Submitted to *Operations Research*

manuscript (Please, provide the manuscript number!)

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Ambiguity Updating Robust Architecture:

Towards Time Consistency in Multi-period DRO

Wenxin Liu, Yu Zhang

School of Business Administration, Southwestern University of Finance and Economics, Sichuan 611130, China, liuwenxin921@gmail.com, waltyuzhang@gmail.com

Aakil Caunhye

The University of Edinburgh Business School, 29 Buccleuch Place, Edinburgh EH8 9JS, United Kingdom

Gar Goei Loke

Department of Management and Marketing, Durham University Business School, Waterside, Durham DH1 1SL, United Kingdom, gar.g.loke@durham.ac.uk

Multi-period distributionally robust optimization (DRO) is in general not time consistent. For a given multi-period DRO problem with generalized moments ambiguity, we propose an Ambiguity Updating Robust Architecture (AURA) that constructs the ambiguity set of all conditional distributions-to-go, given the realized uncertainty. We prove that AURA is time consistent, but is more conservative than the original DRO and is equal if and only if DRO was originally time consistent. AURA's robust counterpart is a deterministic dynamic program. When the support of the uncertainty is continuous, AURA contains all deterministic distributions, thus is as conservative as parameter robustness. We term this "ambiguity collapse". We trace the origins of ambiguity collapse to distributions that become deterministic after some particular history is traversed. We term these "conditionally clairvoyant" distributions. We propose a *regularization* of the original DRO ambiguity set to remove such distributions; doing so allows ambiguity collapse to be avoided. These results are illustrated using a stylized numerical experiment on a water management setting.

Key words: distributionally robust optimization, multi-period optimization, water management

1. Introduction

We consider the stochastic and robust optimization settings for optimization under uncertainty,

$$\min_{\boldsymbol{x}} \mathbb{E}_{\mathbb{P}}\left[c(\boldsymbol{x},\boldsymbol{w})\right] \qquad \text{or} \qquad \min_{\boldsymbol{x}} \sup_{\mathbb{Q}\in\mathcal{U}} \mathbb{E}_{\mathbb{Q}}\left[c(\boldsymbol{x},\boldsymbol{w})\right],$$

where $\boldsymbol{x}$ are decision variables, $\mathbb{P}$ is the true distribution generating uncertainty $\boldsymbol{w}$, and c the cost function to minimize. When true $\mathbb{P}$ is not known, its distributionally robust analogue optimizes against a worst-case distribution $\mathbb{Q}$ in an ambiguity set $\mathcal{U}$ (Wiesemann et al. 2014), such as the

Kullback-Leibler (Calafiore 2007), Wasserstein (Mohajerin Esfahani and Kuhn 2018), and scenario-wise (Chen et al. 2020) ambiguity sets. We study the multi-period problem

$$\min_{\boldsymbol{x}_t(\cdot):0\leqslant t\leqslant T} \quad c_0(\boldsymbol{x}_0)+\sup_{\mathbb{Q}\in\mathcal{U}}\mathbb{E}_{\mathbb{Q}}\left[\sum_{t=1}^{T} c_t\Big(\boldsymbol{x}_t(\boldsymbol{w}_1,\ldots,\boldsymbol{w}_t),\boldsymbol{w}_1,\ldots,\boldsymbol{w}_t\Big)\right],$$

where adaptive non-anticipative decisions $\boldsymbol{x}_t(\cdot)$ are made sequentially to ambiguities revealed over time $\boldsymbol{w}_t$ (Ben-Tal et al. 2004). Multi-period distributionally robust optimization (DRO) can be difficult to solve and intractable (Shapiro and Nemirovski 2005). Often, optimal policies are sought within restricted classes, such as linear decision rules (*e.g.* Bertsimas et al. 2019); and under certain assumptions, their optimality can be proven (Bertsimas et al. 2009, Long et al. 2024).

In many practical situations, as uncertainties materialize, they provide information about subsequent uncertainties through temporal dependence. The decision-maker stands to improve subsequent decision quality, by focusing on only the remnant uncertainties and distributions that can occur, *given* the materialized uncertainty. In this work, we examine how this can be done.

1.1. Capturing time dynamics in multi-period DRO

Capturing temporal dependence can lead to more meaningful solutions in multi-period DRO. In early works, the ambiguity set is specified in a way that indirectly captures some dependence, for example covariance in a moments ambiguity set (Delage and Ye 2010), or the empirical distribution in data-driven ambiguity sets. However, ambiguities over time often arise through some underlying stochastic process or dynamics (Shapiro et al. 2021). Additionally, the rise of machine learning has allowed researchers to model ambiguity in more precise ways and this sparked the present exploration into how this precision can be incorporated into optimization (Bertsimas and Kallus 2020, Elmachtoub and Grigas 2022). We also see such trends in DRO, where ambiguity sets incorporate structure from learning models (Zhu et al. 2022, Sim et al. 2025), though this is still largely confined to simpler learners. These forces motivate the search for better ways to incorporate time dynamics in ambiguity sets.

Markov chains in Markov Decision Processes (MDPs) is one way to define time dynamics. MDPs can be sensitive to transition probability estimates (Xu and Mannor 2006), which provides justification for incorporating robustness. Wiesemann et al. (2013) study robust MDPs, by constructing an ambiguity set of the confidence region containing the true transition probabilities. Except in specific settings (*e.g.* linear quadratic control, Taskesen et al. 2023), MDPs or their robust analogues are not necessarily tractable, leading to efforts to improve algorithmic efficiency, such as Georghiou et al. (2019) who propose deterministic lower and upper bounds on future costs and updating them recursively. A wholly separate approach that retains generality in the structure of

the time dynamics, while appealing to stochastic processes, is Xin and Goldberg (2022). Here, the authors examine the newsvendor context and propose an ambiguity set for the demand distribution containing only martingales.

In these works, the ambiguity sets are defined around general models of stochastic processes, but their tractability may not be guaranteed – at least not as tractable as the usual ambiguity sets that do not explicitly model time dynamics. It is also unclear if they permit tractable approximations, such as linear decision rules. Thus, the value obtained through modeling precision of the time dynamics is not necessarily offset by costs to tractability.

To circumvent tractability limitations, an alternative approach works directly with conditional distributions. Conditional distributions can be modeled through copulas in the DRO context (Gao and Kleywegt 2017), and is often invoked in the contextual optimization setting (Esteban-Pérez and Morales 2022). In multi-period DRO, Hu et al. (2025) propose a time series ambiguity set where the coefficients of the time series are known, but its shock terms are not. Such models are often more tractable, however, they require precise specification of the conditional distribution. For example, Hu et al. (2025) assume that the conditional distribution of the uncertainty is known given the distribution of the shock term in the previous period. Such information is unavailable in reality, or might lead to compound errors, where errors in estimation carry over into the definition of the ambiguity set (Kannan et al. 2024).

In this work, we examine if ambiguity sets of conditional distributions can be constructed without restricting the form of the conditional distributions or assuming independence. Central to our approach is that we utilize Bayesian updating within the ambiguity set itself. Bayesian updating does feature in the context of multi-period DRO, however, it is often applied to redefine the problem (Lin et al. 2022), or in rolling horizon implementations where the model is optimized without knowledge that Bayesian updating will subsequently be undertaken (Shapiro et al. 2023).

1.2. Rolling horizon and time consistency

In practice, when multi-period DRO models are solved, the decision policy $\boldsymbol{x}(\cdot)$ need not be implemented *faithfully*, but in a *rolling horizon* fashion (Chand et al. 2002) – that is, only implementing here-and-now decisions and then re-solving the problem-to-go with the resultant state, after the uncertainty is revealed, as the initial state. Based on *empirical and practical experience*, this often results in better decision quality. Here lies the problem – if one obtains a better solution by resolving the model with knowledge of the revealed uncertainty, this solution is the very *certificate of sub-optimality* of the original decision. This problem is termed *time consistency*, *i.e.*, does optimal policy $\boldsymbol{x}(\cdot)$ coincide with optimal policies of problems-to-go post-revelation of uncertainties? Time

consistency has been studied for stochastic programs or MDPs (Shapiro 2009, Lin et al. 2022), mainly focusing on necessary and sufficient conditions for time consistency.

Specific to the DRO setting, a more perverse situation can occur, where the problems-to-go can become infeasible or unbounded, despite the goal of robustness! For example, Shapiro and Xin (2020) show in the inventory setting, the existence of infinitely many optimal policies that are neither time consistent nor achieve the target inventory level, when the risk measure corresponding to the ambiguity set lacks strict monotonicity. Rectangularity has been proposed as a sufficient condition for time consistency (Shapiro 2016), and has been utilized to prove time consistency (such as in Xin and Goldberg 2022). Temporal independence is a special case that guarantees time consistency, for example, in Yu and Shen (2022), though their focus is on decision-dependence over time, and Gao et al. (2024), where their proposed multistage DRO model with nested distance is time consistent if the nominal distribution is independent in time.

There is a general recognition that characterizing the conditions for time consistency in multi-period DRO is not easy. In this paper, we take a different approach. Instead of determining if a given multi-period DRO problem is time consistent, we attempt to construct the smallest "completion" of the ambiguity set that would make the DRO problem time consistent.

1.3. Contributions

In this paper, we propose a way to construct a time-consistent version of any given multi-period DRO problem by updating the ambiguity set as uncertainty is revealed sequentially.

a) We propose the Ambiguity Updating Robust Architecture (AURA) that solves the worst-case problem-to-go, imbued with an *ambiguity updating series* (LAUS) – a family of ambiguity sets of conditional distributions that vary with revealed uncertainties of previous time periods.

b) If the ambiguity updating series is consistent with conditional distributions from the original ambiguity set, then (i) (AURA) admits a time consistent optimal solution, (ii) this comes at the price of poorer in-sample performance that is no larger than the conservativeness of a parameter robust solution, and (iii) it coincides with the original DRO problem if and only if the latter were also time consistent (Theorem 1).

c) Under suitable Slater's conditions, strong duality reduces the robust counterpart of (AURA) to a deterministic dynamic program (Theorem 2).

d) When the support set of the uncertainty is continuous, (LAUS) can and often will contain all deterministic distributions, causing (AURA) to coincide with the parameter robust model, in a phenomenon we term "ambiguity collapse" (Theorem 3).

e) Ambiguity collapse is caused by distributions that become deterministic given some sample path of history, which we term "conditionally clairvoyant" (CC) (Theorem 4). We argue that CC distributions are unrealistic and guarding against them is needlessly conservative.

f) Removing CC distributions from the ambiguity set averts collapse with a strictly positive gap in objective values between (AURA) and the parameter robust model (Theorem 5).

g) To solve (AURA) on the new ambiguity set, we propose a discretized model, whose robust counterpart reduces to a deterministic dynamic program under strong duality (Theorem 6).

In Section 2, we propose our model termed *Ambiguity Updating Robust Architecture* (AURA) and its corresponding *ambiguity updating series*, illustrating its construction for a generalized moments ambiguity set. In Section 3, we analyze how ambiguity collapse occurs in the continuous setting and propose our fix. In Section 4, we illustrate our results on a stylized numerical experiment on a water management problem. All proofs are deferred to Appendix A.

Notation

Let $\mathcal{P}(\mathcal{I})$ be the set of probability measures with outcomes in $\mathcal{I}$. For a set $A \subset \mathcal{I}$ and $\mathbb{P} \in \mathcal{P}(\mathcal{I})$, we adopt the convention that when $\mathbb{P}[\boldsymbol{p} \in A, \cdot] = 0$, $\mathbb{P}[\cdot | \boldsymbol{p} \in A] = 0$. We use $\mathbb{P}_{\boldsymbol{p}}$ and $\mathbb{P}|_{\boldsymbol{p}}$ to denote the marginal and conditional distributions on random variable $\boldsymbol{p}$ and given $\boldsymbol{p}$ respectively. For a series of vectors $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T$ in time, the notation $\boldsymbol{x} := (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T)$ collects them into a single vector. Moreover, we use the notation $\boldsymbol{x}_{<t} := (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{t-1})$ and similarly $\boldsymbol{x}_{\leqslant t}$, and $\boldsymbol{x}_{>t} := (\boldsymbol{x}_{t+1}, \ldots, \boldsymbol{x}_T)$ and similarly $\boldsymbol{x}_{\geqslant t}$. For a matrix $\boldsymbol{G}$, $\boldsymbol{G}_{\leqslant t}$ denotes the first t columns of $\boldsymbol{G}$. Let $B_\epsilon(\boldsymbol{w})$ denote the open ball of radius ϵ centered at $\boldsymbol{w}$, with the metric implicitly assumed. $\text{conv}(A)$, $\text{cl}(A)$ and $\text{vol}(A)$ denote the convex hull, closure and volume (under the Lebesgue measure) of set A respectively.

2. Ambiguity Updating

Our multi-period setting is as follows: For times $t = 0, \ldots, T$, where T denotes the decision time-window, let $(\boldsymbol{x}_t \in \mathcal{X}_t \subseteq \mathbb{R}^{d_t})_{t=0}^T$ denote d_t-dimensional decisions, constrained within feasibility sets $\mathcal{X}_t$ that may be a function of uncertainties. At time $t = 0$, decision $\boldsymbol{x}_0$ is made without further knowledge, incurring cost $c_0(\boldsymbol{x}_0)$. Subsequently, at the start of every time period $t = 1, \ldots, T$, some u_t-dimensional uncertainty $\boldsymbol{w}_t \in \mathcal{W}_t \subseteq \mathbb{R}^{u_t}$ is revealed, then decision $\boldsymbol{x}_t$ is made as a function of manifested uncertainties $\boldsymbol{w}_{\leqslant t}$, incurring cost $c_t(\boldsymbol{x}_t; \boldsymbol{x}_{<t}, \boldsymbol{w}_{\leqslant t})$. For notational convenience, we suppress dependence of c_t on past decisions $\boldsymbol{x}_{<t}$ and uncertainty $\boldsymbol{w}_{\leqslant t}$. Let $\mathbb{P}$ denote the joint distribution of uncertainties $\boldsymbol{w}$ with joint support represented by $\mathcal{W} \subseteq \bigtimes_{t=1}^T \mathcal{W}_t$. Let

$$\hat{\mathcal{W}}_{\leqslant t} := \left\{ \boldsymbol{w}_{\leqslant t} \in \bigtimes_{\tau=1}^{t} \mathcal{W}_\tau \;\middle|\; \exists \boldsymbol{w}_{>t} \in \bigtimes_{\tau=t+1}^{T} \mathcal{W}_\tau, \text{ s.t. } (\boldsymbol{w}_{\leqslant t}, \boldsymbol{w}_{>t}) \in \mathcal{W} \right\} \quad \text{and}$$
$$\hat{\mathcal{W}}_t | \boldsymbol{w}_{<t} := \left\{ \boldsymbol{w}_t \in \mathcal{W}_t \;\middle|\; (\boldsymbol{w}_{<t}, \boldsymbol{w}_t) \in \hat{\mathcal{W}}_{\leqslant t} \right\},$$

denote the set of permissible sample paths (termed "histories") up to time t and the set of permissible realizations at time t for a given history $\boldsymbol{w}_{<t}$. We drop $\boldsymbol{w}_{<t}$ from $\hat{\mathcal{W}}_t | \boldsymbol{w}_{<t}$ whenever the context is clear.

DEFINITION 1 (NON-ANTICIPATORY POLICY). The solution $\boldsymbol{x}$ is *non-anticipatory* if for all time periods $t = 1, \ldots, T$, decision $\boldsymbol{x}_t := \boldsymbol{x}_t(\boldsymbol{w}_{\leqslant t})$ is a function of the history only up to time t, *i.e.*

$$\boldsymbol{x}_t(\boldsymbol{w}_{\leqslant t}, \boldsymbol{w}_{>t}) = \boldsymbol{x}_t(\boldsymbol{w}_{\leqslant t}, \tilde{\boldsymbol{w}}_{>t}), \qquad \forall \boldsymbol{w}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t},\ \forall \boldsymbol{w}_{>t}, \tilde{\boldsymbol{w}}_{>t} : (\boldsymbol{w}_{\leqslant t}, \boldsymbol{w}_{>t}), (\boldsymbol{w}_{\leqslant t}, \tilde{\boldsymbol{w}}_{>t}) \in \mathcal{W}.$$

Thus, we represent the space of non-anticipative and feasible decision policies as

$$\Pi := \left\{ \boldsymbol{x}(\cdot) \,\middle|\, \boldsymbol{x}_t : \hat{\mathcal{W}}_{\leqslant t} \to \mathcal{X}_t \text{ and } \boldsymbol{x}_t \text{ is non-anticipative, } \forall t = 0, \ldots, T \right\}.$$

In a similar fashion, we let Π_t and $\Pi_{t:T}$ represent the space of non-anticipative and feasible decision policies for $\boldsymbol{x}_t$ and $\boldsymbol{x}_{\geqslant t}$ respectively.

ASSUMPTION 1. *The support set $\mathcal{W}$ and all sets of permissible realizations $\hat{\mathcal{W}}_t$ are convex and compact for all $t = 1, \ldots, T$.*

ASSUMPTION 2 **(Relatively complete recourse)**. *We assume* relatively complete recourse, i.e., *for all times $t = 1, \ldots, T$, feasible decisions $\boldsymbol{x}_{<t}$ and histories $\boldsymbol{w}_{\leqslant t}$, there exists some feasible $\boldsymbol{x}_t \in \mathcal{X}_t$.*

The Assumptions above are commonly seen in the literature and simplify analysis of the theoretical content later.

We can write the nominal multi-period DRO model as follows:

$$\min_{\boldsymbol{x} \in \Pi} \quad c_0(\boldsymbol{x}_0) + \sup_{\mathbb{Q} \in \mathcal{U}} \mathbb{E}_{\mathbb{Q}} \left[\sum_{t=1}^{T} c_t\Big(\boldsymbol{x}_t(\boldsymbol{w}_{\leqslant t})\Big) \right] \tag{DRO}$$

for some ambiguity set $\mathcal{U}$ of distributions that is exogenously defined by the decision-maker. We refer to this ambiguity set $\mathcal{U}$ as the *parent* ambiguity set and (DRO) as the *parent* problem. For reference, we also state the (parameter) robust optimization formulation, that optimizes against the worst-case sample path:

$$\min_{\boldsymbol{x}(\cdot) \in \Pi} \quad c_0(\boldsymbol{x}_0) + \sup_{\boldsymbol{w} \in \mathcal{W}} \sum_{t=1}^{T} c_t\Big(\boldsymbol{x}_t(\boldsymbol{w}_{\leqslant t})\Big). \tag{RO}$$

2.1. Loss of time consistency in DRO

EXAMPLE 1. Consider the $T = 2$ period setting. Suppose the parent ambiguity set is:

$$\mathcal{U} := \left\{ \mathbb{Q} \in \mathcal{M}\left(\mathcal{W} := \{0,1\}^2, \mathbb{R}\right) \,\middle|\, \begin{array}{l} (w_1, w_2) \sim \mathbb{Q} : \\ \quad \mathbb{E}_{\mathbb{Q}}[w_1 + w_2] = \mu \\ \quad \mathbb{Q}\left[(w_1, w_2) \in \mathcal{W}\right] = 1 \end{array} \right\}.$$

Suppose now, uncertainty $w_1 = \hat{w}_1$ manifests and one wishes to solve the restricted problem on w_2, by capturing the conditional distributions of $w_2 \,|\, (w_1 = \hat{w}_1)$. This is non-trivial. The attempt

$$\mathcal{V}(\hat{w}_1; \mu) := \left\{ \mathbb{P} \in \mathcal{M}\left(\mathcal{W}_2 := \{0,1\}, \mathbb{R}\right) \,\middle|\, \begin{array}{l} w_2 \sim \mathbb{P} : \\ \quad \mathbb{E}_{\mathbb{P}}[\hat{w}_1 + w_2] = \mu \\ \quad \mathbb{P}[w_2 : (\hat{w}_1, w_2) \in \mathcal{W}] = 1 \end{array} \right\},$$

which simply substitutes w_1 with $\hat{w}_1$, is evidently wrong. To give a counterexample, suppose $1 < \mu \leqslant 2$. Here, $\mathcal{U} \neq \varnothing$, but if $\hat{w}_1 = 0$ is revealed, then $\mathcal{V}(\hat{w}_1; \mu) = \varnothing$, since $\mathbb{E}_{\mathbb{P}}[0 + w_2] \leqslant 1 < \mu$! In fact, we can prove the stronger claim: Even if we use $\mathcal{V}(\hat{w}_1; \nu)$ where $\nu \neq \mu$ and $\nu := \nu(\hat{w}_1) \in [0, 2]$ that is even allowed to depend on $\hat{w}_1$, there is no such function ν such that $\mathcal{V}(\hat{w}_1; \nu)$ contains all and exactly all conditional distributions of $w_2|(w_1 = \hat{w}_1)$. (Proof in Appendix A.)

The issue at hand is *time consistency*. We adopt the following definition that is consistent with Xin and Goldberg (2021):

Definition 2 (Time consistency). Let $\mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U})$ and $\mathscr{J}_{>t}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U})$ denote the sets of all possible conditional marginal distributions of $\boldsymbol{w}_{t+1}$ and $\boldsymbol{w}_{>t}$ respectively, given $\hat{\boldsymbol{w}}_{\leqslant t}$ that arise from some parent distribution in $\mathcal{U}$. For every time period t and history $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$,

i. Define the Bellman recursion:

$$\begin{aligned} F_T(\boldsymbol{x}_T; \hat{\boldsymbol{w}}) &:= c_T(\boldsymbol{x}_T), \\ F_t(\boldsymbol{x}_{\geqslant t}; \hat{\boldsymbol{w}}_{\leqslant t}) &:= c_t(\boldsymbol{x}_t) + \sup_{\mathbb{Q}_{t+1} \in \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U})} \mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_{t+1}}} \left[F_{t+1}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) \right]. \end{aligned} \tag{1}$$

We say that an optimal policy $\boldsymbol{x}^*(\cdot)$ of (DRO) is *type-I time consistent* if $\boldsymbol{x}^*(\cdot) \in \arg\min F_0(\boldsymbol{x})$.

ii. Define the distributional worst-case problem-to-go by

$$\begin{aligned} \tilde{F}_T(\boldsymbol{x}_{<T}, \hat{\boldsymbol{w}}_{\leqslant T}) &= \min_{\boldsymbol{x}_T \in \mathcal{X}_T} c_T(\boldsymbol{x}_T), \\ \tilde{F}_t(\boldsymbol{x}_{<t}, \hat{\boldsymbol{w}}_{\leqslant t}) &:= \min_{\boldsymbol{x}_{\geqslant t}(\cdot) \in \Pi_{t:T}} c_t(\boldsymbol{x}_t) + \sup_{\mathbb{Q} \in \mathscr{J}_{>t}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U})} \mathbb{E}_{\mathbb{Q}} \left[\sum_{t'=t+1}^{T} c_{t'}(\boldsymbol{x}_{t'}(\boldsymbol{w}_{\leqslant t'})) \right]. \end{aligned} \tag{2}$$

We say that an optimal policy $\boldsymbol{x}^*(\cdot)$ of (DRO) is *type-II time consistent* if, for every time period t and every history $\hat{\boldsymbol{w}}_{\leqslant t}$, its continuation policy $\boldsymbol{x}^*_{\geqslant t}(\hat{\boldsymbol{w}}_{\leqslant t}, \cdot)$ is optimal for $\tilde{F}_t(\boldsymbol{x}_{<t}, \hat{\boldsymbol{w}}_{\leqslant t})$.

