors are applied to the historical flow time-series and do not take into account possible shifts in the hydrologic regime; they uniformly shift the intensity of flows (i.e. both low and high flows are equally augmented) and do not change the frequency or pattern of severe events (Prudhomme et al., 2010). Furthermore, the scaling of the flow factors in the IGDT analysis results in a compression or extension of the range of seasonal ranges of flows resulting in a further departure from the historical flow hydrology regime. Further research into methods that produce hydrologically consistent future flow time series is recommended and would strengthen the validity of the IGDT and RDM analyses. Uncertainty about the reliability of the River Severn Transfer should be considered, but without a water resource model of that area we could not consider it in this study. This would add a further dimension of uncertainty and signifi-

cantly increase the computational burden of the planning problem. Further work could include this source of uncertainty and use sparser sampling methods for RDM. These improvements would increase the planning relevance of study results but would likely not significantly change methodological implications, the focus of this paper.

7.4. Complementarity of RDM and IGDT

Both RDM and IGDT helped assess the robustness of 20 different water supply infrastructure portfolios (IPs) to Knightian or 'deep' uncertainty in hydrological inflows, water demands and energy costs for a 2035 planning horizon. Due to differences in their approaches they reached similar but not entirely matching results. IGDT promotes IP 18 as the infrastructure portfolio of preference whereas RDM initially suggests IP 14. This observation leads to the main finding of this paper: because RDM and IGDT provide planners with different and complementary information they can be beneficially used together. We recommend joint use of IGDT and RDM for the planning of water resource systems as each method can help clarify the results of the other.

IGDT begins with a best estimate for each uncertain system input and then sequentially chooses values increasingly farther away from expected inputs. The best estimate of future conditions we used has a greater disparity between winter and summer flows, higher water demands and higher energy costs than the historical average. As a result the IGDT analysis, even at a zero interval of uncertainty, begins with the system under considerable stress before it considers less or more favourable conditions.

RDM samples from all combinations of the uncertain system parameters to identify conditions of system failure. RDM samples the 'extreme' futures as in IGDT but also more benign ones. In both IGDT and RDM, the robustness of simulated futures is assessed using multiple criteria such that IPs that are less infrastructure intensive (less costly and energy intensive) outperform more infrastructure intensive IPs in benign future states. These benign future states are taken into account during RDM's initial regret analysis and reduce the multi-criteria regret values of less infrastructure intensive options that perform reasonably well in storage reliability, storage susceptibility and environmental categories. IGDT does not consider these more benign futures and therefore less infrastructure intensive IPs do not perform well. For more benign futures, IP 14 has a lower median regret than IP 18, the more infrastructure intensive alternative while for more dire futures (Fig. 7), which more closely relate to the futures sampled in IGDT, IP 18 performs better. When only considering harsh futures, IGDT and RDM recommend the same future plan. Indeed, RDM's vulnerability scenario 1 (Table 6) shows IP 14 is vulnerable to failure under conditions that are only slightly harsher and even sometimes more benign (in the demand uncertainty dimension) than those that define the best estimate for the IGDT analysis.

Because IGDT assesses robustness of each option under progressively more dire futures, there is a tendency to favour IPs with infrastructure-intensive solutions. Using IGDT, it is difficult to jointly consider information about an IP's robustness to uncertainty and opportuneness for a performance reward. The regret analysis in our RDM implementation produced a single plot to characterise performance across all futures (Fig. 4). In this way RDM provides complementary information to an IGDT analysis on the suitability of future plans across a range of dire, benign and opportune futures. Our RDM application found solutions that are robust over a wide range of future states whilst our IGDT analysis identified solutions that are robust over dire futures and opportunistic over favourable futures. Taking into account the robustness of IPs for mild futures in water planning studies helps prevent decision makers from automatically favouring conserva-

tive (in this case infrastructure intensive) options. At the same time IGDT underlined the fact that our preferred option under RDM was quite vulnerable to dire futures.