In general, optimal solutions of (DRO) are not time consistent (Shapiro 2016). Our explanation, and what Example 1 shows, is because the geometry of the ambiguity set is not preserved under conditioning. Hence, solving (1) or (2) after the uncertainties $\boldsymbol{w}_{\leqslant t}$ have materialized often results in stronger performance than *faithfully* following the optimal decision rule in (DRO) for all times t. This is otherwise referred to as a *rolling horizon* implementation. As we know, Bellman recursion naturally produces time-consistent policies, whereas rolling horizon provides no guarantee of optimality. The above definitions therefore formalize time consistency from both Bellman recursion and rolling horizon perspectives.

Proposition 1 (**Shapiro and Xin (2020)**). *An optimal policy $\boldsymbol{x}^*(\cdot)$ of (DRO) is Type-I time consistent if and only if it is Type-II time consistent.*

REMARK 1. Bellman's Principle of Optimality works even when we sequentially take supremums within the expectation on each time period, while (DRO) takes the supremum for the worst-case distribution all at once. This mechanics is better illustrated in the proof of Theorem 1i. later.

Proposition 1 indicates that both definitions are equivalent when discussing about optimal policies. For convenience, we will adopt the definition of Type-I time consistency (1) since defining uncertainties to-go as a function of past observations is more natural than defining possible pasts conditional on given futures. Hereon, we will not differentiate which type of time consistency we are referring to for optimal solutions. Note that there is also a set of definitions for whether the problem (DRO) is time consistent or not – we are not concerned about that in this manuscript.

In practice, one cannot solve (1) in Definition 2 just yet, because it is unclear how to define the ambiguity set $\mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$ originating from the parent ambiguity set $\mathcal{U}$, given revealed uncertainty $\hat{\boldsymbol{w}}_{\leqslant t}$. Naïve attempts to update the constraint RHS's in the ambiguity set (which often is done in practice) can lead to disaster, as Example 1 illustrates. Avoiding this altogether, by directly prescribing information about $\boldsymbol{w}_t \,|\, \boldsymbol{w}_{<t}$ in $\mathcal{U}$ (such as in Hu et al. 2025), or by assuming temporal independence, can lead to over-specification of the ambiguity beyond what the data supports or incur estimation error. Instead, we directly tackle the problem of defining $\mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$, without assuming more distributional information than available in $\mathcal{U}$, and return to solve (DRO), but with knowledge of how ambiguity sets $\mathcal{U}_t(\hat{\boldsymbol{w}}_{\leqslant t})$ are updated in subsequent problems-to-go.

2.2. Ambiguity updating robust architecture

Motivated by the definition of (1), we proposed our model, the Ambiguity Updating Robust Architecture (AURA), in the Bellman recursion form:

$$
\begin{aligned}
\min_{\boldsymbol{x}(\cdot)\in\Pi} \quad & c_0(\boldsymbol{x}_0) + \sup_{\mathbb{Q}_{\boldsymbol{w}_1}\in\mathscr{J}_1} \mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_1}}\left[F_1(\boldsymbol{x}_{\geqslant 1};\boldsymbol{w}_1)\right] && \text{(AURA)}\\
\text{s.t.} \quad & F_t(\boldsymbol{x}_{\geqslant t};\hat{\boldsymbol{w}}_{<t},\boldsymbol{w}_t) := c_t(\boldsymbol{x}_t) + \sup_{\substack{\mathbb{Q}_{\boldsymbol{w}_{t+1}|(\hat{\boldsymbol{w}}_{<t},\boldsymbol{w}_t)}\\ \in\mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{<t},\boldsymbol{w}_t)}} \mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_{t+1}|(\hat{\boldsymbol{w}}_{<t},\boldsymbol{w}_t)}}\left[F_{t+1}(\boldsymbol{x}_{\geqslant t+1};(\hat{\boldsymbol{w}}_{<t},\boldsymbol{w}_t),\boldsymbol{w}_{t+1})\right], \\
& \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall t = 1,\ldots,T-1\\
& F_T(\boldsymbol{x}_T;\hat{\boldsymbol{w}}_{<T},\boldsymbol{w}_T) := c_T(\boldsymbol{x}_T)
\end{aligned}
$$

where $F_t : \Pi_{t:T}\times\hat{\mathcal{W}}_{\leqslant t}\to\mathbb{R}$ are the problems-to-go for optimizing decisions-to-go $\boldsymbol{x}_{\geqslant t}(\cdot)$ given prior uncertainty manifested $\hat{\boldsymbol{w}}_{<t}\in\hat{\mathcal{W}}_{<t}$ with uncertainty-to-go $\boldsymbol{w}_t\in\hat{\mathcal{W}}_t$. For notational convenience, we shorthand $\mathbb{Q}_{\boldsymbol{w}_{t+1}|\boldsymbol{w}_{\leqslant t}}$ as $\mathbb{Q}_{t+1}$.

Critically, we aim to define a family of ambiguity sets $\mathscr{U} := \{\mathscr{J}_1,\ldots,\mathscr{J}_T\}$, where for all $t = 1,\ldots,T$ and for all $\hat{\boldsymbol{w}}_{<t}\in\hat{\mathcal{W}}_{<t}$, then $\mathscr{J}_t(\hat{\boldsymbol{w}}_{<t})\subseteq\mathcal{M}(\hat{\mathcal{W}}_t,\mathbb{R})$ contain all conditional distributions of $\boldsymbol{w}_t|(\boldsymbol{w}_{<t}=\hat{\boldsymbol{w}}_{<t})$ from the parent ambiguity set $\mathcal{U}$. More specifically,

Definition 3 (Ambiguity updating series). Assuming axiom of choice, the sequence of ambiguity sets $\mathscr{U} := \{\mathcal{U}_1, \ldots, \mathcal{U}_T\}$ is an *ambiguity updating series*, if for all $t = 1, \ldots, T-1$ and $\boldsymbol{w}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$, the ambiguity set $\mathcal{U}_{t+1}(\boldsymbol{w}_{\leqslant t}) \subseteq \mathcal{P}(\hat{\mathcal{W}}_{t+1})$ contains only distributions defined on $\boldsymbol{w}_{t+1}$.

Given parent uncertainty set $\mathcal{U}$, we say that, with respect to $\mathcal{U}$, ambiguity updating series $\mathscr{U}$ is:

i. Inclusive: For all $\mathbb{P} \in \mathcal{U}$, and for all $t = 1, \ldots, T$, for all $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$ such that $\mathbb{P}[\boldsymbol{w}_t \in B_\epsilon(\hat{\boldsymbol{w}}_{\leqslant t})] > 0$, $\forall \epsilon > 0$, there exists some $\mathbb{Q}_{t+1} \in \mathcal{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$ such that $\mathbb{Q}_{t+1}[\cdot] = \mathbb{P}[\boldsymbol{w}_{t+1} = \cdot \,|\, \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}]$; and,

ii. Traceable: For all $t = 1, \ldots, T$, for all $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$, and for all $\mathbb{Q}_{t+1} \in \mathcal{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$, there exists some $\mathbb{P} \in \mathcal{U}$ such that $\mathbb{Q}_{t+1}[\cdot] = \mathbb{P}[\boldsymbol{w}_{t+1} = \cdot \,|\, \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}]$.

If ambiguity updating series $\mathscr{U}$ is both inclusive and traceable with respect to $\mathcal{U}$, we say it is *exact*.

The Definition demands that the ambiguity sets in the series indeed contain marginal distributions of the uncertainties at time t, conditional on past revealed uncertainties $\boldsymbol{w}_{\leqslant t}$. 'Inclusiveness' requires that no parent distribution that was contained in $\mathcal{U}$ was lost; while 'traceability' requires that the series does not contain distributions that cannot be retraced to some parent distribution in $\mathcal{U}$. Their combination mimics a 'weak equivalence' between parent $\mathcal{U}$ and series $\mathscr{U}$.

Theorem 1. *If $\mathscr{U}$ is an exact ambiguity updating series for parent ambiguity set $\mathcal{U}$, then*

i. (DRO) $\leqslant$ (AURA) $\leqslant$ (RO).

ii. There exists some optimal solution $\boldsymbol{x}^{AURA}$ for (AURA) *that is time consistent.*

iii. There exists some optimal solution to (DRO) *that is time consistent if and only if* (DRO) $=$ (AURA).

Result i. explains that (AURA) has to pay an additional premium, bounded above by the conservativeness of the parameter RO solution, because of the time consistency afforded in Result ii.. The latter is core because (AURA) can always recover time consistent solutions, as long as it is imbued with a suitable ambiguity updating series $\mathscr{U}$, whereas there is no guarantee that (DRO) generates an optimal solution that is time consistent. Result ii. holds essentially since (AURA) is equivalent to the Bellman recursion form (Proposition 1), thus it must be time consistent (Carpentier et al. 2012). Here, note that degeneracy is expected in (AURA), since there are feasible solutions defined on non-worst-case sample paths that do not threaten the optimality on the worst-case sample path. Should (DRO) generate a time consistent optimal solution, then Result iii. shows that (AURA) does no poorer. Since we will be able to solve (AURA) later, this provides a computational approach to determining the time consistency of (DRO) that did not previously exist.

Curiously, while 'inclusive' and 'traceable' conditions in Definition 3 seems to mimic sufficient and necessary conditions, a gap between (DRO) and (AURA) remains. To understand why, consider the reconstruction $\mathbb{Q}^*[\boldsymbol{w}] = \prod_{t=1}^{T} \mathbb{Q}^*_{\boldsymbol{w}_{\leqslant t}}$ from worst-case conditional distributions $\mathbb{Q}^*_{\hat{\boldsymbol{w}}_{\leqslant t}} \in \mathcal{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$. This need not lie in the parent set $\mathbb{Q}^* \notin \mathcal{U}$. In fact, in the proof of Result iii., time consistency occurs

if and only if worst-case distributions $\mathbb{Q}^*_{\hat{\boldsymbol{w}}_{\leqslant t}}$ are common across histories $\hat{\boldsymbol{w}}_{\leqslant t}$. Thus, (AURA) is protected against more distributions than (DRO). Hence, we view (AURA) as a kind of "completion" of (DRO) under conditioning.

2.3. Constructing an ambiguity updating series

To see how an ambiguity updating series can be constructed, consider the following structure:

ASSUMPTION 3. *The parent ambiguity set is a generalized moments ambiguity set*

$$\mathcal{U} := \left\{ \mathbb{Q} \in \mathcal{M}\left(\bigtimes_{t=1}^{T} \mathbb{R}^{u_t}, \mathbb{R} \right) \;\middle|\; \boldsymbol{w} \sim \mathbb{Q}: \quad \begin{array}{l} \mathbb{E}_{\mathbb{Q}}\left[\boldsymbol{g}(\boldsymbol{w})\right] = \boldsymbol{0}, \\ \mathbb{Q}\left[\boldsymbol{w} \in \mathcal{W}\right] = 1. \end{array} \right\}, \tag{3}$$

with support $\mathcal{W}$ *and constraints* $\boldsymbol{g}(\cdot): \mathcal{W} \to \mathbb{R}^b$.

For all times $t = 0, \ldots, T-1$ and histories $\hat{\boldsymbol{w}}_{\leqslant t}$, define the *lifted ambiguity updating series* (LAUS):

$$\begin{aligned} \mathscr{J}_{t+1}\left(\hat{\boldsymbol{w}}_{\leqslant t}\right) &:= \left\{ \mathbb{Q}_{t+1} \in \mathcal{M}\left(\mathbb{R}^{u_{t+1}}, \mathbb{R}\right) \;\middle|\; \begin{array}{l} \boldsymbol{w}_{t+1} \sim \mathbb{Q}_{t+1}: \\ \quad \exists \mathbb{P} \in \mathcal{U}: \mathbb{Q}_{t+1}[\cdot] = \mathbb{P}\left[\boldsymbol{w}_{t+1} = \cdot \,|\, \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right] \\ \quad \mathbb{Q}_{t+1}\left[\boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1} | \hat{\boldsymbol{w}}_{\leqslant t}\right] = 1 \end{array} \right\} \\ &:= \left\{ \mathbb{Q}_{t+1} \in \mathcal{M}\left(\mathbb{R}^{u_{t+1}}, \mathbb{R}\right) \;\middle|\; \begin{array}{l} \boldsymbol{w}_{t+1} \sim \mathbb{Q}_{t+1}, \exists \mathbb{P} \in \mathcal{M}\left(\bigtimes_{t=1}^{T} \mathbb{R}^{u_t}, \mathbb{R} \right), \boldsymbol{w} \sim \mathbb{P}: \\ \quad \mathbb{E}_{\mathbb{P}}\left[\boldsymbol{g}(\boldsymbol{w})\right] = \boldsymbol{0} \\ \quad \mathbb{P}\left[\boldsymbol{w} \in \mathcal{W}\right] = 1 \\ \quad \mathbb{P}[\boldsymbol{w}_t \in B_\epsilon(\hat{\boldsymbol{w}}_{\leqslant t})] > 0, \qquad \forall \epsilon > 0 \\ \quad \mathbb{Q}[\boldsymbol{w}_{t+1}] = \dfrac{\mathbb{P}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1} = \boldsymbol{w}_{t+1}\right]}{\mathbb{P}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right]} \end{array} \right\}. \end{aligned} \tag{4}$$

(LAUS)

(4) describes that $\mathscr{J}_{t+1}\left(\hat{\boldsymbol{w}}_{\leqslant t}\right)$ must contain all marginal distribution of $\mathbb{P} \in \mathcal{U}$ on $\boldsymbol{w}_{t+1}$, conditional on history $\hat{\boldsymbol{w}}_{\leqslant t}$. Notice that the distribution $\mathbb{P}[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}]$ (or essentially $\mathbb{P}[\boldsymbol{w}_t \in B_\epsilon(\hat{\boldsymbol{w}}_{\leqslant t})]$ in the continuous case) remains unobserved even if the manifestation $\hat{\boldsymbol{w}}_{\leqslant t}$ is observed – that is the whole point of distributional ambiguity! The key to construct an ambiguity updating series $\mathscr{U}$ is lifting, specifically $\mathbb{P} \in \mathcal{U}$ itself. Hereon, we specifically use the notation of $\mathscr{J}_{t+1}\left(\hat{\boldsymbol{w}}_{\leqslant t}\right)$ to refer to (LAUS).

We first note that if the parent ambiguity set $\mathcal{U}$ is convex and compact, then this transfers to (LAUS) under mild conditions.

PROPOSITION 2. *Suppose the parent ambiguity set* $\mathcal{U}$ *of the form* (3) *is convex and compact. Then for all times* $t = 1, \ldots, T$ *and histories* $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$, *the ambiguity set* $\mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$ *is convex.*

PROPOSITION 3 **(LAUS is exact)**. (LAUS) *is an exact ambiguity updating series for* (3). *Consequently, the objective value of* (AURA) *under* (LAUS) *is an upper bound for that of* (DRO) *under* (3), *and a time consistent optimal solution exists.*

(LAUS) involves a linear fractional relationship between $\mathbb{P}$ and $\mathbb{Q}$. By leveraging fractional programming techniques (Cooper and Charnes 1962), its robust counterpart can be computed:

LEMMA 1 **(Reformulation)**. *For all times $t = 0, \ldots, T-1$ and histories $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$, where $\mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}) \neq \varnothing$, the worst-case problem over the ambiguity set (LAUS) has upper bound,*

$$\sup_{\mathbb{Q}_{t+1} \in \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})} \mathbb{E}_{\mathbb{Q}_{t+1}} \left[F_{t+1}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) \right] \leqslant D_t(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}),$$

where $D_t(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t})$ is defined as the problem:

$$\inf_{\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant t}}, \theta^{\hat{\boldsymbol{w}}_{\leqslant t}} \text{ free}} \quad \theta^{\hat{\boldsymbol{w}}_{\leqslant t}} \tag{5}$$

$$\text{s.t.} \quad {\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant t}}}^{\top} \boldsymbol{g}(\boldsymbol{w}) + \theta^{\hat{\boldsymbol{w}}_{\leqslant t}} \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right] \tag{6}$$

$$\geqslant F_{t+1}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right], \qquad \forall \boldsymbol{w} \in \mathcal{W}$$

Prima facie, (AURA) is a distributionally robust stochastic dynamic program, which can be very difficult to solve. However, Lemma 1 reformulates each problem-to-go F_t into a deterministic form D_t. This allows a full reformulation of (AURA). For notational convenience, let $\hat{\boldsymbol{w}}_{\leqslant 0}$ refer to the initial state, thus any variable indexed by $\hat{\boldsymbol{w}}_{\leqslant 0}$ has no dependence on $\hat{\boldsymbol{w}}$.

ASSUMPTION 4 **(Slater's condition)**. *Define*

$$\mathscr{J}_t^{\boldsymbol{\alpha}}(\hat{\boldsymbol{w}}_{\leqslant t-1}) = \left\{ \mathbb{Q}_t \in \mathcal{M}\left(\hat{\mathcal{W}}_t, \mathbb{R}\right) \;\middle|\; \begin{array}{l} \boldsymbol{w}_t \sim \mathbb{Q}_t, \exists \mathbb{P} \in \mathcal{M}\left(\mathcal{W}, \mathbb{R}\right), \boldsymbol{w} \sim \mathbb{P}: \\ \quad \mathbb{E}_{\mathbb{P}}\left[\boldsymbol{g}(\boldsymbol{w})\right] = \boldsymbol{\alpha} \\ \quad \mathbb{P}\left[\boldsymbol{w}_{\leqslant t-1} \in B_\epsilon(\hat{\boldsymbol{w}}_{\leqslant t-1})\right] > 0, \ \forall \epsilon > 0 \\ \quad \mathbb{Q}[\boldsymbol{w}_t] = \dfrac{\mathbb{P}\left[\boldsymbol{w}_{\leqslant t-1} = \hat{\boldsymbol{w}}_{\leqslant t-1}, \boldsymbol{w}_t = \boldsymbol{w}_t\right]}{\mathbb{P}\left[\boldsymbol{w}_{\leqslant t-1} = \hat{\boldsymbol{w}}_{\leqslant t-1}\right]} \end{array} \right\},$$

where $\boldsymbol{\alpha}$ is in a small neighborhood of $\mathbf{0}$. There exists a tight and weakly closed subset of probability $\hat{\mathscr{J}}_t(\hat{\boldsymbol{w}}_{\leqslant t-1}) \subset \mathcal{M}\left(\hat{\mathcal{W}}_t, \mathbb{R}\right)$ such that $\mathscr{J}_t^{\boldsymbol{\alpha}}(\hat{\boldsymbol{w}}_{\leqslant t-1}) \subset \hat{\mathscr{J}}_t(\hat{\boldsymbol{w}}_{\leqslant t-1})$ for all $\boldsymbol{\alpha}$ close to zero. "Tight" means that $\forall \epsilon > 0$, there exists a compact set $\hat{\mathcal{W}}_t^\epsilon \subset \mathbb{R}^{u_t}$, such that $\mathbb{Q}_t(\hat{\mathcal{W}}_t^\epsilon) > 1 - \epsilon, \forall \mathbb{Q}_t \in \mathscr{J}_t$.

THEOREM 2. *Under Assumptions 3 and 4, for any fixed policy $\boldsymbol{x}(\cdot) \in \mathcal{X}$, (AURA) has strong dual:*

$$\sup_{\mathbb{Q}_{\boldsymbol{w}_1} \in \mathscr{J}_1} \quad \mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_1}}\left[F_1(\boldsymbol{x}_{\geqslant 1}; \boldsymbol{w}_1)\right] \tag{7}$$

$$= \min_{\boldsymbol{\lambda}, \boldsymbol{\theta} \text{ free}} \quad \theta^{\hat{\boldsymbol{w}}_{\leqslant 0}} \tag{AURA-RC}$$

$$\text{s.t.} \quad {\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant t}}}^{\top} \boldsymbol{g}(\boldsymbol{w}) + \theta^{\hat{\boldsymbol{w}}_{\leqslant t}} \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right] \geqslant \left(c_{t+1}(\boldsymbol{x}_{t+1}) + \theta^{(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})} \mathbf{1}[t < T-1]\right) \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right],$$

$$\forall t = 0, \ldots, T-1, \forall \boldsymbol{w} \in \mathcal{W}, \forall \hat{\boldsymbol{w}} \in \hat{\mathcal{W}}_t, \tag{8}$$

(AURA) reformulates into a deterministic dynamic program. Ignoring uncertainty and replacing expectations with nominal values already requires a dynamic program to solve in general, thus in some sense, akin to many robust optimization results, we obtained distributional robustness (and in our case, time consistency as well) 'for free'.

3. Ambiguity collapse in the continuous space

There is a disparity between two settings: If $\mathcal{W}$ is discrete, then all is theoretically fine and while (AURA-RC) may have a large number of constraints, there are also associated techniques to deal with these large-scale problems, such as constraint generation. The goal of our paper is not about the computational methods to solve dynamic programs; and we leave this discussion here.