IGDT robustness curves provide decision makers with useful visualisations that complement an RDM analysis. By plotting the performance of each option for each metric at each increment of uncertainty, trends in performance under increasingly dire or favourable conditions are revealed. This was seen for example in the service reliability performance criteria for IP 18 (Fig. 8e): initially IP 16 performs better but as conditions get worse IP 18 improves revealing it performs increasingly better under harsher conditions relative to IP 16. The IGDT plots show choosing IP 18 over IP 16 results in a cheaper and less energy intensive (Fig. 8c and d) system in harsher futures. IGDT robustness analysis also shows which performance criterion each IP fails at first whilst RDM does not track which criterion caused the failure. In this way we might consider environmental performance as a 'soft' performance rule which in reality could be reduced temporarily during droughts. IGDT tells us that under very harsh conditions, IP 18 fails environmental performance first, but if this is a shifting goal-post during droughts, IP 18 would still be acceptable (i.e. service reliability is not the first requirement to fail under harsh conditions).

The sampling in RDM meant a wide range of dire, benign and favourable conditions were considered during plan evaluation. RDM's scenario discovery approach identified the vulnerable sets of uncertain conditions for the preferred plan. These vulnerability scenarios were then used to check if other infrastructure portfolios could actually perform better under the rough conditions. Our RDM application showed IP 18 was the most robust option for more severe conditions; this was the preferred option from our IGDT analysis.

Our RDM implementation required 311,080 simulations and IGDT 1600 simulations. The supplementary runs required for IGDT suggest it is worthwhile to consider joint application. For example IGDT could be used along-side regret analysis to identify the initial candidate strategy, or to propose an alternative one for analysis in RDM's scenario discovery phase. Also, IGDT can be used to thin out the option sets considered in RDM; this could be useful when there are many possible initial options or plans. As discussed in Section 2.2. Info-Gap assumes analysts can provide a best estimate around which the uncertain parameter is assumed to lie and therefore does not explore the full uncertainty space. We found that complementing an IGDT approach with another such as RDM which samples future conditions more widely is an effective guard against potentially missing critical design conditions. Info-Gap can effectively be used to explore local robustness around selected scenarios in the uncertainty space. The overarching benefit of using RDM and IGDT methods is that they provide a structured approach to simulating systems under Knightian uncertainty. The methods select relevant system conditions to simulate and reveal how different plans perform under these conditions. Joint use widens the insights gained and helps verify and deepen the understanding that arises from using either method.

8. Conclusion

In this paper we formulated and applied the IGDT and RDM planning frameworks to a water resource management problem: expanding London's water supply system in the Thames basin, UK. We show that although the methods initially provide different capacity expansion recommendations, they produce broadly similar results but provide complementary information about the performance of different proposed infrastructure portfolios. Joint use of IGDT and RDM for the planning of water resource systems helps better understand the results of each method.

IGDT efficiently evaluates system performance under different options considering the most dire and opportune future conditions. The framework identifies which criteria cause failure of a future plan and allows to plot the performance of each option as future system inputs progressively deviate from an initial estimate. It may however ignore parts of a Knightian uncertainty space if the initial estimate ends up not being a good guess. RDM samples a wider set of combinations of the uncertain variables allowing a more holistic, albeit more computational intensive assessment of the future performance of different plans. Through RDM's scenario discovery combinations of uncertain inputs that lead to system vulnerabilities are identified. Regret analysis efficiently quantifies how different future plans fare over the vulnerable scenarios considering multiple performance criteria. RDM provides a structured framework for iterative refinement of future plans.

RDM and IGDT offer structured and insightful frameworks to plan complex systems with multiple requirements subject to multiple unknown future stresses. Joint use of the methods can make them more computationally efficient and maximise understanding by revealing how each method can skew results towards particular future plans.

Acknowledgements

The authors thank Ke Zhou for his help with early RDM work. The UK EPSRC funded the work through the ARCC-Water project (EP/G028540/1); A. Woods was partially supported by AECOM. Steven Wade and Brigitte von Christierson from HR Wallingford provided the climate change perturbation ratios. The authors thank Rob Wilby, the A.E. and a reviewer for comments that strengthened the article; the usual disclaimer applies.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.jhydrol.2013.03.006>.