If $\mathcal{W}$ is continuous, it is possible that (LAUS) may simplify, for all $t>0$ and for all $\hat{\boldsymbol{w}}_{\leqslant t}\in\hat{\mathcal{W}}_{\leqslant t}$, into $\left\{\mathbb{Q}\in\mathcal{M}\left(\mathbb{R}^{u_{t+1}},\mathbb{R}\right)\,\middle|\,\mathbb{Q}\left[\boldsymbol{w}_{t+1}\in\hat{\mathcal{W}}_{t+1}\right]=1\right\}$, the set of all distributions that are supported on the possible remaining trajectories $\hat{\mathcal{W}}_{t+1}$. To see how this occurs, consider the following illustration:

Example 2 (Ambiguity collapse). Consider the example where $\mathcal{W}_1=\mathcal{W}_2=[0,1]$ and w_1 does not restrict the allowed uncertainty of w_2, *i.e.* $\hat{\mathcal{W}}_2(\hat{w}_1)\equiv\mathcal{W}_2$ for all $\hat{w}_1\in[0,1]$. For simplicity, consider a parent ambiguity set in the form of (3) where $g(\boldsymbol{w})=(1,1)^\top\boldsymbol{w}-1.5$ is linear. Note that the curve $g(\boldsymbol{w})=0$ must intersect the square $[0,1]^2$, otherwise the parent ambiguity set is empty. For any $\hat{w}_1,w_2\in[0,1]$, consider the distribution:

$$\mathbb{P}[\boldsymbol{w}]=\begin{cases}\zeta & \text{when } \boldsymbol{w}=(\hat{w}_1,w_2)^\top\\ 1-\zeta & \text{when } \boldsymbol{w}=(1,1)^\top\\ 0 & \text{otherwise}\end{cases}\quad\text{or}\quad\mathbb{P}[\boldsymbol{w}]=\begin{cases}\zeta & \text{when } \boldsymbol{w}=(\hat{w}_1,w_2)^\top\\ 1-\zeta & \text{when } \boldsymbol{w}=(0,0)^\top\\ 0 & \text{otherwise}\end{cases},$$

depending on whether $g(\hat{w}_1,w_2)<0$ (former) or $g(\hat{w}_1,w_2)>0$ (latter). Regardless of $\hat{w}_1,w_2\in[0,1]$, there always exists some $\zeta\in[0,1]$ such that $\mathbb{E}_{\mathbb{P}}[g(\boldsymbol{w})]=0$ (see Figure 1). Indeed, pick $\zeta=\frac{1}{2(2-\hat{w}_1-w_2)}$ and $\zeta=\frac{3}{2(\hat{w}_1+w_2)}$ respectively (in both cases, $\zeta\in[0,1]$; for the former, $g(\hat{w}_1,w_2)<0$ implies $\hat{w}_1+w_2<1.5$, hence $\zeta\in[0,1]$; for the latter, $g(\hat{w}_1,w_2)>0$ implies $\hat{w}_1+w_2>1.5$, hence $\zeta\in[0,1]$). Under $\mathbb{P}$, w_2 occurs almost surely given $w_1=\hat{w}_1$, *i.e.* for every $w_2\in\mathcal{W}_2$, there exists some parent distribution $\mathbb{P}$ that yields outcome w_2 w.p. 1 given $w_1=\hat{w}_1$. Thus, all of the one-point distributions on each $w_2\in\mathcal{W}_2$ lies in $\mathscr{J}_2(\hat{w}_1)$. Since $\mathscr{J}_2$ is convex, $\mathscr{J}_2$ contains all distributions supported on $\mathcal{W}_2$.

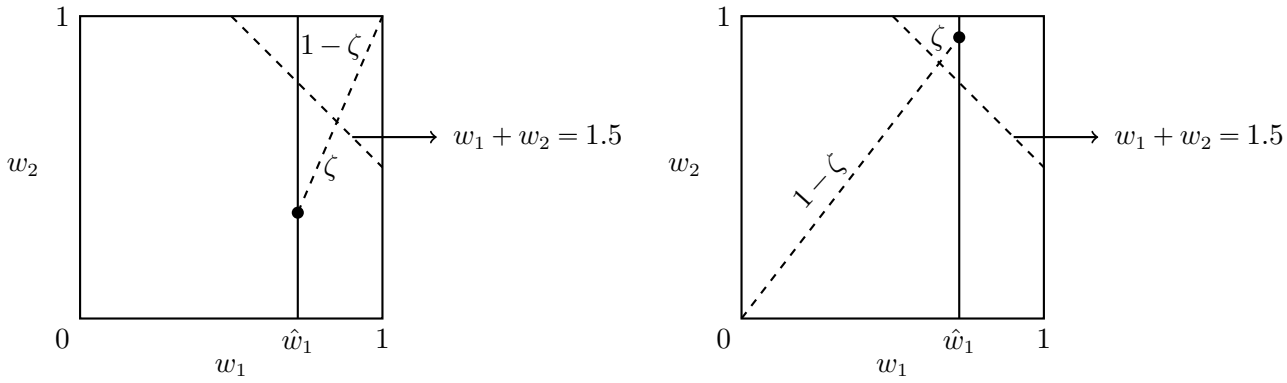


Figure 1 How $\mathbb{P}$ is constructed for given history $\hat{w}_1$, if the intended w_2 lies below (left) or above (right) $g(w)=0$.

The result of Example 2 turns out to be more general.

DEFINITION 4. For a given parent ambiguity set $\mathcal{U}$,

i. If $\mathbf{0}$ lies in the strict interior of the convex hull $\text{conv}\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\}$, then we say that $\mathcal{U}$ satisfies the *containment* property,

ii. If for every admissible realization $\hat{\boldsymbol{w}}_1 \in \hat{\mathcal{W}}_1$, the prescribed moment vector $\mathbf{0}$ is attainable by a probability distribution over its feasible continuations, *i.e.*,

$$\mathbf{0} \in \text{conv}\left\{\boldsymbol{g}(\boldsymbol{w}_1, \boldsymbol{w}_{>1}) : \boldsymbol{w}_1 = \hat{\boldsymbol{w}}_1, \boldsymbol{w} \in \mathcal{W}\right\}, \qquad \forall \hat{\boldsymbol{w}}_1 \in \hat{\mathcal{W}}_1, \tag{9}$$

then we say that $\mathcal{U}$ satisfies the *recoverable* property.

THEOREM 3 **(Price of time consistency)**. *Suppose $\mathcal{W}$ is a continuous space and parent ambiguity set $\mathcal{U}$ is of the generalized moments form (3). Let $\mathscr{U} := \{\mathscr{U}_1, \ldots, \mathscr{U}_T\}$ be some exact ambiguity updating series for $\mathcal{U}$. Denote the ambiguity updating series of (LAUS) by $\{\mathscr{J}_1, \ldots, \mathscr{J}_T\}$. Then:*

i. $\mathscr{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}) \supseteq \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$, for all $t > 0$ and almost surely for all $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$; and,

ii. Suppose $\mathcal{U}$ obeys containment. Then, under (LAUS), the problems-to-go of (AURA) and (RO) coincide at all times $t = 1, \ldots, T$ and for all histories $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$. Additionally, if $\mathcal{U}$ is recoverable, we have (AURA) = (RO).

Theorem 3ii. proves that, under mild conditions of containment and recoverability, (LAUS) will always degenerate to the set of all possible distributions, if the parent ambiguity set is a generalized moments ambiguity set of the form (3). This means that the price of solving (AURA) is the conservativeness of parameter robustness. We term this *"ambiguity collapse"*. Generalized moments ambiguity sets cover a very general group of ambiguity sets that we use to model practical problems today. Together with Theorem 1iii., this means that if these multi-period DRO models that we have been solving thus far are not at least as conservative as their parameter robust counterparts, then they are certainly not time consistent. Note that we do not show the non-existence of models more conservative than (DRO), but less conservative than (AURA), that can generate a time consistent solution. Nonetheless, Theorem 3i. guarantees that amongst exact ambiguity updating series, where time consistency is guaranteed, (LAUS) is already the smallest construction, hence least conservative. Later in the simulations, we will see that other methods for generating completions of the parent ambiguity set can be extremely large and needlessly conservative. Despite so, the price of pursuing time consistency through (LAUS) is already parameter robustness.

It is worthwhile to note that the discrete setting is spared this fate – there, in our numerical simulations later, we see that (i) the gaps between (DRO) < (AURA) < (RO) can be large; and, (ii) the gap between (AURA) and (RO) shrinks as we increase the number of points sampled to mimic a convergence of the discrete setting to the continuous case.

3.1. Sufficient conditions for ambiguity collapse and its avoidance

Before proceeding, note that if (DRO) = (RO), then by Theorem 1, there is no way for (AURA), which sits between them, to produce a different objective value. Hence, we impose

ASSUMPTION 5 **(Positive distance)**. *There are no single-point distributions amongst the worst-case distributions of* $\mathcal{U}$ *for the given optimization problem* (DRO)*. Consequently,* (DRO) $<$ (RO)*.*

ASSUMPTION 6 **(RO non-degeneracy)**. *For any given time* t *and history* $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_t$*, under the worst-case policy* $\boldsymbol{x}^{RO}$ *for the* (RO) *problem, its resulting problem-to-go is non-degenerate,* i.e.*, there exists some* $\boldsymbol{w}^*_{>t}$ *such that* $\hat{\boldsymbol{w}} := (\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}^*_{>t}) \in \mathcal{W}$ *and for all* $\tilde{\boldsymbol{w}} := (\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{>t})$ *such that*

$$\sum_{t=t+1}^{T} c_t\Big(\boldsymbol{x}_t^{RO}(\hat{\boldsymbol{w}}_{\leqslant t})\Big) < \sum_{t=t+1}^{T} c_t\Big(\boldsymbol{x}_t^{RO}(\tilde{\boldsymbol{w}}_{\leqslant t})\Big).$$

Assumption 6 is a further assumption to guarantee on the other side that (RO)'s worst-case sample path is non-degenerate. Thus, this means that the worst-case distribution in (DRO) if we do not impose any distributional constraints in parent ambiguity set $\mathcal{U}$ is uniquely some deterministic distribution. This simplifies the attempt to characterize ambiguity collapse.

Theorem 3 is a tragic result in the pursuit of time-consistent solutions. While the class of generalized moments parent ambiguity sets is large, it is key to differentiate that not all parent ambiguity sets might lead to ambiguity collapse under (AURA), as we will later see. This leads to an important realization that ambiguity collapse is inherently a property of the parent ambiguity set $\mathcal{U}$, and not of the (AURA) framework. We will also subsequently argue that there is something structurally 'poor' about generalized moments that do not work well, and is in fact unrealistic, in uncertainties with genuine temporal dependence.

To achieve this, we aim to obtain sufficient conditions for ambiguity collapse and for its avoidance, that are based on properties of the parent ambiguity set $\mathcal{U}$.

DEFINITION 5. For any time $t = 1, \ldots, T$, we say that a metric $d(\cdot, \cdot)$ defined on the weak topology of $\mathcal{P}(\hat{\mathcal{W}}_t)$ is *point distribution Lipschitz (PDL)* if there exists some constant $D > 0$ such that

$$d\left(\delta_{\boldsymbol{w}_t}, \delta_{\boldsymbol{w}'_t}\right) \leqslant D\|\boldsymbol{w}_t - \boldsymbol{w}'_t\|, \qquad \forall \boldsymbol{w}_t, \boldsymbol{w}'_t \in \hat{\mathcal{W}}_t. \tag{10}$$

Trivially, the p-Wasserstein distance for $p \geqslant 1$ is PDL, since $d_{W_p}\left(\delta_{\boldsymbol{w}_t}, \delta_{\boldsymbol{w}'_t}\right) \leqslant \|\boldsymbol{w}_t - \boldsymbol{w}'_t\|_p$, and all p-norms are Lipschitz-equivalent in Euclidean space.

DEFINITION 6. Let $t = 1, ..., T-1$ be a given time. We say that a history $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$ up till time t,

i. *Collapses* $\mathcal{U}$ (denoted $\hat{\boldsymbol{w}}_{\leqslant t} \downarrow\downarrow \mathcal{U}$) if the set of conditional distributions $\mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U})$ is exactly the set of all distributions supported on $\hat{\mathcal{W}}_{t+1}$.

ii. *Buttresses* $\mathcal{U}$ (denoted $\hat{\boldsymbol{w}}_{\leqslant t} \uparrow\uparrow \mathcal{U}$) if there exists some $r := r(\hat{\boldsymbol{w}}_{\leqslant t}) > 0$ (uniform over $\boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1}$), such that the set of conditional distributions $\mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U})$ contains no neighborhood $B_r(\delta_{\boldsymbol{w}_{t+1}})$ of any deterministic distribution on the support $\hat{\mathcal{W}}_{t+1}$, *i.e.*,

$$B_r(\delta_{\boldsymbol{w}_{t+1}}) \cap \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U}) = \varnothing, \qquad \forall \boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1},$$

where $B_r(\delta_{\boldsymbol{w}_{t+1}}) := \{\mathbb{Q} \in \mathcal{P}(\hat{\mathcal{W}}_{t+1}) : d(\mathbb{Q}, \delta_{\boldsymbol{w}_{t+1}}) < r\}$ is the ball of radius r about the point distribution $\delta_{\boldsymbol{w}_{t+1}}$, for any given PDL metric $d(\cdot, \cdot)$.

Furthermore, we say that the time t is a *collapse* (resp. *buttress*) *time* of $\mathcal{U}$ if for all $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$, $\hat{\boldsymbol{w}}_{\leqslant t} \downarrow\downarrow \mathcal{U}$ (resp. $\hat{\boldsymbol{w}}_{\leqslant t} \uparrow\uparrow \mathcal{U}$). We collect all collapse and buttress times in the sets $\mathscr{C}(\mathcal{U})$ and $\mathscr{B}(\mathcal{U})$ respectively. Moreover, if $\mathscr{C}(\mathcal{U}) = \{1, \ldots, T\}$ contains all times, we say $\mathcal{U}$ is *perfectly collapsible*.

Definition 6i. formalizes what it means for ambiguity collapse to occur at a particular time t and on a particular sample path with revealed history $\hat{\boldsymbol{w}}_{\leqslant t}$. If every possible history $\hat{\boldsymbol{w}}_{\leqslant t}$ up till time t can collapse parent ambiguity set $\mathcal{U}$ (as we have seen in Example 2), then a global collapse has occurred and we record that time t as a collapse time. The following Corollary indicates that our definition of collapse is consistent with ambiguity collapse seen earlier in Theorem 3.

COROLLARY 1. *If the generalized moments parent ambiguity sets $\mathcal{U}$ satisfies containment and its constraints $\boldsymbol{g}(\cdot)$ are continuous and finite, it is perfectly collapsible.*

Definition 6ii. describes the opposite scenario where there is an absence of any deterministic distributions in the set of conditional distributions $\mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U})$. As we shall soon see, this avoids ambiguity collapse. The more technical point of note here is that, without stronger assumptions, the exclusion of these deterministic distributions alone is not enough, since one may still arbitrarily closely approach these deterministic distributions should (i) a ball around them not be excluded, and (ii) there is no *uniform* minimal finite radius keeping the set of conditional distributions away from these deterministic distributions.

THEOREM 4 **(Conditions for collapse)**. *If any given time $t \geqslant 1$ obeys for all $\tau \geqslant t, \tau \in \mathscr{C}(\mathcal{U})$, then the problems-to-go of* (AURA) *and* (RO) *at time t given any history $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_t$ coincide. Additionally, if $\mathcal{U}$ is perfectly collapsible and satisfies the recoverable property, then* (AURA)=(RO).

The key conclusion from Theorem 4 is that distributions that abruptly become deterministic given some history or sample path, which we term *"conditionally clairvoyant (CC) distributions"*, are the culprits behind ambiguity collapse. In a sense, CC distributions are highly unlikely in reality – uncertainties never ever completely vanish after some time, even without the intervention of decisions. Hence, to plan against such distributions is to be needlessly conservative in practice. What is interesting is that many common ambiguity sets, such as the generalized moments ambiguity

set of form (3), do not exclude CC distributions by construct, thus leading to ambiguity collapse. Critically, CC distributions are also those that play a role in the time consistency of (DRO), since they cause (AURA) = (RO), yet without considering conditioning, we retain (DRO) < (AURA).

The concept of buttressing histories provides a way forward to address this issue – simply exclude these CC or 'near'-CC distributions that would land in this ball around a deterministic distribution. This can be done for any given parent ambiguity set and we imminently present a way to do this.

3.2. Avoiding collapse through removal of CC distributions

A possible approach to removing CC distributions is through requiring density *continuity*.

Definition 7 (CVD-regularized ambiguity). For a given parent ambiguity set $\mathcal{U}$, with generalized moments, define its *continuously varying distributed (CVD)-regularized ambiguity set* with given parameter $L > 0$ as

$$\mathcal{U}^*(L) := \left\{ \mathbb{P} \in \mathcal{M}\left(\bigtimes_{t=1}^{T} \mathbb{R}^{u_t}, \mathbb{R}\right) \;\middle|\; \boldsymbol{w} \sim \mathbb{P}: \quad \begin{array}{l} \mathbb{E}_{\mathbb{P}}\left[\boldsymbol{g}(\boldsymbol{w})\right] = \mathbf{0}, \\ \left\|p_{\mathbb{P}}(\boldsymbol{w}) - p_{\mathbb{P}}(\boldsymbol{w}')\right\| \leqslant L\left\|\boldsymbol{w} - \boldsymbol{w}'\right\|, \quad \forall \boldsymbol{w}, \boldsymbol{w}' \in \mathcal{W}, \\ \mathbb{P}\left[\boldsymbol{w} \in \mathcal{W}\right] = 1. \end{array} \right\}, \tag{11}$$

where $p_{\mathbb{P}}$ is the density function of $\mathbb{P}$ with respect to the Lebesgue measure.

The critical addition is the (Lipschitz) continuity requirement that prohibits the distribution from varying too wildly across nearby uncertainties. By definition, since CC distributions become deterministic given some history, therefore, about that history, their density cannot possibly be Lipschitz continuous. We will imminently prove this mathematically. We have termed this CVD-regularization, as we are pruning the parent ambiguity set $\mathcal{U}$, so that it is well-behaved and devoid of pathological exceptions, much like the concept of regularization in machine learning.

An alternative explanation for the continuity requirement is as follows. Distributions that vary wildly within a small neighbourhood of points describe immense sensitivity, and therefore instability, of the likelihood to specific uncertainties. These unstable distributions decrease the usefulness of the propagation of information over time – a slight measurement error in the uncertainty can lead to vastly different outcomes, which would ultimately require the decision-maker to be extremely conservative despite knowing the present manifestation of the uncertainty. This interpretation is consistent with the broader literature of how conditional distributions are handled in uncertain contexts (for example, Esteban-Pérez and Morales 2022, who employed the notion of *trimmings* to avert the issue of degeneracy of the event $\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}$).

We show that CVD-regularization fixes the problem of collapse, *i.e.*, for a well-chosen Lipschitz parameter L, the regularized $\mathcal{U}^*(L)$ has a non-empty set of buttress times $\mathscr{B}(\mathcal{U}^*(L)) \neq \varnothing$, and this leads to a positive bound between the difference in objective values of (AURA) and (RO).

DEFINITION 8. For given $\boldsymbol{\varepsilon} \in \mathbb{R}^T$ and $\boldsymbol{\eta} \in (0,1)^T$, we say that the support set $\mathcal{W}$ is $(\boldsymbol{\varepsilon}, \boldsymbol{\eta})$-dispersed, if for all times $t = 1, \ldots, T$,

$$\sup_{\substack{\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t} \\ \boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1}}} \frac{\operatorname{vol}\left(\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1}) \cap \hat{\mathcal{W}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})\right)}{\operatorname{vol}\left(\hat{\mathcal{W}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})\right)} < 1 - \eta_t. \tag{12}$$

This is a technical definition to ensure that set of permissible realizations $\hat{\mathcal{W}}_{t+1}$ does not concentrate only at any particular point $\boldsymbol{w}_{t+1}$.

THEOREM 5. *Suppose $\mathcal{W}$ is $(\boldsymbol{\varepsilon}, \boldsymbol{\eta})$-dispersed. Define*

$$\mathcal{D}_t := \mathcal{D}_t(\varepsilon_t, \eta_t) := \left\{ \mathbb{P} \in \mathcal{U} \;\middle|\; \begin{array}{l} \boldsymbol{w} \sim \mathbb{P}, \exists \hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}, \exists \boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1} : \\ \quad \mathbb{P}[\boldsymbol{w}_{\leqslant t} \in \mathcal{B}_\epsilon(\hat{\boldsymbol{w}}_{\leqslant t})] > 0, \qquad \forall \epsilon > 0 \\ \quad \mathbb{P}\left[\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1}) \mid \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right] \geqslant 1 - \eta_t \end{array} \right\} \qquad \textit{and}$$

$$L_t(\varepsilon_t, \eta_t) := \inf_{\mathbb{P} \in \mathcal{D}_t(\varepsilon_t, \eta_t)} \sup_{\substack{\boldsymbol{w}, \boldsymbol{w}' \in \mathcal{W}, \\ \boldsymbol{w} \neq \boldsymbol{w}'}} \frac{|p_{\mathbb{P}}(\boldsymbol{w}) - p_{\mathbb{P}}(\boldsymbol{w}')|}{\|\boldsymbol{w} - \boldsymbol{w}'\|}.$$

Then,

i. $L_t(\varepsilon_t, \eta_t) > 0$ for all times t; and,

ii. For all times t where $L_t(\varepsilon_t, \eta_t) > L$, for all of its histories $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$,

a) $\hat{\boldsymbol{w}}_{\leqslant t} \uparrow\uparrow \mathcal{U}^*(L)$, hence $t \in \mathscr{B}(\mathcal{U}^*(L))$ is a buttress time; and,

b) The objective gap between the problems-to-go of (AURA) and (RO) is at least

$$F^{RO}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}^*_{t+1}) - \sup_{\mathbb{Q} \in \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U}^*(L))} \mathbb{E}_{\mathbb{Q}}[F^{AURA}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})] \geqslant \eta_t \Delta_t(\hat{\boldsymbol{w}}_{\leqslant t}, \varepsilon_t) > 0, \tag{13}$$

where

$$\Delta_t(\hat{\boldsymbol{w}}_{\leqslant t}, \varepsilon_t) := F^{RO}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}^*_{t+1}) - \max_{\boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1} \setminus \mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})} F^{RO}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}),$$

is the gap of (RO) excluding worst-case continuation $\boldsymbol{w}^*_{t+1}$ and its neighbourhood. Moreover, $\Delta_t(\hat{\boldsymbol{w}}_{\leqslant t}, \varepsilon_t) > 0$.

Consequently, if L is chosen such that $L < \max_t L_t(\varepsilon_t, \eta_t)$, then $\mathscr{B}(\mathcal{U}^*(L)) \neq \varnothing$, and (AURA) $<$ (RO). If L is chosen such that $L < \min_t L_t(\varepsilon_t, \eta_t)$, then all problems-to-go of (AURA) achieve strictly smaller objective than that of (RO) at any time t and any history $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$.

Theorem 5 shows that CVD-regularization can indeed avoid ambiguity collapse. This can be done in a "weak" fashion by choosing $L < \max_t L_t(\epsilon_t, \eta_t)$ and only guaranteeing that some time period avoids collapse; or in a "strong" fashion by choosing $L < \min_t L_t(\epsilon_t, \eta_t)$, where collapse is avoided at every time and under every possible history realized, with the certified gap of at least $\eta_t \Delta_t$. Note that we will not lose time consistency, because we will still apply (AURA) upon $\mathcal{U}^*(L)$, treating it like a new parent ambiguity set, thereby reaping the guarantees from Theorem 1ii..

3.3. Computational approach, reformulation and tractability

Solving (DRO) with (11) as its ambiguity set is non-trivial, much less (AURA). To avert this, we propose a discretized approximation for the Lipschitz constraint in (11):

$$\left|\int_{\bar{B}_k} d\mathbb{P}(\boldsymbol{w}) - \int_{\underline{B}_k} d\mathbb{P}(\boldsymbol{w})\right| \leqslant L \cdot D(\bar{B}_k, \underline{B}_k), \quad \forall k = 1, \ldots, K, \tag{14}$$

Here $((\bar{B}_k, \underline{B}_k)_{k=1}^K)$ are pairs of bounded sets used to discretize the uncertainty space $\mathcal{W}$. Then $D(\cdot,\cdot)$ defines a distance measure between two sets. Notice that if the size of the sets shrinks to 0, then we will recover $\|p(\boldsymbol{w}) - p(\boldsymbol{w}')\| \leqslant L\|\boldsymbol{w} - \boldsymbol{w}'\|$. For convenience, we choose the following distance metric:

$$D(\bar{B}, \underline{B}) := W_1(U[\bar{B}], U[\underline{B}]),$$

where $W_1(\cdot,\cdot)$ is the 1-Wasserstein metric between two distributions and $U[B]$ is the uniform distribution over the set B. In Appendix A.12, we present the result that this indeed defines a pseudo metric. This is deliberately chosen as opposed to the more usual $D(\bar{B}_k, \underline{B}_k) = \inf_{\boldsymbol{x}\in\bar{B}, \boldsymbol{y}\in\underline{B}} ||\boldsymbol{x} - \boldsymbol{y}||$, so as to ensure that even if $\bar{B} \cap \underline{B} \neq \varnothing$, $D(\bar{B}_k, \underline{B}_k)$ can still be non-zero. With this approximation, we formally propose the discretized CVD-regularized ambiguity set:

$$\mathcal{U}^{\text{CVD}}(L) := \left\{ \mathbb{P} \in \mathcal{M}\left(\bigtimes_{t=1}^{T} \mathbb{R}^{u_t}, \mathbb{R}\right) \;\middle|\; \begin{array}{l} \boldsymbol{w} \sim \mathbb{P}: \\ \quad \mathbb{E}_{\mathbb{P}}[\boldsymbol{g}(\boldsymbol{w})] = \boldsymbol{0}, \\ \quad \left|\int_{\bar{B}_k} \mathbb{P}(d\boldsymbol{w}) - \int_{\underline{B}_k} \mathbb{P}(d\boldsymbol{w})\right| \leqslant L \cdot D(\bar{B}_k, \underline{B}_k), \quad k = 1, \ldots, K, \\ \quad \mathbb{P}[\boldsymbol{w} \in \mathcal{W}] = 1. \end{array} \right\}. \tag{CVD}$$

Proposition 4. *Suppose the discretized CVD-regularized (DRO) ambiguity set obeys Slater's condition (Assumption 7), then for any fixed policy $\boldsymbol{x}(\cdot) \in \mathcal{X}$, it has the strong dual:*

$$\begin{aligned} &\sup_{\mathbb{P}\in\mathcal{U}^{CVD}(L)} \quad \mathbb{E}_{\mathbb{P}}\left[\sum_{t=1}^{T} c_t\left(\boldsymbol{x}_t(\boldsymbol{w}_{\leqslant t})\right)\right] \\ = &\min_{\substack{\bar{\boldsymbol{\rho}},\underline{\boldsymbol{\rho}}\geqslant 0 \\ \theta,\boldsymbol{\lambda}\ free}} \quad \theta + \sum_{k=1}^{K} (\bar{\rho}_k + \underline{\rho}_k) \cdot L \cdot D(\bar{B}_k, \underline{B}_k) \\ &\quad s.t. \quad \boldsymbol{\lambda}^\top \boldsymbol{g}(\boldsymbol{w}) + \theta + \sum_{k=1}^{K} \left(\mathbf{1}[\boldsymbol{w} \in \bar{B}_k] - \mathbf{1}[\boldsymbol{w} \in \underline{B}_k]\right)(\bar{\rho}_k - \underline{\rho}_k) \geqslant \sum_{t=1}^{T} c_t\left(\boldsymbol{x}_t(\boldsymbol{w}_{\leqslant t})\right), \qquad \forall \boldsymbol{w} \in \mathcal{W} \end{aligned}$$

where $\bar{\boldsymbol{\rho}}, \underline{\boldsymbol{\rho}}$ are the dual variables corresponding to the CVD constraints.

Under (CVD) as the parent ambiguity set, both (DRO) and (AURA) can be reformulated and the latter's dual possesses a similar structure as before in Theorem 2. For brevity, we present the specific Slater's conditions for both models in Appendix A.13.

THEOREM 6. *Under Assumption 7, for any fixed policy* $\boldsymbol{x}(\cdot) \in \mathcal{X}$, (AURA) *under* (CVD) *has the strong dual:*

$$
\begin{aligned}
&\sup_{\mathbb{Q}_{\boldsymbol{w}_1} \in \mathcal{U}_1} && \mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_1}}\left[F_1(\boldsymbol{x}_{\geqslant 1}; \boldsymbol{w}_1)\right] \\
=\ &\min_{\substack{\bar{\boldsymbol{\rho}}, \underline{\boldsymbol{\rho}} \geqslant 0 \\ \boldsymbol{\lambda}, \boldsymbol{\theta}, \eta \text{ free}}} && \theta^{\hat{\boldsymbol{w}}_{\leqslant 0}} && \text{(Cont-RC)} \\
&\quad s.t. && \boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant 0}\top} \boldsymbol{g}(\boldsymbol{w}) + \theta^{\hat{\boldsymbol{w}}_{\leqslant 0}} \geqslant c_1(\boldsymbol{x}_1) + \theta^{(\hat{\boldsymbol{w}}_{\leqslant 0}, \boldsymbol{w}_1)}, && \forall \boldsymbol{w} \in \mathcal{W} \\
& && \boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant t}\top} \boldsymbol{g}(\boldsymbol{w}) + \eta^{\hat{\boldsymbol{w}}_{\leqslant t}} + \sum_{k=1}^{K} \left(\mathbf{1}\left[\boldsymbol{w} \in \bar{B}_k\right] - \mathbf{1}\left[\boldsymbol{w} \in \underline{B}_k\right]\right)\left(\bar{\rho}_k^{\hat{\boldsymbol{w}}_{\leqslant t}} - \underline{\rho}_k^{\hat{\boldsymbol{w}}_{\leqslant t}}\right) \\
& && \quad + \theta^{\hat{\boldsymbol{w}}_{\leqslant t}} \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right] \geqslant \left(c_{t+1}(\boldsymbol{x}_{t+1}) + \theta^{(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})} \mathbf{1}[t < T-1]\right) \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right], \\
& && && \forall t = 1, \ldots, T-1, \forall \boldsymbol{w} \in \mathcal{W}, \forall \hat{\boldsymbol{w}} \in \hat{\mathcal{W}}_t \\
& && -\eta^{\hat{\boldsymbol{w}}_{\leqslant t}} - \sum_{k=1}^{K} (\bar{\rho}_k + \underline{\rho}_k) \cdot L \cdot D(\bar{B}_k, \underline{B}_k) \geqslant 0, && \forall t = 1, ..., T-1, \forall \hat{\boldsymbol{w}} \in \hat{\mathcal{W}}_t.
\end{aligned}
$$

where $\bar{\boldsymbol{\rho}}, \underline{\boldsymbol{\rho}}$ *are the dual variables corresponding to the CVD constraints.*

Theorem 6 states that, like with Theorem 2, (AURA) under (CVD) has a robust counterpart that is a deterministic dynamic program. The proof mirrors that of the latter and is omitted, with the intermediate one-stage reformulation, akin to Lemma 1, presented in Appendix A.13.

4. Numerical Experiments

We demonstrate the proposed models and derived results using the example of multi-period water management, where rainfall and demand manifests at every time period and is related across time. Here, we shall present two settings, where the uncertainty is discrete or continuous.

4.1. Water management problem

In the water management problem, the decision-maker optimizes production and transfer of water between water bodies and water types, namely rainwater and potable water (that is safe for drinking), to satisfy water demand and minimize the chance of floods (the volume of water exceeding the capacity of any water body). The system contains different water bodies and production facilities represented by nodes and a network of pipes adjoining them represented by edges. These include:

1. **Water bodies**: Nodes v that hold a certain amount of water (of different types); its state at the end of any time period t is represented as s_v^t. They have fixed lower and upper capacities ι_v and κ_v respectively. A flooding event occurs whenever the state of the water body at the end of the time period exceeds the upper capacity $s_v^t > \kappa_v$; while a drought event occurs whenever the state falls below the lower capacity $s_v^t < \iota_v$. For all water bodies, note that it is possible for the throughput of water to exceed the actual capacity of the water body without resulting in a flooding event.

(a) Reservoirs: Denoted by index $r \in \mathcal{R}$, the set of all reservoirs, these are water bodies that hold s_r^t amount of rainwater that is untreated and unsafe for drinking, and receive an uncertain amount of rain w_r^t at every time period t.

(b) Tanks: Denoted by index $h \in \mathcal{H}$, the set of all tanks, these are water bodies that hold s_h^t amount of potable water and are used to serve an uncertain amount of demand w_h^t for potable water at every time period t.

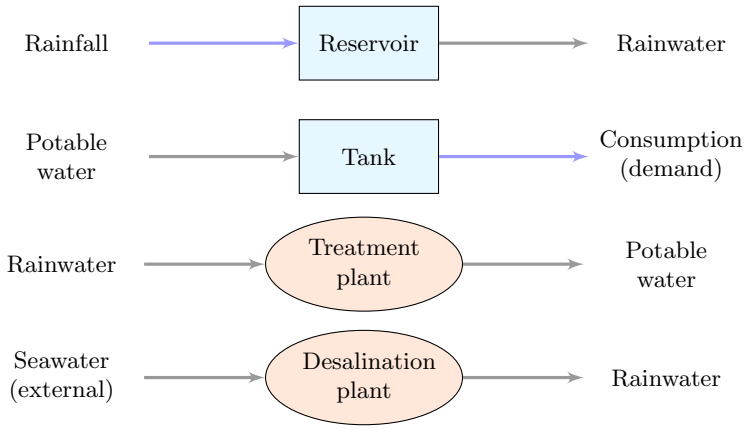
2. **Production facilities**: Nodes v that do not have any storage capacity (thus, have no associated state), but instead transform water from a water body into another type and released into another water body. They all have maximum production capacities τ_v. The cost of production c_v^t is modeled as throughput, *i.e.* on the cost of flows in and out of these facilities.

(a) Treatment plants: Denoted by index $p \in \mathcal{P}$, the set of all treatment plants, these are plants that transform a quantity of rainwater $x_{r,p}^t$ from a reservoir r into an equivalent quantity of potable water $x_{p,h}^t = x_{r,p}^t$ to a tank h at time period t, that is to be decided.

(b) Desalination plants: Denoted by index $d \in \mathcal{D}$, the set of all desalination plants, these are plants that transform a quantity of seawater $x_{0,d}^t$ from the external node $v = 0$ into an equivalent quantity of rainwater $x_{d,r}^t = x_{0,d}^t$ to a reservoir r at time period t, that is to be decided. It is customary to dilute desalinated water with rainwater in practice.

3. **External node**: We model a sink/source node for convenience and label it vertex $v = 0$.

Figure 2 Water bodies (boxes; gray arrows represent decisions and blue arrows represent uncertainties) and production facilities (ovals; gray arrows represent decisions).



Nodes are connected via a series of pipes, which could be bi- or uni-directional, represented as directed edges $(m, n) \in \mathcal{E}$ from node m to node n. Some edges represent uncertainties $\boldsymbol{w}$, such as rainfall and water consumption (demand), while the rest represent decision variables $x_{m,n}^t$, flows from node m to node n at time t. Floods or droughts can occur, and the system is not a closed

system. However, to guarantee complete recourse, we allow all water bodies to flow to and from the sink/source node $v = 0$, with infinite capacity and at very high costs.

4.1.1. Sequence of events. Our planning horizon has T periods, denoted $\mathcal{T} = \{1, \ldots, T\}$. Each period has two stages. At the start of period t, water bodies v possess a state s_v^{t-1} of water. Before uncertainties in that period $\boldsymbol{w}_t$ materialize, the decision-maker makes here-and-now decisions on flows, denoted $x_{m,n}^{t,\leftharpoonup}$, from node m to node n at the period t. Let $\mathcal{E}^{\leftharpoonup} \subseteq \mathcal{E}$ denote the set of all edges involved in here-and-now (first-stage) decisions. These decisions have upper capacity $K_{m,n}^{\leftharpoonup}$ and unit cost $c_{m,n}^{t,\leftharpoonup}$.

1. Production: Flows incident to a production facility are decided, $x_{m,v}^{t,\leftharpoonup}$ and $x_{v,n}^{t,\leftharpoonup}$ for all $v \in \mathcal{P} \cup \mathcal{D}$ and $(m,v),(v,n) \in \mathcal{E}$. Production is committed ahead of the uncertainties as it takes time.
2. Transfer: Flows between reservoirs and between reservoir and the source/sink node are decided, $x_{m,n}^{t,\leftharpoonup}$ for $m \in \mathcal{R}$, $n \in \mathcal{R} \cup \{0\}$, $(m,n) \in \mathcal{E}^{\leftharpoonup}$. Large volumes of water transfer take time.

After this, uncertainties in rainfall w_r^t and consumption demand w_h^t realize. The decision-maker reacts through additional highly capacitated $K_{m,n}^{\rightharpoonup}$ and expensive $c_{m,n}^{t,\rightharpoonup}$ transfers between nodes, denoted as $x_{m,n}^{t,\rightharpoonup}$. The set of edges involved in these *recourse decisions* is denoted $\mathcal{E}^{\rightharpoonup} \subseteq \mathcal{E}$.

1. Balancing: Last minute balancing between reservoirs and between reservoir and the source/sink node is conducted, $x_{m,n}^{t,\rightharpoonup}$ for $m \in \mathcal{R}$, $n \in \mathcal{R} \cup \{0\}$, $(m,n) \in \mathcal{E}^{\rightharpoonup}$. The allowed water volume is limited, $K_{m,n}^{\rightharpoonup} \ll K_{m,n}^{\leftharpoonup}$, and the cost is pricier than here-and-now decisions, $c_{m,n}^{t,\rightharpoonup} > c_{m,n}^{t,\leftharpoonup}$.
2. Flooding events: To maintain complete recourse, we permit flooding events, modelled by separate edges from any water body $v \in \mathcal{R} \cup \mathcal{C}$ to the sink/source node 0, $x_{v,0}^{t,\rightharpoonup}$. They have unlimited capacity and incur astronomical costs $c_{v,0}^{t,\rightharpoonup} \gg 0$ to discourage its occurrence.
3. Drought events: Similarly, we permit drought events, modelled by separate edges from the sink/source node 0 to any water body $v \in \mathcal{R} \cup \mathcal{C}$, $x_{0,v}^{t,\rightharpoonup}$, representing a last recourse injection of water. They also have unlimited capacity and incur astronomical costs $c_{0,v}^{t,\rightharpoonup} \gg 0$.

At the end of the time period t, the new state of all water bodies s_v^t, $v \in \mathcal{R} \cup \mathcal{C}$ is updated:

$$s_v^t = s_v^{t-1} + \sum_{(m,v)\in\mathcal{E}^{\leftharpoonup}} x_{m,v}^{t,\leftharpoonup} - \sum_{(v,n)\in\mathcal{E}^{\leftharpoonup}} x_{v,n}^{t,\leftharpoonup} + \sum_{(m,v)\in\mathcal{E}^{\rightharpoonup}} x_{m,v}^{t,\rightharpoonup} - \sum_{(v,n)\in\mathcal{E}^{\rightharpoonup}} x_{v,n}^{t,\rightharpoonup} + w_v^t 1_{\{v\in\mathcal{R}\}} - w_v^t 1_{\{v\in\mathcal{H}\}}.$$

In the simulations, we record if any flooding or drought event has occurred (if their decision variables are non-zero). Table 2 in Appendix B.1 summarizes the notation. To ensure non-anticipativeness, first-stage $\boldsymbol{x}^{\leftharpoonup}$ are a function of uncertainties only in the previous time periods $\boldsymbol{w}_{<t}$, whereas second-stage $\boldsymbol{x}^{\rightharpoonup}$ can be a function of uncertainties including in that time period $\boldsymbol{w}_{\leqslant t}$. In total, the uncertainty has dimensions $\boldsymbol{w} \in \mathbb{R}^{T(|\mathcal{R}|+|\mathcal{H}|)}$, comprising rainfall into each reservoir $r \in \mathcal{R}$ and demand for water from each tank $h \in \mathcal{H}$, for every time period $t \in \mathcal{T}$.

4.1.2. Optimization formulations. The sequence of events leads to the following nominal stochastic optimization formulation for the water management problem (WMP):

$$\min_{\boldsymbol{x}^{\leftharpoonup}(\cdot),\boldsymbol{x}^{\rightharpoonup}(\cdot)} \quad \mathbb{E}_{\mathbb{Q}}\left[\sum_{t=1}^{T}\left(\boldsymbol{c}^{t,\leftharpoonup\top}\boldsymbol{x}^{t,\leftharpoonup}(\boldsymbol{w}_{<t})+\boldsymbol{c}^{t,\rightharpoonup\top}\boldsymbol{x}^{t,\rightharpoonup}(\boldsymbol{w}_{\leqslant t})\right)\right] \tag{WMP}$$

$$\text{s.t.} \quad \sum_{(v,n)\in\mathcal{E}^{\leftharpoonup}} x^{t,\leftharpoonup}_{v,n} \leqslant \tau_v, \qquad \forall v\in\mathcal{P}\cup\mathcal{D}, t\in\mathcal{T} \tag{15}$$

$$\sum_{(m,p)\in\mathcal{E}^{\leftharpoonup}} x^{t,\leftharpoonup}_{m,p} = \sum_{(p,n)\in\mathcal{E}^{\leftharpoonup}} x^{t,\leftharpoonup}_{p,n}, \qquad \forall p\in\mathcal{P}, t\in\mathcal{T} \tag{16}$$

$$\iota_v \leqslant s^t_v + \sum_{(m,v)\in\mathcal{E}^{\leftharpoonup}} x^{t,\leftharpoonup}_{m,v} - \sum_{(v,m)\in\mathcal{E}^{\leftharpoonup}} x^{t,\leftharpoonup}_{v,m} + \sum_{(i,v)\in\mathcal{E}^{\rightharpoonup}} x^{t,\rightharpoonup}_{i,v} \tag{17}$$

$$- \sum_{(v,j)\in\mathcal{E}^{\rightharpoonup}} x^{t,\rightharpoonup}_{v,j} + w^t_v 1_{\{v\in\mathcal{R}\}} - w^t_v 1_{\{v\in\mathcal{H}\}} \leqslant \kappa_v, \qquad \forall v\in\mathcal{R}\cup\mathcal{H}, t\in\mathcal{T}$$

$$0 \leqslant x^{t,\leftharpoonup}_{m,n} \leqslant K^{\leftharpoonup}_{m,n}, \qquad \forall (m,n)\in\mathcal{E}^{\leftharpoonup}, t\in\mathcal{T} \tag{18}$$

$$0 \leqslant x^{t,\rightharpoonup}_{m,n} \leqslant K^{\rightharpoonup}_{m,n}, \qquad \forall (m,n)\in\mathcal{E}^{\rightharpoonup}, t\in\mathcal{T} \tag{19}$$

Here, the objective comprises the costs incurred in every time period, with contributions from the here-and-now and recourse decisions respectively. There are no costs incurred directly by the state of the system at the end of time period $\boldsymbol{s}^t$ or explicitly as a function of the uncertainties. Constraint (15) limits the maximum production capacity of production facilities and (16) ensures that inflow and outflow at these facilities are balanced. (17) contains the final state of a water body within its upper and lower capacity limits. Finally, (18) and (19) constrain the capacities on the flows for the here-and-now and recourse decisions respectively.

To set up the robust optimization formulation, we use the parent ambiguity set:

$$\mathcal{U} := \left\{ \mathbb{Q}\in\mathcal{M}\left(\mathbb{R}^{T(|\mathcal{R}|+|\mathcal{H}|)},\mathbb{R}\right) \;\middle|\; \boldsymbol{w}\sim\mathbb{Q}: \begin{array}{l} \mathbb{E}_{\mathbb{Q}}\left[\boldsymbol{G}\boldsymbol{w}\right]\leqslant\boldsymbol{\mu}, \\ \mathbb{Q}\left[\boldsymbol{w}\in\mathcal{W}\right]=1. \end{array} \right\}. \tag{20}$$

For narrative brevity, the corresponding (WMP-DRO) and (WMP-RO) versions are provided in Appendix B.2. Here, $\boldsymbol{G}\in\mathbb{R}^{b\times T(|\mathcal{R}|+|\mathcal{H}|)}$ is a given matrix of coefficients and $\boldsymbol{\mu}\in\mathbb{R}^b$ is the constraint vector, where b represents the number of constraints imposed. This induces the following models. For brevity, when we say that a model is solved in *rolling horizon* fashion, we mean that its corresponding problem-to-go is re-optimized at each stage using the present state as initial conditions, and only first-stage decisions are implemented.

i. RO: We solve (WMP-RO) with support set $\mathcal{W}$. It is implemented both as a static model and updated in rolling-horizon fashion.

ii. DRO: We solve (WMP-DRO) under ambiguity set (20). It is implemented statically (*i.e.*, decisions $\boldsymbol{x}^{\leftharpoonup}(\cdot)$ and $\boldsymbol{x}^{\rightharpoonup}(\cdot)$ are implemented with no updating or resolving), or with naïve updating. Without time consistency guarantees, a rolling-horizon implementation requires updating (20), by updating the right-hand side of the distributional constraint, to obtain ambiguity set

$$\ddot{\mathcal{U}}_{\epsilon}(\hat{\boldsymbol{w}}_{<t}) := \left\{ \mathbb{Q}\in\mathcal{M}\left(\mathbb{R}^{(T-t+1)(|\mathcal{R}|+|\mathcal{C}|)},\mathbb{R}\right) \;\middle|\; \begin{array}{l} \boldsymbol{w}_{\geqslant t}\sim\mathbb{Q}: \\ \quad \mathbb{E}_{\mathbb{Q}}\left[\boldsymbol{G}_{\geqslant t}\boldsymbol{w}_{\geqslant t}\right]\leqslant\boldsymbol{\mu}_t := \boldsymbol{\mu}-\hat{\boldsymbol{G}}_{<t}\hat{\boldsymbol{w}}_{<t}+\epsilon \\ \quad \mathbb{Q}\left[\boldsymbol{w}_{\geqslant t}:(\hat{\boldsymbol{w}}_{<t},\boldsymbol{w}_{\geqslant t})\in\mathcal{W}\right]=1. \end{array} \right\}, \tag{21}$$

where $\boldsymbol{\mu} - \hat{\boldsymbol{G}}_{<t}\hat{\boldsymbol{w}}_{<t}$ is the expected value-to-go of the constraint and ϵ loosens the constraint. When ϵ is small, $\ddot{\mathcal{U}}_\epsilon(\hat{\boldsymbol{w}}_{<t}) = \varnothing$ possibly for particular histories $\hat{\boldsymbol{w}}_{<t}$. This can occur despite complete recourse where decisions $\boldsymbol{x}^{\curvearrowleft}$ and $\boldsymbol{x}^{\curvearrowright}$ are always feasible, whereupon we term it "infeasibility". $\ddot{\mathcal{U}}_\epsilon(\hat{\boldsymbol{w}}_{<t})$ can be made nonempty for ϵ large enough, though at the cost of additional conservativeness. Specifically, we will use $\epsilon \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$.

iii. AURA: (AURA) is always solved under the ambiguity updating series (LAUS) induced by parent ambiguity set (20), and always implemented in rolling-horizon fashion.