References

- BBC, 2007. Water Firm may Lift Hosepipe Ban. <<http://news.bbc.co.uk/1/hi/uk/6259763.stm>> (accessed 15.07.11).
- Ben-Haim, Y., 2001. Information-Gap Decision Theory. Academic Press, San Diego, CA.
- Ben-Haim, Y., 2010. Info-Gap Economics: An Operational Introduction. Palgrave Macmillan, Basingstoke.
- Beven, K., 2006. A manifesto for the equifinality thesis. *J. Hydrol.* 320 (1–2), 18–36.
- Beven, K., Binley, A., 1992. The future of distributed models – model calibration and uncertainty prediction. *Hydrol. Process.* 6 (3), 279–298.
- Bryant, B.P., Lempert, R.J., 2010. Thinking inside the box: a participatory, computer-assisted approach to scenario discovery. *Technol. Forecast. Soc. Change* 77 (1), 34–49.
- Christierson, B.V., Vidal, J.-P., Wade, S.D., 2012. Using UKCP09 probabilistic climate information for UK water resource planning. *J. Hydrol.* 424, 48–67.
- Cole, G.A., Marsh, T.J., 2006. An historical analysis of drought in England and Wales. In: Demuth, S., Gustard, A., Planos, E., Scatena, F., Servat, E. (Eds.), *Climate Variability and Change – Hydrological Impacts*. Iahs Publication, pp. 483–489.
- Cookson, A., Weston, A., 2008. Thames Water Modelling Project: Model 1a Final Report. Environment Agency.
- Dessai, S., Hulme, M., 2007. Assessing the robustness of adaptation decisions to climate change uncertainties: a case study on water resources management in the East of England. *Global Environ. Change – Human Policy Dimensions* 17 (1), 59–72.
- Eurostat, 2011. Electricity Prices for Industrial Consumers. <<http://epp.eurostat.ec.europa.eu/tgm/table.do?tab=table&init=1&plugin=1&language=en&pcode=ten00114>> (accessed 06.02.11).
- Fredrich, A.J., 1975. *Hydrologic Engineering Methods for Water Resources Development*. Davis, California.
- Friedman, J.H., Fisher, N.I., 1999. Bump hunting in high-dimensional data. *Stat. Comput.* 9 (2), 123–143.