CVD-regularization is always implemented on the parent ambiguity set (20),

$$\mathcal{U}^{\text{CVD}}(L) := \left\{ \mathbb{Q} \in \mathcal{M}\left(\mathbb{R}^{T(|\mathcal{R}|+|\mathcal{C}|)}, \mathbb{R}\right) \;\middle|\; \begin{array}{l} \boldsymbol{w} \sim \mathbb{Q}: \\ \quad \mathbb{E}_{\mathbb{Q}}[\boldsymbol{G}\boldsymbol{w}] \leqslant \boldsymbol{\mu}, \\ \quad \left| \int_{\bar{B}_k} \mathbb{P}(d\boldsymbol{w}) - \int_{\underline{B}_k} \mathbb{P}(d\boldsymbol{w}) \right| \leqslant L \cdot D(\bar{B}_k, \underline{B}_k), k = 1, \ldots, K, \\ \quad \mathbb{Q}[\boldsymbol{w} \in \mathcal{W}] = 1. \end{array} \right\}. \tag{22}$$

For details on how sets $(\bar{B}_k, \underline{B}_k)$ are chosen, please refer to Appendix B.3.2. When discussing CVD-regularization, we will only be comparing AURA against DRO(static) (both solved under $\mathcal{U}^{\text{CVD}}(L)$) and RO (which is unaffected by CVD-regularization). It is unclear how to update CVD-regularized DRO under rolling horizon, so we do not test it.

4.1.3. Experimental set-up. For simplicity, we consider a simple water network with 2 reservoirs, 1 treatment plant, 1 desalination plant and 1 tank. Figure 9 in Appendix B.1 illustrates how these nodes are connected and the edges that are active in the here-and-now decisions before uncertainties materialize $\mathcal{E}^{\curvearrowleft}$ and in the recourse decisions after uncertainties materialize $\mathcal{E}^{\curvearrowright}$. For simplicity, we consider $T = 2$ in this illustration, hence it is a two-period, two-stage problem.

We assume that $\boldsymbol{G}$ and $\mathcal{W}$ are known and without error, *i.e.*, they are consistent in- and out-of-sample. The matrix $\boldsymbol{G}$ has $b = |\mathcal{R}| + |\mathcal{H}| + 1$ constraints on (i) the average rainfall expected for each reservoir $r \in \mathcal{R}$ over the T time periods, $\mathbb{E}[\sum_t w_r^t] \leqslant \mu_r$, (ii) the average demand of water expected for each tank $h \in \mathcal{H}$ over the T time periods, $\mathbb{E}[\sum_t w_h^t] \leqslant \mu_h$, and (iii) the average net inflow into the system over the T time periods, $\mathbb{E}[\sum_r \sum_t w_r^t - \sum_h \sum_t w_h^t] \leqslant \mu_{\text{net}}$. True distribution $\mathbb{P}$ need not lie in the ambiguity set (20) used to solve the models. We simulate this by assuming that $\mathbb{E}_{\mathbb{P}}[\boldsymbol{G}\boldsymbol{w}] = \boldsymbol{\mu}_{\text{OOS}}$, which is elementwise either (i) equal $\boldsymbol{\mu}_{\text{OOS}} = \boldsymbol{\mu}$ (exactly specified), (ii) smaller $\boldsymbol{\mu}_{\text{OOS}} < \boldsymbol{\mu}$ (lying strictly within $\mathcal{U}$, hence is over-specified) or larger $\boldsymbol{\mu}_{\text{OOS}} > \boldsymbol{\mu}$ (lying outside $\mathcal{U}$, hence is under-specified) than the assumed vector $\boldsymbol{\mu}$. This is achieved by letting $\boldsymbol{\mu}_{\text{oos}}$ take values in $\{0.99\,\boldsymbol{\mu},\ \boldsymbol{\mu},\ 1.01\,\boldsymbol{\mu}\}$. Uncertainties are scaled so that rainfall is supported on $[0, 1]$. Correspondingly, demand is scaled to fall within $[1, 3]$ (it is larger than total rainfall so that water production is necessary, in expectation). The resulting support is $\mathcal{W} \subseteq [0, 1]^{|\mathcal{R}| \times T} \times [1, 3]^{|\mathcal{H}| \times T}$. In the base case, we specify $\boldsymbol{\mu} := (\mu_{r_1}, \mu_{r_2}, \mu_d, \mu_{\text{net}}) = (0.7,\ 0.8,\ 4,\ -3)$, where DRO is feasible and DRO $<$ RO.

To undertake out-of-sample testing, we generate 20 test distributions for each case of $\boldsymbol{\mu}_{\text{OOS}}$ tested. All results presented are averages over the 20 distributions unless otherwise stated. For more details on generation of these distributions, please see Appendix B.3.1.

4.2. Discrete uncertainties

In this section, we simplify the experimental setting by assuming rainfall realizations within $\{0, 0.25, 0.5, 0.75, 1\}$, while water demand is discretized into {1, 1.5, 2, 2.5, 3}. The support set of the uncertainties is assumed to be known and is constructed from 300 sampled historical paths.

4.2.1. Infeasibility and conservativeness of Naïve updating. In our experiments, (dual) infeasibility arising from time inconsistency in DRO does occur. This motivates updating the ambiguity set. Table 1 reports the infeasibility rates for DRO(naïve) under different values of ϵ.

Table 1 Rate of infeasibility of DRO(naïve) under different RHS relaxation ϵ

After uncertainty revelation in time period t	ϵ					
	0.0	0.2	0.4	0.6	0.8	1.0
$t=1$	81.82%	69.70%	48.48%	30.30%	12.12%	0.0%
$t=2$	81.82%	69.70%	48.48%	30.30%	12.12%	0.0%

When ϵ is small enough, time inconsistency presents as infeasibility. In 12.12% of sample paths tested, the problems-to-go were dual infeasible, *i.e.* DRO(naïve) model failed to return a solution, even after relaxing the RHS of the distributional constraint by $\epsilon \leqslant 0.8$. Feasibility is only recovered after ϵ reaches 1.0, a value exceeding the maximum expected rainfall $(0.7, 0.8)$ of the reservoirs! Relaxing constraint by this much necessarily comes at the cost of increased conservativeness. In Figure 3, we plot the out-of-sample costs, obtained by varying $\boldsymbol{\mu}_{\text{OOS}}$, incurred by the four models, with DRO(naïve) is fixed at $\epsilon = 1.0$, where feasibility is guaranteed.

Naïve updating pays the price of conservativeness for maintaining feasibility under rolling horizon. The relaxation is so large that DRO(naïve) performs almost as poorly as RO(rolling). In contrast, AURA, which directly models the remaining set of conditional distributions, achieves the best out-of-sample performance. We also recover the common observation in the literature that a rolling horizon implementation of DRO often performs better than the static version.

4.2.2. DRO vs AURA vs RO. We temporarily leave Figure 3 behind to delve into the relationship between DRO, RO and AURA. In-sample, Theorem 1 guarantees (DRO) $\leqslant$ (AURA) $\leqslant$ (RO). We examine how the *time consistency premium* TCP $:= \frac{\text{AURA}-\text{DRO}}{\text{DRO}}$, which represents the (in-sample) price that AURA pays for making DRO time consistent, as a fraction of DRO's cost,

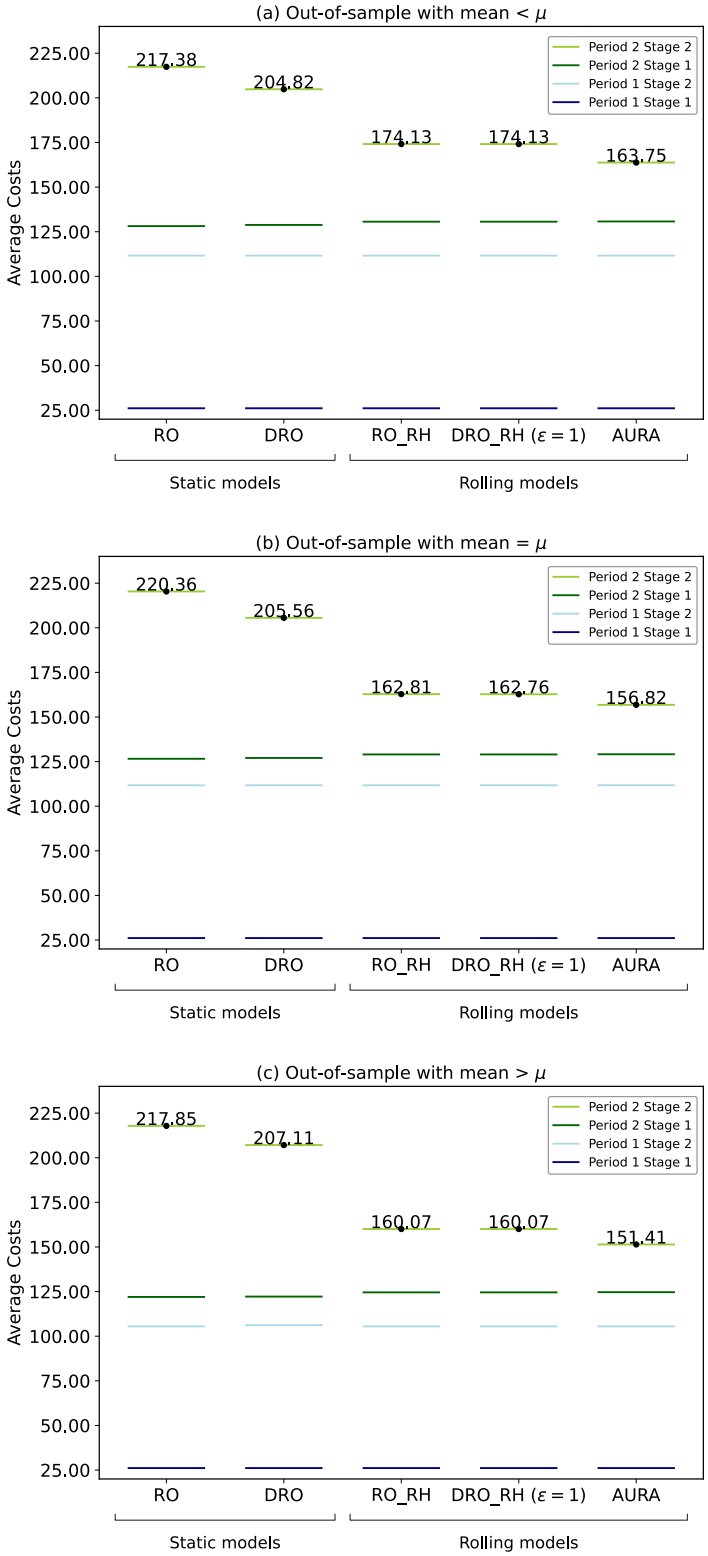


Figure 3 Average out-of-sample costs of different models, broken down by period and stage, under different evaluation means $\boldsymbol{\mu}_{\text{OOS}}$.

and the *collapse gap* $\text{CG} := \frac{\text{RO}-\text{AURA}}{\text{RO}}$, which represents the savings upon RO obtained for solving the distributionally robust AURA, as a fraction of RO's cost, vary with tightness of the parent ambiguity set. We do so by fixing the maximum expected rainfall at $\mu_{r_1} = 0.7$ and $\mu_{r_2} = 0.8$ for both reservoirs, while varying the maximum expected demand over $\mu_d \in [3.7, 5]$ and net inflow over $\mu_{\text{net}} \in [-5.0, -2.5]$. Results are reported in Figures 4. Notably, we recover (DRO) $\leqslant$ (AURA) $\leqslant$ (RO) in all cases, which verifies Theorem 1i..

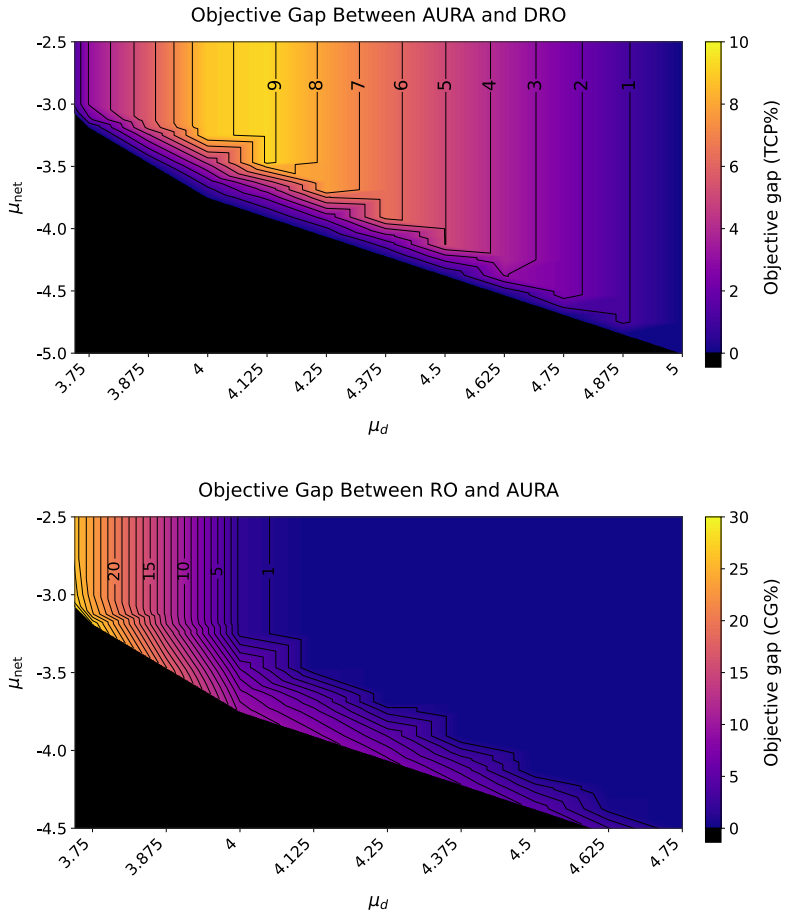


Figure 4 TCP (above) and CG (below) as maximum expected demand μ_d and net inflow μ_{net} varies. Black regions relate to values without feasible distributions for DRO or AURA.

The TCP is largest when the parent ambiguity set is neither too large (moving rightwards or upwards in the chart) or too small (moving leftwards or downwards). When the parent ambiguity set is small, there are few feasible distributions and both AURA and DRO will identify very similar worst-case distributions; when the set is large, both models approach RO as the ambiguity set comes close to all distributions defined on the support (*i.e.* the distributional constraint becomes so loose as to be redundant), and thus agree in that limit. The intermediate range where the parent ambiguity set is neither large or small, presents a large boundary where seeking "closure" under conditioning would lead to many new distributions considered in AURA but not in DRO.

Instead, the CG is monotone non-increasing as the ambiguity set grows. This behavior is expected because RO is independent of the ambiguity set, whereas a larger ambiguity set makes AURA more conservative thus its objective value increases.

4.2.3. In-sample vs out-of-sample. While AURA pays a TCP in-sample in Figure 4, it gains in time consistency out-of-sample in relation to DRO, as seen by the reversal in performance in Figure 3. Note that AURA's final stage final period costs are the least among all the models. Barely visible, AURA incurs the highest costs in the first-stage of period 2. This explains how Theorem 1 works – by being more conservative, AURA reaps long-term dividends, where customizing to the revealed history through updating and the avoidance of time inconsistency penalties depress recourse costs. These findings are also robust whether the true distribution $\mathbb{P}$ lies inside, on the boundary of, or outside the parent ambiguity set.

At this point, it is important to explain the nuance in how Figure 3 is generated. Crucially, the optimal solution of (WMP-DRO) has degeneracy for the same reasons that we discussed for AURA in Theorem 1ii. – decisions on non-worst-case sample paths do not influence the overall objective value, thus are degenerate. However, different degenerate solutions do result in vastly different out-of-sample costs. We plot this in Figure 5. Here, we search among the degenerate solutions of DRO out-of-sample, for each testing distribution, to find the worst-performing, best-performing and default solver solutions, and compare them against AURA.

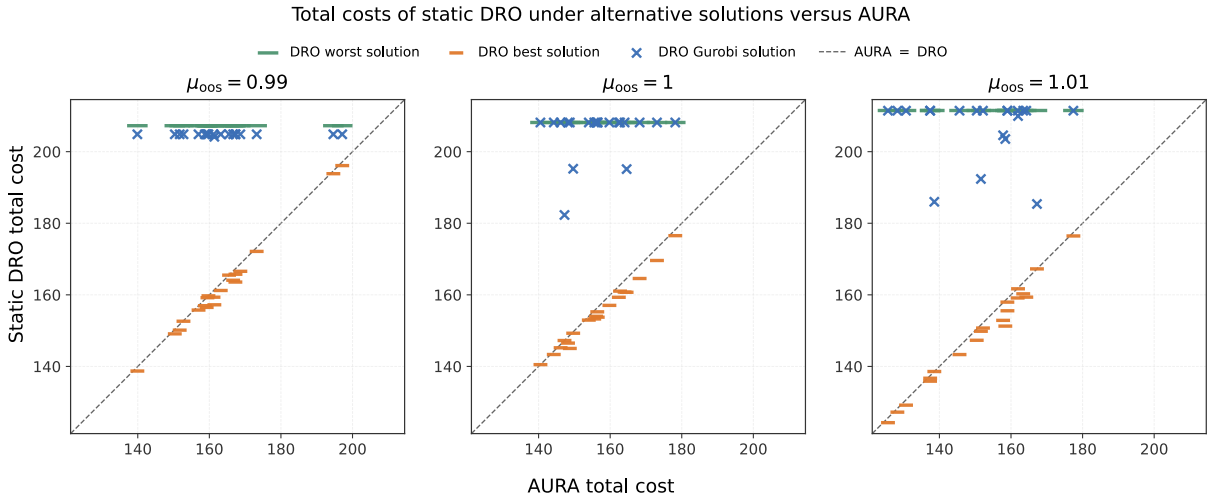


Figure 5 Total Costs of Static DRO under Alternative Solutions versus AURA across Test Distributions.

The results are illuminating. Best-performing DRO solutions may even slightly beat AURA's out-of-sample costs. This makes sense because AURA is indeed more conservative. Unfortunately, the worst-performing DRO solutions can be very much worse than AURA (as much as 68% costlier),

especially when the true distribution lies outside of the parent ambiguity set. Gravely, to the best of our knowledge, we do not know of any method, computationally or analytically, to differentiate these degenerate solutions in-sample for DRO since without ambiguity updating, DRO does not have a well-defined problem-to-go! AURA, on the other hand, has a well-defined problem-to-go that always selects the time consistent solution. Consequently, it becomes a gamble as to which solutions is picked for DRO. Incidentally, solvers have a tendency to select the worst-performing DRO solutions, as Figure 5 indicates! This little analysis stresses the importance of time consistency and the theoretical backing that underpins AURA's construction.

4.2.4. Ambiguity collapse. Finally, we also illustrate the presence of ambiguity collapse. To do so, we track the CG as the discretization of $\mathcal{W}$ (rainfall between $[0,1]$ and demand between $[1,3]$) is progressively made finer. We use $\boldsymbol{\mu} = (0.7, 0.8, 4.29, -3)$ for this illustration so as to maximize the starting CG (*c.f.* Figure 4 below). We plot this in Figure 6. Note that RO's objective value remains unchanged throughout, since the worst-case uncertainty $\boldsymbol{w}^{\mathrm{RO}}$ is an extreme point.

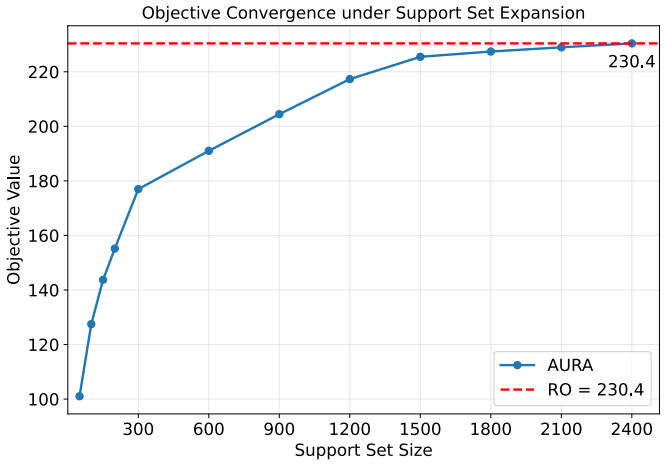


Figure 6 Objective convergence of AURA to RO (dotted) as the support set size increases.

From the Figure, the CG gradually decreases as the support set becomes denser. As more points are added to the support set, the likelihood that the worst-case uncertainty $\boldsymbol{w}^{\mathrm{RO}}$ is sampled increases. This eventually occurs at $2,400$ sample paths, whereby AURA collapses to RO.

4.3. Continuous uncertainty

In this subsection, we approximate continuous support $\mathcal{W}$, using the $2,400$ sample paths of Section 4.2.4, retaining the parameters of that experiment. As we have seen, AURA = RO at that specification. Here, we consider the effect of CVD-regularization. We specify two pairs of sets $(\bar{B}_k, \underline{B}_k)$, $k = 1, 2$, and vary the Lipschitz bounds L. Figure 7 reports the resulting in-sample objectives.

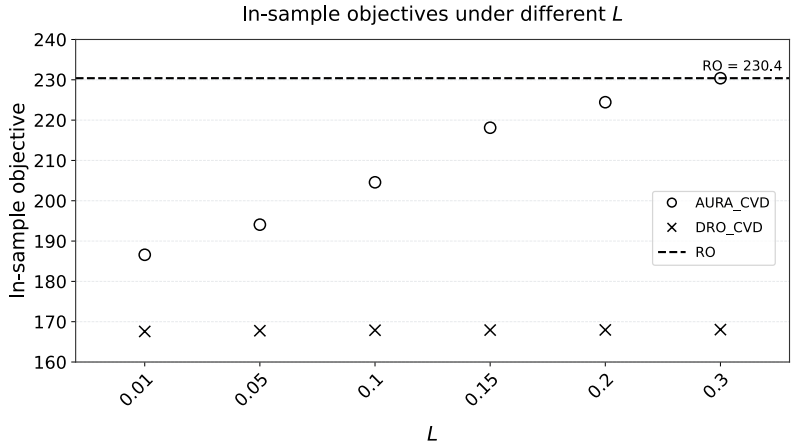


Figure 7 In-sample objective values of RO, AURA and DRO by varying Lipscthiz parameter L.

AURA's objective value increases as L increases and approaches that of RO at $L = 0.3$. This is expected: when L is large, few CC distributions are eliminated and the un-regularized AURA is recovered, which collapses to RO. This empirically validates the result of Theorem 5 that CVD-regularization averts collapse. While slight, the objective value of DRO does also increase with L. This is consequential as it confirms that the removal of CC distributions from the parent ambiguity set does not come at a high penalty to conservativeness, indicating that CC distributions do not necessarily play a role in deciding the worst-case distributions when updating is not considered.

In Figure 8, we study the out-of-sample costs of AURA and DRO, as a proportion of RO's out-of-sample cost, which is unaffected by CVD-regularization. There are three notable observations when focusing on the exactly specified case of $\boldsymbol{\mu}_{\text{oos}} = \boldsymbol{\mu}$. First, at large L, AURA's performance approaches RO's. This again is consistent with the notion of collapse. More interestingly, we notice a threshold value, whereupon the performance of AURA splits into two regimes, a collapse regime, where L is not small enough to eliminate enough CC distributions, and a buttress regime, where AURA achieves around 10% lower costs than AURA. Further analysis indicates that there is in fact a continuum between the two. Second, the minor changes in DRO's objective value in Figure 7 are amplified in this out-of-sample result, however, the change in DRO's performance with L is still small compared to AURA. This again validates our earlier conclusion that CVD-regularization does not play a fundamental part in the conservativeness of DRO solutions. Finally, we note that in the buttress regime, AURA out-performs DRO, indicating that once CC distibutions are sufficiently controlled and collapse is avoided, we successfully recover the gains from updating.

We also examine the over- and under-specified cases. The performance when $\boldsymbol{\mu}_{\text{oos}} < \boldsymbol{\mu}$ is very similar to case when $\boldsymbol{\mu}_{\text{oos}} = \boldsymbol{\mu}$. When the ambiguity set is under-specified, *i.e.* $\boldsymbol{\mu}_{\text{oos}} < \boldsymbol{\mu}$, the true distribution will lie outside the ambiguity set. Here, CVD-regularization loses its robustness rapidly with both DRO and AURA's performance degrading at roughly the same degree. AURA maintains its edge over RO, but DRO could potentially perform even worse that RO (under rolling horizon).

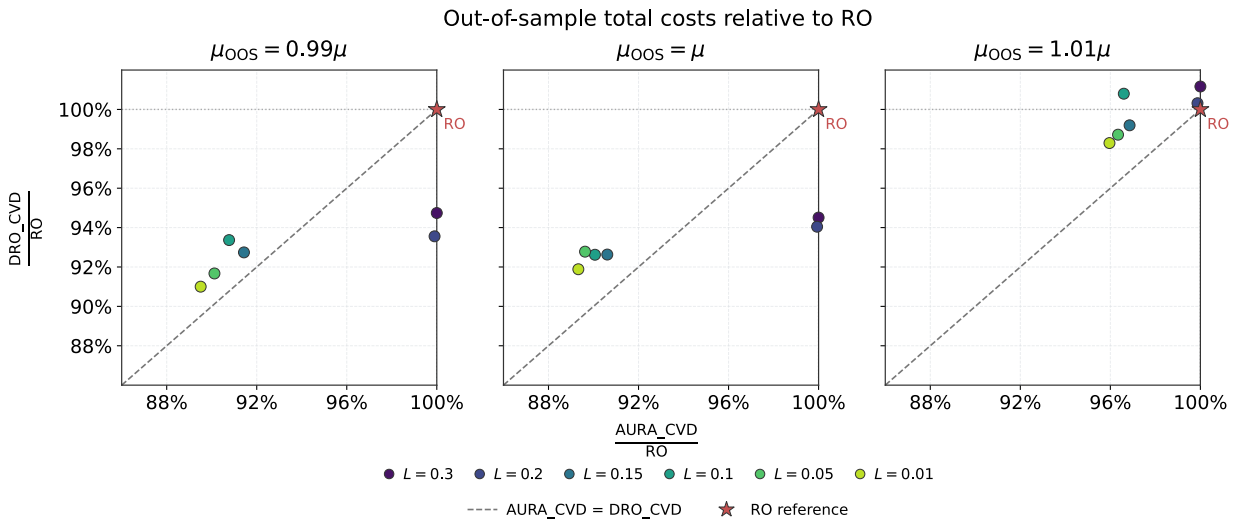


Figure 8 Out-of-sample total costs of AURA and DRO with respect to RO by varying Lipschitz parameter L.

5. Conclusion

This paper is the first to construct an ambiguity set of conditional distributions in multi-period DRO and its reformulation, via our proposed AURA framework. It recovers time consistency. When the support is continuous, AURA collapses to the RO problem, driven by the presence of CC distributions, which can be eliminated via CVD-regularization to avert collapse. Our journey from ambiguity collapse then back tells an important moral – ambiguity sets can contain distributions that do not respect temporal rules! While we examined CC distributions here, other temporal structures could still be present in ambiguity sets. An example is time reversibility; many ambiguity sets are isometric and do not prohibit substituting $\boldsymbol{w}_t$ with $\boldsymbol{w}_{t'}$ say. Thus, we conclude our treatise with a call-to-arms for more study on defining structurally consistent yet tractable ambiguity sets.

References

Ben-Tal, Aharon, Alexander Goryashko, Elana Guslitzer, Arkadi Nemirovski. 2004. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming* **99**(2) 351–376.

Bertsimas, Dimitris, Dan A Iancu, Pablo A Parrilo. 2009. Optimality of affine policies in multi-stage robust optimization. *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*. IEEE, 1131–1138.

Bertsimas, Dimitris, Nathan Kallus. 2020. From predictive to prescriptive analytics. *Management Science* **66**(3) 1025–1044.

Bertsimas, Dimitris, Melvyn Sim, Meilin Zhang. 2019. Adaptive distributionally robust optimization. *Management Science* **65**(2) 604–618.

Calafiore, Giuseppe C. 2007. Ambiguous risk measures and optimal robust portfolios. *SIAM Journal on Optimization* **18**(3) 853–877.

Carpentier, Pierre, Jean-Philippe Chancelier, Guy Cohen, Michel De Lara, Pierre Girardeau. 2012. Dynamic consistency for stochastic optimal control problems. *Annals of Operations research* **200**(1) 247–263.

Chand, Suresh, Vernon Ning Hsu, Suresh Sethi. 2002. Forecast, solution, and rolling horizons in operations management problems: A classified bibliography. *Manufacturing & Service Operations Management* **4**(1) 25–43.

Chen, Zhi, Melvyn Sim, Peng Xiong. 2020. Robust stochastic optimization made easy with RSOME. *Management Science* **66**(8) 3329–3339.

Cooper, A Charnes WW, W Charnes. 1962. Programming with linear fractional functionals. *Naval Research logistics quarterly* **9**(3) 181–186.

Delage, Erick, Yinyu Ye. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research* **58**(3) 595–612.

Elmachtoub, Adam N, Paul Grigas. 2022. Smart "predict, then optimize". *Management Science* **68**(1) 9–26.

Esteban-Pérez, Adrián, Juan M Morales. 2022. Distributionally robust stochastic programs with side information based on trimmings. *Mathematical Programming* **195**(1) 1069–1105.

Gao, Rui, Rohit Arora, Yizhe Huang. 2024. Data-driven multistage distributionally robust linear optimization with nested distance. *arXiv preprint arXiv:2407.16346* .

Gao, Rui, Anton J Kleywegt. 2017. Distributionally robust stochastic optimization with dependence structure. *arXiv preprint arXiv:1701.04200* .

Georghiou, Angelos, Angelos Tsoukalas, Wolfram Wiesemann. 2019. Robust dual dynamic programming. *Operations Research* **67**(3) 813–830.

Hu, Jie, Zhi Chen, Shuming Wang. 2025. Budget-driven multiperiod hub location: A robust time-series approach. *Operations Research* **73**(2) 613–631.

Kannan, Rohit, Güzin Bayraksan, James R Luedtke. 2024. Residuals-based distributionally robust optimization with covariate information. *Mathematical Programming* **207**(1) 369–425.

Lin, Yifan, Yuxuan Ren, Enlu Zhou. 2022. Bayesian risk markov decision processes. *Advances in Neural Information Processing Systems* **35** 17430–17442.

Long, Daniel Zhuoyu, Jin Qi, Aiqi Zhang. 2024. Supermodularity in two-stage distributionally robust optimization. *Management Science* **70**(3) 1394–1409.

Mohajerin Esfahani, Peyman, Daniel Kuhn. 2018. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming* **171**(1) 115–166.

Shapiro, Alexander. 2001. On duality theory of conic linear problems. *Nonconvex Optimization and its Applications* **57** 135–155.

Shapiro, Alexander. 2009. On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters* **37**(3) 143–147.

Shapiro, Alexander. 2016. Rectangular sets of probability measures. *Operations Research* **64**(2) 528–541.

Shapiro, Alexander, Darinka Dentcheva, Andrzej Ruszczynski. 2021. *Lectures on stochastic programming: modeling and theory*. SIAM.

Shapiro, Alexander, Arkadi Nemirovski. 2005. On complexity of stochastic programming problems. *Continuous optimization: Current trends and modern applications* 111–146.

Shapiro, Alexander, Linwei Xin. 2020. Time inconsistency of optimal policies of distributionally robust inventory models. *Operations Research* **68**(5) 1576–1584.

Shapiro, Alexander, Enlu Zhou, Yifan Lin. 2023. Bayesian distributionally robust optimization. *SIAM Journal on Optimization* **33**(2) 1279–1304.

Sim, Melvyn, Qinshen Tang, Minglong Zhou, Taozeng Zhu. 2025. The analytics of robust satisficing: predict, optimize, satisfice, then fortify. *Operations Research* **73**(5) 2708–2728.

Taskesen, Bahar, Dan Iancu, Çağıl Koçyiğit, Daniel Kuhn. 2023. Distributionally robust linear quadratic control. *Advances in Neural Information Processing Systems* **36** 18613–18632.

Wiesemann, Wolfram, Daniel Kuhn, Berç Rustem. 2013. Robust markov decision processes. *Mathematics of Operations Research* **38**(1) 153–183.

Wiesemann, Wolfram, Daniel Kuhn, Melvyn Sim. 2014. Distributionally robust convex optimization. *Operations research* **62**(6) 1358–1376.

Xin, Linwei, David A Goldberg. 2021. Time (in) consistency of multistage distributionally robust inventory models with moment constraints. *European Journal of Operational Research* **289**(3) 1127–1141.

Xin, Linwei, David Alan Goldberg. 2022. Distributionally robust inventory control when demand is a martingale. *Mathematics of Operations Research* **47**(3) 2387–2414.

Xu, Huan, Shie Mannor. 2006. The robustness-performance tradeoff in markov decision processes. *Advances in Neural Information Processing Systems* **19**.

Xu, Huifu, Yongchao Liu, Hailin Sun. 2018. Distributionally robust optimization with matrix moment constraints: Lagrange duality and cutting plane methods. *Mathematical Programming* **169**(2) 489–529.

Yu, Xian, Siqian Shen. 2022. Multistage distributionally robust mixed-integer programming with decision-dependent moment-based ambiguity sets. *Mathematical Programming* **196**(1) 1025–1064.

Zhu, Taozeng, Jingui Xie, Melvyn Sim. 2022. Joint estimation and robustness optimization. *Management Science* **68**(3) 1659–1677.

A. Proofs

We do not prove Proposition 1 and the reader is referred to Shapiro and Xin (2020).

A.1. Claim in Example 1

We illustrate this for $\hat{w}_1 = 0$ as the case of $\hat{w}_1 = 1$ is analogous. Consider $\mu \in [1,2)$. Consider the following two parent distributions for some $\alpha : 0 < \alpha < \frac{2-\mu}{2}$:

$$\mathbb{P}_1[\boldsymbol{w}] = \begin{cases} \alpha & \text{if } \boldsymbol{w} = (0,0) \\ 2(1-\alpha) - \mu & \text{if } \boldsymbol{w} = (1,0) \\ \mu - (1-\alpha) & \text{if } \boldsymbol{w} = (1,1) \end{cases}$$

$$\mathbb{P}_2[\boldsymbol{w}] = \begin{cases} \alpha & \text{if } \boldsymbol{w} = (0,1) \\ 2(1-\alpha) - (\mu-\alpha) & \text{if } \boldsymbol{w} = (1,0) \\ \mu - 1 & \text{if } \boldsymbol{w} = (1,1) \end{cases}$$

In both cases, it is not difficult to verify that $\mathbb{E}_{\mathbb{P}_1}[w_1 + w_2] = \mu = \mathbb{E}_{\mathbb{P}_2}[w_1 + w_2]$, thus they both lie in the parent set $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{U}$. However, under $\mathbb{P}_1$ and $\mathbb{P}_2$, w_2 takes values 0 and 1 respectively with probability 1, conditional on the event $w_1 = \hat{w}_1 = 0$. Thus, $\mathbb{E}_{\mathbb{P}_1}[w_2 | w_1 = \hat{w}_1] = 0 \neq 1 = \mathbb{E}_{\mathbb{P}_2}[w_2 | w_1 = \hat{w}_1]$. Hence, there does not exist a ν such that both $\mathbb{P}_1|_{w_1=\hat{w}_1}, \mathbb{P}_2|_{w_1=\hat{w}_1} \in \mathcal{V}(\hat{w}_1; \nu)$. □

Remark 2. The proof also indicates that the best we can do, if we are to preserve the structure of the moments constraint and still require all conditional distributions of parent distributions to be contained in the ambiguity set, is

$$\mathcal{V}(\hat{w}_1) := \left\{ \mathbb{P} \in \mathcal{M}\left(\mathcal{W}_2 := \{0,1\}, \mathbb{R}\right) \;\middle|\; \begin{array}{l} w_2 \sim \mathbb{P}: \\ \quad \mathbb{E}_{\mathbb{P}}[w_2] \in [0,1] \\ \quad \mathbb{P}[w_2 : (\hat{w}_1, w_2) \in \mathcal{W}] = 1 \end{array} \right\}.$$

But this makes the moment constraint $\mathbb{E}_{\mathbb{P}}[w_2] \in [0,1]$ utterly redundant, since $w_2 \in \{0,1\}$ anyway. In other words, we need to be far more conservative than required, if we wish to only 'update the RHS' to preserve the geometry of the ambiguity set, something that is also concurred in our numerical experiments.

A.2. Theorem 1

For brevity, we suppress the dependence of $\boldsymbol{x}$ (and $\boldsymbol{w}$) on $\boldsymbol{c}$. Let AURA($\boldsymbol{x}$) and DRO($\boldsymbol{x}$) denote the objective values of (AURA) and (DRO) for a given policy $\boldsymbol{x}$, with optimal policies $\boldsymbol{x}^*_{\text{AURA}}$ and $\boldsymbol{x}^*_{\text{DRO}}$ respectively.

Part i. We first show the relationship between (AURA) and (DRO): We aim to show that AURA($\boldsymbol{x}$) $\geqslant$ DRO($\boldsymbol{x}$) for any feasible decision policy $\boldsymbol{x}(\cdot)$, then (AURA) $= \sup_{\boldsymbol{x}}$ AURA($\boldsymbol{x}$) $\geqslant$ $\sup_{\boldsymbol{x}}$ DRO($\boldsymbol{x}$) $=$ (DRO). Fix policy $\boldsymbol{x}(\cdot)$. For any $t = 2, ..., T$,

$$\sup_{\mathbb{Q}_t^{(\cdot)}} \mathbb{E}_{\mathbb{Q}_{t-1}}\left[c_{t-1} + \mathbb{E}_{\mathbb{Q}_t^{(\boldsymbol{w}_{\leqslant t-1})}}[F_t(\boldsymbol{x}_{\geqslant t}; \hat{\boldsymbol{w}}_{\leqslant t-1}, \boldsymbol{w}_t)] \right] \tag{23}$$

$$\geqslant \mathbb{E}_{\mathbb{Q}_{t-1}}\left[c_{t-1}+\sup_{\overline{\mathbb{Q}}_t^{(\boldsymbol{w}_{\leqslant t-1})}\in\mathcal{U}_t(\boldsymbol{w}_{\leqslant t-1})}\mathbb{E}_{\overline{\mathbb{Q}}_t^{(\boldsymbol{w}_{\leqslant t-1})}}[F_t(\boldsymbol{x}_{\geqslant t};\hat{\boldsymbol{w}}_{\leqslant t-1},\boldsymbol{w}_t)]\right].$$

$\mathbb{Q}_t^{(\boldsymbol{w}_{\leqslant t-1})}:=\mathbb{Q}_t^{*(\boldsymbol{w}_{\leqslant t-1})}$, where $\mathbb{Q}_t^{*(\boldsymbol{w}_{\leqslant t-1})}\in\operatorname{argsup}_{\overline{\mathbb{Q}}_t^{(\boldsymbol{w}_{\leqslant t-1})}\in\mathcal{U}_t(\boldsymbol{w}_{\leqslant t-1})}\mathbb{E}_{\overline{\mathbb{Q}}_t^{(\boldsymbol{w}_{\leqslant t-1})}}[F_t(\boldsymbol{x}_{\geqslant t};\hat{\boldsymbol{w}}_{\leqslant t-1},\boldsymbol{w}_t)]$ for all $\boldsymbol{w}_{\leqslant t-1}\in\hat{\mathcal{W}}_{\leqslant t-1}$, is always a feasible solution for the problem on the LHS of the inequality (23). Moreover, it attains the same objective value on both sides of the inequality if and only if $\mathbb{Q}_t^{*(\boldsymbol{w}_{\leqslant t-1})}$ is a measurable function in $\boldsymbol{w}_{\leqslant t-1}$, which holds due to the condition 'traceable', which requires that $\mathbb{Q}_t^{*(\boldsymbol{w}_{\leqslant t-1})}$ originates from some distribution in the parent ambiguity set. Hence, by induction,

$$\begin{aligned}\mathrm{AURA}(\boldsymbol{x}) &= \sup_{\mathbb{Q}_1}\mathbb{E}_{\mathbb{Q}_1}\left[c_1+\sup_{\mathbb{Q}_2^{(\boldsymbol{w}_1)}}\mathbb{E}_{\mathbb{Q}_2^{(\boldsymbol{w}_1)}}\left[c_2+\cdots+\sup_{\mathbb{Q}_T^{(\boldsymbol{w}_{\leqslant T-1})}}\mathbb{E}_{\mathbb{Q}_T^{(\boldsymbol{w}_{\leqslant T-1})}}[c_T]\cdots\right]\right]\\ &= \sup_{\substack{\mathbb{Q}_1,\ldots,\mathbb{Q}_T^{(\cdot)}:\\ \mathbb{Q}_t^{(\boldsymbol{w}_{\leqslant t-1})}\in\mathcal{U}_t(\boldsymbol{w}_{\leqslant t-1}),\forall t,\forall \boldsymbol{w}_{\leqslant t-1}\in\hat{\mathcal{W}}_{\leqslant t-1}}}\mathbb{E}_{\mathbb{Q}_1}\left[c_1+\mathbb{E}_{\mathbb{Q}_2^{(\boldsymbol{w}_1)}}\left[c_2+\cdots+\mathbb{E}_{\mathbb{Q}_T^{(\boldsymbol{w}_{\leqslant T-1})}}[c_T]\cdots\right]\right]\\ &\geqslant \sup_{\mathbb{P}\in\mathcal{U}}\mathbb{E}_{\mathbb{P}_1}\left[c_1+\mathbb{E}_{\mathbb{P}_2}\left[c_2+\cdots+\mathbb{E}_{\mathbb{P}_T}[c_T]\cdots\right]\right]=\mathrm{DRO}(\boldsymbol{x}),\end{aligned}\tag{24}$$

where $\mathbb{P}_t=\mathbb{P}[\boldsymbol{w}_t\,|\,\boldsymbol{w}_{\leqslant t-1}]$ are the conditional distributions of $\boldsymbol{w}_t$ given history $\boldsymbol{w}_{\leqslant t-1}$ for all $t=2,\ldots,T$ and $\mathbb{P}_1=\mathbb{P}[\boldsymbol{w}_1]$ is marginal distribution of $\boldsymbol{w}_1$. Here, the inequality (24) holds as the 'inclusive' condition implies that for the optimal distribution $\mathbb{P}^*$ of the RHS of (24), there is some feasible family belonging to the ambiguity updating series, defined as $\tilde{\mathbb{Q}}_t^{(\boldsymbol{w}_{\leqslant t-1})}[\boldsymbol{w}_t=\cdot]=\mathbb{P}^*[\boldsymbol{w}_t=\cdot\,|\,\boldsymbol{w}_{\leqslant t-1}]$ for all $t=2,...,T$ and for all $\boldsymbol{w}_{\leqslant t-1}\in\hat{\mathcal{W}}_{\leqslant t-1}$.