- Groves, D.G., Lempert, R.J., 2007. A new analytic method for finding policy-relevant scenarios. *Global Environ. Change – Human Policy Dimensions* 17 (1), 73–85.
- Hall, J. et al., 2011. Robust climate policies under uncertainty: a comparison of robust decision-making and info-gap methods. *Risk Anal.*
- Halrow Ltd., 2010. Thames Valley Stored Water System. <<http://www.halcrow.com/Our-projects/Project-details/Thames-Stored-Water-System-England>> (accessed 28.09.10).
- Hine, D., Hall, J.W., 2010. Information gap analysis of flood model uncertainties and regional frequency analysis. *Water Resour. Res.*, 46.
- Hipel, K.W., Ben-Haim, Y., 1999. Decision making in an uncertain world: information-gap modeling in water resources management. *IEEE Trans. Syst. Man Cybern. Part C – Appl. Rev.* 29 (4), 506–517.
- Johns, T.C. et al., 2003. Anthropogenic climate change for 1860 to 2100 simulated with the HadCM3 model under updated emissions scenarios. *Clim. Dyn.* 20 (6), 583–612.
- Jones, A.E., 1983. Flow Gauging on the River Thames – The First 100 Years. National River Flow Archive, UK.
- Jones, M.A., Harris, S.J., Baxter, K.M., Anderson, M., 2005. The Streatham groundwater source: an analogue for the development of recharge enhanced groundwater resource management in the London basin. In: *International Symposium on Management of Aquifer Recharge*. United Nations Educational, Scientific and Cultural Organization, Berlin, Germany.
- Knight, F.H., 1921. *Risk, Uncertainty, and Profit*. Houghton Mifflin, Boston, MA.
- Lempert, R., Groves, D., Popper, S., Bankes, S., 2006. A general, analytic method for generating robust strategies and narrative scenarios. *Manage. Sci.* 52 (4), 514–528.
- Lempert, R.J., Collins, M.T., 2007. Managing the risk of uncertain threshold responses: comparison of robust, optimum, and precautionary approaches. *Risk Anal.* 27 (4), 1009–1026.
- Lempert, R.J., Groves, D.G., 2010. Identifying and evaluating robust adaptive policy responses to climate change for water management agencies in the American west. *Technol. Forecast. Soc. Change* 77 (6), 960–974.
- Lempert, R.J., Popper, S.W.C.B.S., 2003. *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. RAND Corporation, Santa Monica, CA.
- Marsh, T., Cole, G., Wilby, R., 2007. Major droughts in England and Wales, 1800–2006. *Weather* 62 (4).
- Matrosov, E.S., Harou, J.J., Loucks, D.P., 2011. A computationally efficient open-source water resource system simulator – application to London and the Thames Basin. *Environ. Modell. Software* 26 (12), 1599–1610.
- McCarthy, M.A., Lindenmayer, D.B., 2007. Info-gap decision theory for assessing the management of catchments for timber production and urban water supply. *Environ. Manage.* 39 (4), 553–562.
- Merrett, S., 2007. The Thames catchment: a river basin at the tipping point. *Water Policy* 9 (4), 393–404.
- Moore, R.J., 1985. The probability-distributed principle and runoff production at point and basin scales. *Hydrol. Sci. J. – J. Des Sci. Hydrol.* 30 (2), 273–297.
- Murphy, J.M. et al., 2007. A methodology for probabilistic predictions of regional climate change from perturbed physics ensembles. *Philos. Trans. R. Soc. A – Math. Phys. Eng. Sci.* 365 (1857), 1993–2028.
- Murphy, J.M. et al., 2009. *UK Climate Projections Science Report: Climate Change Projections*. Met Office, Exeter, UK.
- Nakicenovic, N., Swart, R. (Eds.), 2000. *Special Report on Emissions Scenarios*. A Special Report of Working Group III of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK.
- Office for National Statistics, 2008. London: Population and Migration. <<http://www.statistics.gov.uk/ci/nugget.asp?id=2235>> (accessed 10.06.11).
- OFWAT, 2010. *A Study on Potential Benefits of Upstream Markets in the Water Sector in England and Wales*. Birmingham, England.
- Prudhomme, C., Wilby, R.L., Crooks, S., Kay, A.L., Reynard, N.S., 2010. Scenario-neutral approach to climate change impact studies: application to flood risk. *J. Hydrol.* 390 (3–4), 198–209.
- Raluy, R.G., Serra, L., Uche, J., 2005a. Life cycle assessment of water production technologies – Part 1: life cycle assessment of different commercial desalination technologies (MSF, MED, RO). *Int. J. Life Cycle Assess.* 10 (4), 285–293.
- Raluy, R.G., Serra, L., Uche, J., Valero, A., 2005b. Life cycle assessment of water production technologies – Part 2: reverse osmosis desalination versus the Ebro river water transfer. *Int. J. Life Cycle Assess.* 10 (5), 346–354.
- Savage, L.J., 1954. *The Foundation of Statistics*. Dover Publications.
- Simon, H., 1959. Theories of decision-making in economic and behavioral science. *Am. Econ. Rev.* 49 (3), 253–283.
- Stokes, J., Horvath, A., 2006. Life cycle energy assessment of alternative water supply systems. *Int. J. Life Cycle Assess.* 11 (5), 335–343.
- Thames Water, 2010. *Revised Water Resources Management Plan*. Thames Water, UK.
- WaterBriefing, 2006. *Hopepipe Ban Remains while Drought Order Goes*. <<http://www.waterbriefing.org/index.php/home/regulation-and-legislation/item/3291-hosepipe-ban-remains-while-drought-order-goes>> (accessed 01.06.11).
- Wilby, R., Greenfield, B., Glenny, C., 1994. A coupled synoptic hydrological model for climate-change impact assessment. *J. Hydrol.* 153 (1–4), 265–290.