We next show the relationship between (AURA) and (RO): Given any decision policy $\boldsymbol{x}(\cdot)\in\mathcal{X}$, for any time t and any history $\hat{\boldsymbol{w}}_{\leqslant t}\in\hat{\mathcal{W}}_{\leqslant t}$, we have, for every probability distribution $\mathbb{Q}_{\boldsymbol{w}_{t+1}}$, the inner expectation of problem-to-go for (AURA)

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_{t+1}|\hat{\boldsymbol{w}}_{\leqslant t}}}[F_{t+1}(\boldsymbol{x}_{\geqslant t+1};\hat{\boldsymbol{w}}_{\leqslant t},\boldsymbol{w}_{t+1})] &\leqslant \int_{\hat{\mathcal{W}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})}\sup_{\boldsymbol{w}_{t+1}\in\hat{\mathcal{W}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})}F_{t+1}(\boldsymbol{x}_{\geqslant t+1};\hat{\boldsymbol{w}}_{\leqslant t},\boldsymbol{w}_{t+1})\,d\mathbb{Q}_{\boldsymbol{w}_{t+1}}[\boldsymbol{w}_{t+1}]\\ &= \sup_{\boldsymbol{w}_{t+1}\in\hat{\mathcal{W}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})}F_{t+1}(\boldsymbol{x}_{\geqslant t+1};\hat{\boldsymbol{w}}_{\leqslant t},\boldsymbol{w}_{t+1})\end{aligned}$$

Since this holds for all $\mathbb{Q}_{\boldsymbol{w}_{t+1}}\in\mathcal{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$, we have the conclusion

$$\sup_{\mathbb{Q}_{\boldsymbol{w}_{t+1}}\in\mathcal{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})}\mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_{t+1}|\hat{\boldsymbol{w}}_{\leqslant t}}}[F_{t+1}(\boldsymbol{x}_{\geqslant t+1};\hat{\boldsymbol{w}}_{\leqslant t},\boldsymbol{w}_{t+1})]\leqslant\sup_{\boldsymbol{w}_{t+1}\in\hat{\mathcal{W}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})}F_{t+1}(\boldsymbol{x}_{\geqslant t+1};\hat{\boldsymbol{w}}_{\leqslant t},\boldsymbol{w}_{t+1})$$

The conclusion follows since the above inequality holds for every time t, every history $\hat{\boldsymbol{w}}_{\leqslant t}$, and every admissible policy $\boldsymbol{x}(\cdot)$.

Part ii. We prove this by contradiction. Denote the set of optimal solutions of (AURA) as

$$\mathcal{X}^*=\left\{\boldsymbol{x}^*\,\middle|\,c_0+\sup_{\mathbb{Q}_{\boldsymbol{w}_1}\in\mathcal{U}_1}\mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_1}}\left[F_1(\boldsymbol{x}_{\geqslant 1}^*;\boldsymbol{w}_1)\right]=\min_{\boldsymbol{x}\in\mathcal{X}}\left\{c_0+\sup_{\mathbb{Q}_{\boldsymbol{w}_1}\in\mathcal{U}_1}\mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_1}}\left[F_1(\boldsymbol{x}_{\geqslant 1};\boldsymbol{w}_1)\right]\right\}\right\}.$$

Assume that all $\boldsymbol{x}^* \in \mathcal{X}^*$ are not time consistent. For each $\boldsymbol{x}^* \in \mathcal{X}^*$, there is some largest time t, realization $\boldsymbol{w}_{<t} = \hat{\boldsymbol{w}}_{<t}$ (that depends on the solution $\boldsymbol{x}^*$) and policy-to-go $\check{\boldsymbol{x}}_{\geqslant t}$ such that

$$F_t(\boldsymbol{x}^*_{\geqslant t}; \hat{\boldsymbol{w}}_{<t}) > F_t(\check{\boldsymbol{x}}_{\geqslant t}; \hat{\boldsymbol{w}}_{<t}). \tag{25}$$

Let $\bar{t}$ be the largest such time t across all $\boldsymbol{x}^* \in \mathcal{X}^*$, and let $\overline{\mathcal{W}}$ be the set of all histories $\hat{\boldsymbol{w}}_{<\bar{t}}$ for which there exists some policy $\boldsymbol{x}^* \in \mathcal{X}^*$ such that (25) holds. For each $\hat{\boldsymbol{w}}_{<\bar{t}} \in \overline{\mathcal{W}}$, pick some

$$\check{\boldsymbol{x}}^{(\hat{\boldsymbol{w}}_{<\bar{t}})}_{\geqslant \bar{t}}(\cdot) \in \operatorname*{arg\,min}_{\boldsymbol{x}_{\geqslant \bar{t}}(\cdot)} F_{\bar{t}}(\boldsymbol{x}_{\geqslant \bar{t}}(\cdot); \hat{\boldsymbol{w}}_{<\bar{t}}).$$

Now, construct the set of solutions $\overline{\mathcal{X}}$ by replacing the solutions on history $\hat{\boldsymbol{w}}_{<\bar{t}}$ with $\check{\boldsymbol{x}}^{(\hat{\boldsymbol{w}}_{<\bar{t}})}_{\geqslant \bar{t}}(\cdot)$ for all $\boldsymbol{x}^* \in \mathcal{X}^*$. Specifically,

$$\overline{\mathcal{X}} := \left\{ \boldsymbol{x} \;\middle|\; \exists \boldsymbol{x}^* \in \mathcal{X}^* : \boldsymbol{x}(\hat{\boldsymbol{w}}_{<\bar{t}}; \cdot) = \begin{cases} \check{\boldsymbol{x}}^{(\hat{\boldsymbol{w}}_{<\bar{t}})}_{\geqslant \bar{t}}(\cdot) & \text{if } \hat{\boldsymbol{w}}_{<\bar{t}} \in \overline{\mathcal{W}} \\ \boldsymbol{x}^*(\hat{\boldsymbol{w}}_{<\bar{t}}; \cdot) & \text{otherwise} \end{cases} \right\}.$$

By definition, these solutions are feasible. Notice that if $\overline{\mathcal{X}} \nsubseteq \mathcal{X}^*$, in other words, there is some $\check{\boldsymbol{x}} \in \overline{\mathcal{X}}$ such that $\check{\boldsymbol{x}} \notin \mathcal{X}^*$ and some $\boldsymbol{x}^* \in \mathcal{X}^*$ such that

$$\check{\boldsymbol{x}}(\hat{\boldsymbol{w}}_{<\bar{t}}; \cdot) = \begin{cases} \check{\boldsymbol{x}}^{(\hat{\boldsymbol{w}}_{<\bar{t}})}_{\geqslant \bar{t}}(\cdot) & \text{if } \hat{\boldsymbol{w}}_{<\bar{t}} \in \overline{\mathcal{W}} \\ \boldsymbol{x}^*(\hat{\boldsymbol{w}}_{<\bar{t}}; \cdot) & \text{otherwise} \end{cases},$$

then $F_{\bar{t}}(\boldsymbol{x}^*_{\geqslant \bar{t}}; \hat{\boldsymbol{w}}_{<\bar{t}}) \geqslant F_{\bar{t}}(\check{\boldsymbol{x}}_{\geqslant \bar{t}}; \hat{\boldsymbol{w}}_{<\bar{t}}), \forall \hat{\boldsymbol{w}}_{<\bar{t}} \in \hat{\mathcal{W}}_{<\bar{t}}$. Consequently, it also holds that

$$\sup_{\mathbb{Q}_{\bar{t}}} \mathbb{E}_{\mathbb{Q}_{\bar{t}}} F_{\bar{t}}(\boldsymbol{x}^*_{\geqslant \bar{t}}; \hat{\boldsymbol{w}}_{<\bar{t}}) \geqslant \sup_{\mathbb{Q}_{\bar{t}}} \mathbb{E}_{\mathbb{Q}_{\bar{t}}} F_{\bar{t}}(\check{\boldsymbol{x}}_{\geqslant \bar{t}}; \hat{\boldsymbol{w}}_{<\bar{t}}).$$

Since $\boldsymbol{x}^*$ and $\check{\boldsymbol{x}}$ coincide on all decisions prior to time $\bar{t}$, thus $\mathrm{AURA}(\boldsymbol{x}^*) \geqslant \mathrm{AURA}(\check{\boldsymbol{x}})$. But $\check{\boldsymbol{x}} \notin \mathcal{X}^*$, thus $\mathrm{AURA}(\boldsymbol{x}^*) > \mathrm{AURA}(\check{\boldsymbol{x}})$, which contradicts the optimality of $\boldsymbol{x}^*$.

Hence, $\overline{\mathcal{X}} \subseteq \mathcal{X}^*$. We can repeat the above process with $\overline{\mathcal{X}}$ instead of $\mathcal{X}^*$, but now, the largest t must be strictly smaller than $\bar{t}$, due to the assumed maximality of $\bar{t}$. This either arrives at a contradiction if $\overline{\mathcal{X}} \nsubseteq \mathcal{X}^*$, or until $\bar{t} = 0$, at which point, all $\boldsymbol{x} \in \overline{\mathcal{X}}$ are time consistent.

Part iii. ($\Leftarrow$): Suppose (DRO) = (AURA). By part ii., there exists some time consistent optimal solution to (AURA), $\boldsymbol{x}^*_{\mathrm{AURA}}$. Then $\boldsymbol{x}^*_{\mathrm{AURA}}$ is also optimal for (DRO) because $\mathrm{DRO}(\boldsymbol{x}^*_{\mathrm{AURA}}) \leqslant \mathrm{AURA}(\boldsymbol{x}^*_{\mathrm{AURA}}) = \mathrm{DRO}(\boldsymbol{x}^*_{\mathrm{DRO}}) \leqslant \mathrm{DRO}(\boldsymbol{x}^*_{\mathrm{AURA}})$, where the first inequality is due to part i., and the second holds by the optimality of $\boldsymbol{x}^*_{\mathrm{DRO}}$ for (DRO). Hence, all inequalities must hold at equality.

For the first inequality to hold at equality, this means that (24) must hold with equality for $\boldsymbol{x}^*_{\mathrm{AURA}}$. We claim that (24) holds with equality if and only if there is some optimal family $\left\{ \mathbb{Q}^*_1, \ldots, \mathbb{Q}^{*(\cdot)}_T : \mathbb{Q}^{*(\boldsymbol{w}_{\leqslant t-1})}_t \in \mathcal{U}_t(\boldsymbol{w}_{\leqslant t-1}), \forall t, \forall \boldsymbol{w}_{\leqslant t-1} \in \hat{\mathcal{W}}_{\leqslant t-1} \right\}$ and some optimal $\mathbb{P} \in \mathcal{U}$ with $\mathbb{Q}^{*(\boldsymbol{w}_{\leqslant t-1})}_t[\boldsymbol{w}_t = \cdot] = \mathbb{P}[\boldsymbol{w}_t = \cdot \,|\, \boldsymbol{w}_{\leqslant t-1}]$ for all t and for all $\boldsymbol{w}_{\leqslant t-1} \in \hat{\mathcal{W}}_{\leqslant t-1}$ except on a set of measure 0. Indeed, if such a common $\mathbb{P}$ exists, then both sides of (24) attain the same objective value. On

the other hand, if (24) holds with equality, then there is some optimal $\mathbb{P}$ that achieves the same objective as the LHS, itself being feasible for the LHS problem.

Thus, such a $\mathbb{P}$ leads $\boldsymbol{x}^*_{\text{AURA}}$ to be time consistent for (DRO), since $\mathbb{Q}^*$ being optimal for (AURA) under time consistent solution $\boldsymbol{x}^*_{\text{AURA}}$ implies that it is also the certificate for the optimality of $\boldsymbol{x}^*_{\text{AURA}}$ for every problem-to-go.

($\Rightarrow$): We prove this by contradiction. Suppose that (DRO) $<$ (AURA). Let $\boldsymbol{x}^*_{\text{DRO}}$ be a time consistent optimal solution of (DRO), and let $\mathbb{P}^*$ denote the worst-case distribution attained under this solution. If $\text{AURA}(\boldsymbol{x}^*_{\text{DRO}}) = \text{DRO}(\boldsymbol{x}^*_{\text{DRO}})$, then (DRO) $<$ (AURA) $\leqslant \text{AURA}(\boldsymbol{x}^*_{\text{DRO}}) = \text{DRO}(\boldsymbol{x}^*_{\text{DRO}}) =$ (DRO), contradicting our premise. Thus, by part i., $\text{AURA}(\boldsymbol{x}^*_{\text{DRO}}) > \text{DRO}(\boldsymbol{x}^*_{\text{DRO}})$ and there exists some optimal ambiguity updating series $\left\{\mathbb{Q}^*_1, \ldots, \mathbb{Q}^{*(\cdot)}_T \;\middle|\; \mathbb{Q}^{*(\boldsymbol{w}_{\leqslant t-1})}_t \in \mathcal{U}_t(\boldsymbol{w}_{\leqslant t-1}), \forall t, \forall \boldsymbol{w}_{\leqslant t-1} \in \hat{\mathcal{W}}_{\leqslant t-1}\right\}$ for $\text{AURA}(\boldsymbol{x}^*_{\text{DRO}})$ such that

$$c_0 + \mathbb{E}_{\mathbb{Q}^*_1}\left[c_1 + \ldots + \mathbb{E}_{\mathbb{Q}^{*(\boldsymbol{w}_{\leqslant T-1})}_T}[c_T]\right] > c_0 + \mathbb{E}_{\mathbb{P}^*}\left[c_1 + \ldots + c_T\right]. \tag{26}$$

Consider the distributions $\overline{\mathbb{Q}}^*_\tau \in \mathcal{M}\left(\bigtimes_{t=1}^{\tau} \mathcal{W}_t, \mathbb{R}\right)$, $\tau = 1, \ldots, T$, defined by $\overline{\mathbb{Q}}^*_\tau[\boldsymbol{w}_1, \ldots, \boldsymbol{w}_\tau] = \overline{\mathbb{Q}}^*_{\tau-1}[\boldsymbol{w}_1, \ldots, \boldsymbol{w}_{\tau-1}]\,\mathbb{Q}^{*(\boldsymbol{w}_{\leqslant \tau-1})}_\tau$. By (26), there is some smallest time $\bar{t} \geqslant 1$ where $\overline{\mathbb{Q}}^*_\tau[\boldsymbol{w}_1, \ldots, \boldsymbol{w}_\tau]$ and marginal distribution $\mathbb{P}^*[\boldsymbol{w}_1, \ldots, \boldsymbol{w}_\tau]$ differ on a set of measure zero for all $\tau < \bar{t}$; and a set $\overline{\mathcal{W}}_{\bar{t}} \subseteq \hat{\mathcal{W}}_{<\bar{t}}$ of measure larger than 0 on which

$$\mathbb{E}_{\mathbb{Q}^{*(\boldsymbol{w}_{\leqslant \bar{t}-1})}_{\bar{t}}}\left[c_{\bar{t}} + \ldots + \mathbb{E}_{\mathbb{Q}^{*(\boldsymbol{w}_{\leqslant T-1})}_T}[c_T]\right] > \mathbb{E}_{\mathbb{P}^*_{\bar{t}}|_{\boldsymbol{w}_{\leqslant \bar{t}-1}}}\left[c_{\bar{t}} + \ldots + c_T\right], \qquad \forall \boldsymbol{w}_{\leqslant \bar{t}-1} \in \overline{\mathcal{W}}_{\bar{t}}. \tag{27}$$

By (27), for each $\boldsymbol{w}_{\leqslant \bar{t}-1}$, there must exist a set $\overline{\mathcal{W}}(\boldsymbol{w}_{\leqslant \bar{t}-1}) \subseteq \hat{\mathcal{W}}_{\bar{t}}$ of measure larger than 0, where

$$\mathbb{E}_{\mathbb{Q}^{*(\boldsymbol{w}_{\leqslant \bar{t}})}_{\bar{t}+1}}\left[c_{\bar{t}+1} + \ldots + \mathbb{E}_{\mathbb{Q}^{*(\boldsymbol{w}_{\leqslant T-1})}_T}[c_T]\right] > \mathbb{E}_{\mathbb{P}^*_{\bar{t}+1}|_{\boldsymbol{w}_{\leqslant \bar{t}}}}\left[c_{\bar{t}+1} + \ldots + c_T\right], \qquad \forall \boldsymbol{w}_{\leqslant \bar{t}} \in \overline{\mathcal{W}}(\boldsymbol{w}_{\leqslant \bar{t}-1}) \tag{28}$$

Combining (27) and (28), there is thus some set $\overline{\mathcal{W}}_{\bar{t}+1} \subseteq \hat{\mathcal{W}}_{<\bar{t}+1}$ of measure larger than 0, where

$$\mathbb{E}_{\mathbb{Q}^{*(\boldsymbol{w}_{\leqslant \bar{t}})}_{\bar{t}+1}}\left[c_{\bar{t}+1} + \ldots + \mathbb{E}_{\mathbb{Q}^{*(\boldsymbol{w}_{\leqslant T-1})}_T}[c_T]\right] > \mathbb{E}_{\mathbb{P}^*_{\bar{t}+1}|_{\boldsymbol{w}_{\leqslant \bar{t}}}}\left[c_{\bar{t}+1} + \ldots + c_T\right], \qquad \forall \boldsymbol{w}_{\leqslant \bar{t}} \in \overline{\mathcal{W}}_{\bar{t}+1} \tag{29}$$

By induction, we can find a set $\overline{\mathcal{W}}_T \subseteq \hat{\mathcal{W}}_{<T}$ of measure larger than 0, where $\mathbb{E}_{\mathbb{Q}^{*(\boldsymbol{w}_{<T})}_T}[c_T] > \mathbb{E}_{\mathbb{P}^*_T|_{\boldsymbol{w}_{<T}}}[c_T]$, $\forall \boldsymbol{w}_{<T} \in \overline{\mathcal{W}}_T$. Now, by the definition of 'traceable', $\mathbb{Q}^{*(\boldsymbol{w}_{<T})}_T$ must originate from some $\overline{\mathbb{P}} \in \mathcal{U}$ in the parent uncertainty set. Hence,

$$\mathbb{P}^*_T|_{\boldsymbol{w}_{<T}} \notin \operatorname*{argsup}_{\mathbb{P}_T\,:\,\exists \mathbb{P} \in \mathcal{U}, \mathbb{P}_T = \mathbb{P}|_{\boldsymbol{w}_{<T}}} \mathbb{E}_{\mathbb{P}_T}[c_T] \tag{30}$$

Fix any $\boldsymbol{w}_{<T} \in \overline{\mathcal{W}}_T$. Denote $V_T(\boldsymbol{w}_{\leqslant T}) = \min_{\boldsymbol{x}_T \in \mathcal{X}_T} c_T(\boldsymbol{x}_T)$ and the $(T-1)^{\text{th}}$ period problem-to-go as

$$V_{T-1}(\boldsymbol{w}_{<T}) = \min_{\boldsymbol{x}_{T-1}(\cdot) \in \mathcal{X}_{T-1}} \left\{ c_{T-1}(\boldsymbol{x}_{T-1}) + \sup_{\mathbb{P}_T\,:\,\exists \mathbb{P} \in \mathcal{U}, \mathbb{P}_T = \mathbb{P}|_{\boldsymbol{w}_{<T}}} \mathbb{E}_{\mathbb{P}_T}[V_T(\boldsymbol{w}_{\leqslant T})] \right\}. \tag{31}$$

As $\boldsymbol{x}^*_{\text{DRO}}$ is time consistent, $\boldsymbol{x}^*_{\text{DRO},T}(\boldsymbol{w}_{<T},\cdot) \in \arg\min_{\boldsymbol{x}_T} c_T(\boldsymbol{x}_T)$. But (30) implies that

$$V_{T-1}(\boldsymbol{w}_{<T}) > c_{T-1}(\boldsymbol{x}^*_{\text{DRO},T-1}) + \mathbb{E}_{\mathbb{P}^*_T}[V_T(\boldsymbol{w}_{\leqslant T})],$$

contradicting the optimality of $\boldsymbol{x}^*_{\text{DRO},T-1}$ for (31).

A.3. Proposition 2

Fix time t and history $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_t$. For distributions $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$, we aim to show that $\mathbb{Q}^\alpha = \alpha\mathbb{Q}_1 + (1-\alpha)\mathbb{Q}_2$, lies in (LAUS), $\mathbb{Q}^\alpha \in \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$, for any $\alpha \in (0,1)$. It suffices to seek some parent distribution $\mathbb{P} \in \mathcal{U}$, whose conditional distribution is $\mathbb{Q}^\alpha$.

By definition of $\mathbb{Q}_1$ and $\mathbb{Q}_2$, there exist parent distributions $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{U}$ such that, $\mathbb{Q}_1[\cdot] = \mathbb{P}_1[\cdot \mid \hat{\boldsymbol{w}}_{\leqslant t}]$ and $\mathbb{Q}_2[\cdot] = \mathbb{P}_2[\cdot \mid \hat{\boldsymbol{w}}_{\leqslant t}]$. Then we have

$$\begin{aligned}
\mathbb{Q}^\alpha[\cdot] &= \alpha\mathbb{P}_1[\cdot \mid \hat{\boldsymbol{w}}_{\leqslant t}] + (1-\alpha)\mathbb{P}_2[\cdot \mid \hat{\boldsymbol{w}}_{\leqslant t}] = \alpha\frac{\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]}{\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]} + (1-\alpha)\frac{\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]}{\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]} \\
&= \frac{\alpha\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}] + (1-\alpha)\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]}{\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]} \\
&= \frac{\alpha\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}] + (1-\alpha)\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]}{\alpha\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}] + (1-\alpha)\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]} \\
&= \frac{\alpha\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}] + (1-\alpha)\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]}{(1-\alpha)\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}] + \alpha\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]} \Big/ \frac{\alpha\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}] + (1-\alpha)\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]}{(1-\alpha)\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}] + \alpha\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]}.
\end{aligned} \tag{32}$$

Letting $\beta = \frac{\alpha\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]}{(1-\alpha)\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]+\alpha\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]}$, hence $1-\beta = \frac{(1-\alpha)\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]}{(1-\alpha)\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}]+\alpha\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]}$. The RHS of (32) becomes

$$\mathbb{Q}^\alpha[\cdot] = \frac{\beta\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t},\cdot] + (1-\beta)\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]}{\beta\mathbb{P}_1[\hat{\boldsymbol{w}}_{\leqslant t}] + (1-\beta)\mathbb{P}_2[\hat{\boldsymbol{w}}_{\leqslant t}]} = \frac{\mathbb{P}^\alpha[\hat{\boldsymbol{w}}_{\leqslant t},\cdot]}{\mathbb{P}^\alpha[\hat{\boldsymbol{w}}_{\leqslant t}]},$$

where $\mathbb{P}^\alpha = \beta\mathbb{P}_1 + (1-\beta)\mathbb{P}_2$. Moreover, since $\mathbb{P}^\alpha$ is a linear combination of $\mathbb{P}_1$ and $\mathbb{P}_2$, by convexity of $\mathcal{U}$, $\mathbb{P}^\alpha \in \mathcal{U}$. Thus, $\mathbb{P}^\alpha$ is the parent distribution corresponding to $\mathbb{Q}^\alpha$ that we seek.

A.4. Proposition 3

First of all, by definition, for all $t = 1,...,T-1$ and $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$, $\mathscr{J}_{t+1}(\boldsymbol{w}_{\leqslant t}) \subseteq \mathcal{M}(\hat{\mathcal{W}}_{t+1}, \mathbb{R})$. Given any $\mathbb{P} \in \mathcal{U}$, for all $t = 1,...,T$ and histories $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$, the distribution $\mathbb{Q}[\boldsymbol{w}_{t+1}] := {}^{\mathbb{P}[\boldsymbol{w}_{\leqslant t}=\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}=\boldsymbol{w}_{t+1}]}\big/_{\mathbb{P}[\boldsymbol{w}_{\leqslant t}=\hat{\boldsymbol{w}}_{\leqslant t}]} = \mathbb{P}[\boldsymbol{w}_{t+1} = \boldsymbol{w}_{t+1} | \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}]$ lies in $\mathcal{U}_{t+1}$. On the other hand, for all $t = 1,...,T$ and histories $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$, for any $\mathbb{Q}_{t+1} \in \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$, by definition, there exists $\mathbb{P}$ such that $\mathbb{E}_{\mathbb{P}}[\boldsymbol{g}(\boldsymbol{w})] = 0$ and $\mathbb{P}[\boldsymbol{w} \in \mathcal{W}] = 1$, that is $\mathbb{P} \in \mathcal{U}$ and ${}^{\mathbb{P}[\boldsymbol{w}_{\leqslant t}=\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}=\boldsymbol{w}_{t+1}]}\big/_{\mathbb{P}[\boldsymbol{w}_{\leqslant t}=\hat{\boldsymbol{w}}_{\leqslant t}]} = \mathbb{Q}[\boldsymbol{w}_{t+1}]$.

A.5. Lemma 1

We eliminate $\mathbb{Q}_{t+1}$ by writing $\sup_{\mathbb{Q}_{t+1} \in \mathscr{J}_{t+1}(\boldsymbol{w}_{\leqslant t})} \mathbb{E}_{\mathbb{Q}_{t+1}}\left[F_{t+1}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})\right]$ as

$$\begin{aligned}
&\sup_{q \geqslant 0} \int_{\hat{\mathcal{W}}_{t+1}} \frac{\int q\left(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}, \boldsymbol{w}_{>t+1}\right) d\boldsymbol{w}_{>t+1}}{\int q\left(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{>t}\right) d\boldsymbol{w}_{>t}} F_{t+1}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})\, d\boldsymbol{w}_{t+1} \\
&\text{s.t.} \quad \int q(\boldsymbol{w})\boldsymbol{g}(\boldsymbol{w})\, d\boldsymbol{w} = \boldsymbol{0} \quad \text{and} \quad \int q(\boldsymbol{w})\, d\boldsymbol{w} = 1 \quad \text{and} \quad \int q\left(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{>t}\right) d\boldsymbol{w}_{>t} > 0
\end{aligned} \tag{33}$$

where $q \sim \mathbb{P}$ represents the density of $\mathbb{P}$. By a change of variables, $r = \left(\int q\left(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{>t}\right) d\boldsymbol{w}_{>t} \right)^{-1}$ and $\tilde{q} = qr$, (33) becomes

$$\begin{aligned}
\sup_{\tilde{q}, r \geqslant 0} \quad & \int_{\hat{\mathcal{W}}_{t+1}} \int \tilde{q}\left(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}, \boldsymbol{w}_{>t+1}\right) d\boldsymbol{w}_{>t+1} F_{t+1}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})\, d\boldsymbol{w}_{t+1} && (34)\\
\text{s.t.} \quad & \int \tilde{q}(\boldsymbol{w}) \boldsymbol{g}(\boldsymbol{w})\, d\boldsymbol{w} = \mathbf{0} && (\boldsymbol{\lambda})\\
& \int \tilde{q}(\boldsymbol{w})\, d\boldsymbol{w} = r && (\eta)\\
& \int \tilde{q}\left(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{>t}\right) d\boldsymbol{w}_{>t} = 1 && (\theta)
\end{aligned}$$

Its dual forms an upper bound to (34):

$$\begin{aligned}
\inf_{\theta \geqslant 0, \boldsymbol{\lambda}, \theta} \quad & \theta && (35)\\
\text{s.t.} \quad & \boldsymbol{\lambda}^\top \boldsymbol{g}(\boldsymbol{w}) + \eta + \theta \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right] \geqslant F_{t+1}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right], \quad \forall \boldsymbol{w} \in \mathcal{W} && (\tilde{q})
\end{aligned}$$

which is equivalent to $D_t(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t})$ as η must take value 0 at optimality.

A.6. Theorem 2

We first prove weak duality, *i.e.*, (7) $\leqslant$ (AURA-RC). Define problems (S_τ) for $\tau = 0, \ldots, T-1$:

$$\begin{aligned}
\inf_{\substack{\boldsymbol{\lambda}_{\leqslant \tau} \\ \theta_{\leqslant \tau}}} \quad & \theta^{\hat{\boldsymbol{w}}_{\leqslant 0}} && (S_\tau)\\
\text{s.t.} \quad & {\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant \tau'}}}^\top \boldsymbol{g}(\boldsymbol{w}) + \theta^{\hat{\boldsymbol{w}}_{\leqslant \tau'}} \mathbf{1}\left[\boldsymbol{w}_{\leqslant \tau'} = \hat{\boldsymbol{w}}_{\leqslant \tau'}\right] \geqslant \left(c_{\tau'+1}(\boldsymbol{x}_{\tau'+1}) + \theta^{(\hat{\boldsymbol{w}}_{\leqslant \tau'}, \boldsymbol{w}_{\tau'+1})}\right) \mathbf{1}\left[\boldsymbol{w}_{\leqslant \tau'} = \hat{\boldsymbol{w}}_{\leqslant \tau'}\right], \\
& \qquad \forall \tau' = 0, ..., \tau-1;\ \forall \boldsymbol{w} \in \mathcal{W}, \forall \hat{\boldsymbol{w}}_{\leqslant \tau'} \in \hat{\mathcal{W}}_{\leqslant \tau'} \\
& {\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant \tau}}}^\top \boldsymbol{g}(\boldsymbol{w}) + \theta^{\hat{\boldsymbol{w}}_{\leqslant \tau}} \mathbf{1}\left[\boldsymbol{w}_{\leqslant \tau} = \hat{\boldsymbol{w}}_{\leqslant \tau}\right] \geqslant F_{\tau+1}(\boldsymbol{x}_{\geqslant \tau+1}; \hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1}) \mathbf{1}\left[\boldsymbol{w}_{\leqslant \tau} = \hat{\boldsymbol{w}}_{\leqslant \tau}\right], && (36)\\
& \qquad \forall \boldsymbol{w} \in \mathcal{W}, \forall \hat{\boldsymbol{w}}_{\leqslant \tau} \in \hat{\mathcal{W}}_{\leqslant \tau}
\end{aligned}$$

For brevity, we write $\theta_{\leqslant \tau}$ for $\left(\theta^{\hat{\boldsymbol{w}}_{\leqslant 0}}, \theta^{\hat{\boldsymbol{w}}_{\leqslant 1}}, \ldots, \theta^{\hat{\boldsymbol{w}}_{\leqslant \tau}} : \hat{\boldsymbol{w}}_{\leqslant \tau} \in \hat{\mathcal{W}}_{\leqslant \tau}\right)$ and similarly $\boldsymbol{\lambda}_{\leqslant \tau}$. When $\tau = 0$,

$$\begin{aligned}
(S_0) = \inf_{\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant 0}}, \theta^{\hat{\boldsymbol{w}}_{\leqslant 0}}} \quad & \theta^{\hat{\boldsymbol{w}}_{\leqslant 0}} && (37)\\
\text{s.t.} \quad & {\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant 0}}}^\top \boldsymbol{g}(\boldsymbol{w}) + \theta^{\hat{\boldsymbol{w}}_{\leqslant 0}} \geqslant F_1(\boldsymbol{x}_{\geqslant 1}; \boldsymbol{w}_1), \qquad \forall \boldsymbol{w} \in \mathcal{W}
\end{aligned}$$

Hence, (7) $\leqslant (S_0)$ by Lemma 1 applied on $t = 0$. When $\tau = T-1$, $F_T(\boldsymbol{x}_T; \hat{\boldsymbol{w}}_{T-1}, \boldsymbol{w}_T) = c_T(\boldsymbol{x}_T)$ and thus $(S_{T-1}) =$ (AURA-RC). We attempt to prove that $(S_\tau) \leqslant (S_{\tau+1})$ for all $\tau = 0, \ldots, T-2$, which by induction implies (7) $\leqslant$ (AURA-RC). Now, for all $(\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1}) \in \hat{\mathcal{W}}_{\tau+1}$, by Lemma 1,

$$\begin{aligned}
& F_{\tau+1}(\boldsymbol{x}_{\geqslant \tau+1}; \hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1}) - c_{\tau+1}(\boldsymbol{x}_{\tau+1}) \\
= & \sup_{\mathbb{Q}_{\tau+2} \in \mathscr{J}_{\tau+2}} \mathbb{E}_{\mathbb{Q}_{\tau+2}}\left[F_{\tau+2}(\boldsymbol{x}_{\geqslant \tau+2}; (\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1}), \boldsymbol{w}_{\tau+2})\right] \\
\leqslant & \inf_{\substack{\boldsymbol{\lambda}^{(\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1})} \\ \theta^{(\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1})}}} \quad \theta^{(\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1})} && (38)\\
& \text{s.t.} \quad {\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1}}}^\top \boldsymbol{g}(\bar{\boldsymbol{w}}) + \theta^{(\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1})} \mathbf{1}\left[\bar{\boldsymbol{w}}_{\leqslant \tau+1} = (\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1})\right] \\
& \qquad \geqslant F_{\tau+2}(\boldsymbol{x}_{\geqslant \tau+2}; (\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1}), \bar{\boldsymbol{w}}_{\tau+2})\, \mathbf{1}\left[\bar{\boldsymbol{w}}_{\leqslant \tau+1} = (\hat{\boldsymbol{w}}_{\leqslant \tau}, \boldsymbol{w}_{\tau+1})\right], \qquad \forall \bar{\boldsymbol{w}} \in \mathcal{W}
\end{aligned}$$

Consider (36), for all $\hat{\boldsymbol{w}}_{\leqslant\tau}$, as the system of constraints:

$$\begin{cases} {\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant\tau}}}^{\top}\boldsymbol{g}(\boldsymbol{w})+\theta^{\hat{\boldsymbol{w}}_{\leqslant\tau}}\mathbf{1}\left[\boldsymbol{w}_{\leqslant\tau}=\hat{\boldsymbol{w}}_{\leqslant\tau}\right] \\ \qquad\geqslant\left(c_{\tau+1}(\boldsymbol{x}_{\tau+1})+\inf\limits_{\substack{\boldsymbol{\lambda}^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}\\ \theta^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}}}\theta^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}\right)\mathbf{1}\left[\boldsymbol{w}_{\leqslant\tau}=\hat{\boldsymbol{w}}_{\leqslant\tau}\right], \qquad \forall\boldsymbol{w}\in\mathcal{W} \\ {\boldsymbol{\lambda}^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}}^{\top}\boldsymbol{g}(\bar{\boldsymbol{w}})+\theta^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}\mathbf{1}\left[\bar{\boldsymbol{w}}_{\leqslant\tau+1}=(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})\right] \\ \qquad\geqslant F_{\tau+2}(\boldsymbol{x}_{\geqslant\tau+2};(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1}),\bar{\boldsymbol{w}}_{\tau+2})\,\mathbf{1}\left[\bar{\boldsymbol{w}}_{\leqslant\tau+1}=(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})\right], \quad \forall\boldsymbol{w}_{\tau+1}\in\hat{\mathcal{W}}_{\tau+1},\forall\bar{\boldsymbol{w}}\in\mathcal{W} \end{cases} \tag{39}$$

By (38), if the pair $(\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant\tau}},\theta^{\hat{\boldsymbol{w}}_{\leqslant\tau}})$ fulfills (39), it also satisfies (36). Thus the feasibility set for $(\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant\tau}},\theta^{\hat{\boldsymbol{w}}_{\leqslant\tau}})$ admitted by (39) is contained within that for (36), yielding an upper bound for (S_τ):

$$\inf_{\boldsymbol{\lambda}_{\leqslant\tau},\theta_{\leqslant\tau}} \quad \theta^{\hat{\boldsymbol{w}}_{\leqslant 0}} \tag{40}$$

$$\text{s.t.} \quad {\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant\tau'}}}^{\top}\boldsymbol{g}(\boldsymbol{w})+\theta^{\hat{\boldsymbol{w}}_{\leqslant\tau'}}\mathbf{1}\left[\boldsymbol{w}_{\leqslant\tau'}=\hat{\boldsymbol{w}}_{\leqslant\tau'}\right]\geqslant\left(c_{\tau'+1}(\boldsymbol{x}_{\tau'+1})+\theta^{(\hat{\boldsymbol{w}}_{\leqslant\tau'},\boldsymbol{w}_{\tau'+1})}\right)\mathbf{1}\left[\boldsymbol{w}_{\leqslant\tau'}=\hat{\boldsymbol{w}}_{\leqslant\tau'}\right],$$
$$\forall\tau'=0,...,\tau-1;\ \forall\boldsymbol{w}\in\mathcal{W},\forall\hat{\boldsymbol{w}}_{\leqslant\tau'}\in\hat{\mathcal{W}}_{\leqslant\tau'}$$

$${\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant\tau}}}^{\top}\boldsymbol{g}(\boldsymbol{w})+\theta^{\hat{\boldsymbol{w}}_{\leqslant\tau}}\mathbf{1}\left[\boldsymbol{w}_{\leqslant\tau}=\hat{\boldsymbol{w}}_{\leqslant\tau}\right] \tag{41}$$
$$\geqslant\left(c_{\tau+1}(\boldsymbol{x}_{\tau+1})+\inf_{\substack{\boldsymbol{\lambda}^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}\\ \theta^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}}}\theta^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}\right)\mathbf{1}\left[\boldsymbol{w}_{\leqslant\tau}=\hat{\boldsymbol{w}}_{\leqslant\tau}\right], \quad \forall\boldsymbol{w}\in\mathcal{W},\forall\hat{\boldsymbol{w}}_{\leqslant\tau}\in\hat{\mathcal{W}}_{\leqslant\tau'}$$

$${\boldsymbol{\lambda}^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}}^{\top}\boldsymbol{g}(\bar{\boldsymbol{w}})+\theta^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})}\mathbf{1}\left[\bar{\boldsymbol{w}}_{\leqslant\tau+1}=(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})\right] \tag{42}$$
$$\geqslant F_{\tau+2}(\boldsymbol{x}_{\geqslant\tau+2};(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1}),\bar{\boldsymbol{w}}_{\tau+2})\,\mathbf{1}\left[\bar{\boldsymbol{w}}_{\leqslant\tau+1}=(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})\right],$$
$$\forall\boldsymbol{w}_{\tau+1}\in\hat{\mathcal{W}}_{\tau+1},\forall\bar{\boldsymbol{w}}\in\mathcal{W},\forall\hat{\boldsymbol{w}}_{\leqslant\tau}\in\hat{\mathcal{W}}_{\leqslant\tau}$$

(41) can be relaxed to finding feasible $(\boldsymbol{\lambda}^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})},\theta^{(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})})$, while in (42), the composition of $\hat{\boldsymbol{w}}_{\leqslant\tau}\in\hat{\mathcal{W}}_{\leqslant\tau}$ and $\boldsymbol{w}_{\tau+1}\in\hat{\mathcal{W}}_{\tau+1}|\hat{\boldsymbol{w}}_{\leqslant\tau}$ is the same as $\hat{\boldsymbol{w}}_{\leqslant\tau+1}\in\hat{\mathcal{W}}_{\leqslant\tau+1}$. Thus, the change of variables $(\hat{\boldsymbol{w}}_{\leqslant\tau},\boldsymbol{w}_{\tau+1})\mapsto\hat{\boldsymbol{w}}_{\leqslant\tau+1}$ yields (40) $=(S_{\tau+1})$. This completes the induction.

To prove equivalence, we need to prove strong duality. Assumption 4 is the distributionally robust version of the "inf-compactness" condition proposed in Shapiro (2001). If $\hat{\mathcal{W}}_t\subseteq\mathbb{R}^{u_t}$ (not a singleton) is compact, then $\mathscr{J}_t$ is weakly compact, a consequence of Prokhorov's theorem (if $\hat{\mathcal{W}}_t$ is compact, then $\mathscr{J}_t$ is tight and therefore, weakly compact). From Xu et al. (2018), Proposition 2.3, $\sup_{\mathbb{Q}_t\in\mathscr{J}_t}\mathbb{E}_{\mathbb{Q}_t}[c_t(\boldsymbol{x}_t)]$ is lower semicontinuous in $\boldsymbol{x}_t$. From Shapiro (2001), Proposition 2.4, we know that the dual of this model satisfies strong duality. Since $c_{t-1}(\boldsymbol{x}_{t-1})$ is lower semicontinuous, $F_{t-1}=c_{t-1}(\boldsymbol{x}_{t-1})+\sup_{\mathbb{Q}_t\in\mathscr{J}_t}\mathbb{E}_{\mathbb{Q}_t}[c_t(\boldsymbol{x}_t)]$ is lower semicontinuous. It is clear that $\nu(\boldsymbol{x})=c_0(\boldsymbol{x}_0)+\sup\limits_{\mathbb{Q}_{\boldsymbol{w}_1}\in\mathscr{J}_1}\mathbb{E}_{\mathbb{Q}_{\boldsymbol{w}_1}}\left[F_1(\boldsymbol{x}_{\geqslant 1};\boldsymbol{w}_1)\right]$ is lower semicontinuous and strong duality exists.

A.7. Theorem 3i.

We prove this by contradiction. Suppose there exist some time t and a history $\hat{\boldsymbol{w}}_{\leqslant t}$ such that there exists a distribution $\mathbb{Q}_{t+1}\in(\text{LAUS})_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$ but $\mathbb{Q}_{t+1}\notin\mathscr{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$. By the definition of (LAUS), there exists a parent distribution $\mathbb{P}\in\mathcal{U}$ such that $\mathbb{P}[\boldsymbol{w}_{t+1}=\cdot\mid\boldsymbol{w}_{\leqslant t}=\hat{\boldsymbol{w}}_{\leqslant t}]=\mathbb{Q}_{t+1}[\cdot]$. Since $\mathscr{U}$ is

exact, it is *inclusive* with respect to parent set $\mathcal{U}$. Hence the conditional distribution induced by this particular $\mathbb{P}$ on history $\hat{\boldsymbol{w}}_{\leqslant t}$ must belong to $\mathscr{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$. In other words, $\mathbb{Q}_{t+1} \in \mathscr{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$, which contradicts the supposition. Thus, $(\text{LAUS})_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}) \subseteq \mathscr{U}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$.

A.8. Theorem 3ii.

We aim to prove that for all $\hat{\boldsymbol{w}}_{t+1}$ such that $(\hat{\boldsymbol{w}}_{\leqslant t}, \hat{\boldsymbol{w}}_{t+1}) \in \hat{\mathcal{W}}_{\leqslant t+1}$, the distribution $\mathbb{Q}$ where $\mathbb{Q}[\boldsymbol{w}_{t+1} = \hat{\boldsymbol{w}}_{t+1}] = 1$ lies in (LAUS), $\mathbb{Q} \in \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})$. Moreover, if $\mathcal{U}$ has recoverable property, then for all $\hat{\boldsymbol{w}}_1 \in \hat{\mathcal{W}}_1$, there exists a probability distribution $\mathbb{P} \in \mathcal{U}$ such that $\mathbb{E}_{\mathbb{P}}[\boldsymbol{g}(\boldsymbol{w}_1, \boldsymbol{w}_{>1}) \,|\, \boldsymbol{w}_1 = \hat{\boldsymbol{w}}_1] = \boldsymbol{0}$, by which the ambiguity set contains all single-point marginal distributions of $\boldsymbol{w}_1$. Then, (AURA) = (RO).

Fix $\hat{\boldsymbol{w}}_{t+1}$. We aim to find some parent distribution $\mathbb{P} \in \mathcal{U}$ whose conditional distribution given $\hat{\boldsymbol{w}}_t$ gives $\hat{\boldsymbol{w}}_{t+1}$. Denote the filter operator that filters $\mathcal{W}$ by the given history $\hat{\boldsymbol{w}}_{\leqslant t}$ or $\hat{\boldsymbol{w}}_{\leqslant t+1}$ respectively

$$\mathcal{H}_t := \{\boldsymbol{w} \in \mathcal{W} : \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\},$$
$$\mathcal{H}_{t+1} := \{\boldsymbol{w} \in \mathcal{W} : \boldsymbol{w}_{\leqslant t+1} = \hat{\boldsymbol{w}}_{\leqslant t+1}\}.$$

Evidently, $\mathcal{H}_{t+1} \subseteq \mathcal{H}_t$. Here, we eliminate the special case where $\mathcal{W} \backslash \mathcal{H}_t = \varnothing$. In this case, for all $\boldsymbol{w} \in \mathcal{W}$, $\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}$. In other words, there is no uncertainty for all times up to time t. Thus, (AURA) = (RO) if and only if their problems-to-go after $\hat{\boldsymbol{w}}_{\leqslant t}$ is revealed coincide. We can easily deal with this by induction on the time horizon T. Thus, hereon, assume that $\mathcal{W} \backslash \mathcal{H}_t \neq \varnothing$.

Let $\mathbb{P}'$ be any distribution that is supported on $\mathcal{H}_{t+1}$ and therefore, note that $\mathbb{P}'[\boldsymbol{w}_{t+1} = \hat{\boldsymbol{w}}_{t+1} \,|\, \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}] = 1$. It is not guaranteed that $\mathbb{P}' \in \mathcal{U}$ and in particular, let $\boldsymbol{m}' := \mathbb{E}_{\mathbb{P}'}[\boldsymbol{g}(\boldsymbol{w})]$ where in general $\boldsymbol{m}' \neq \boldsymbol{0}$. Our strategy is to find some distribution $\mathbb{P}''$ that is supported on $\mathcal{W} \backslash \mathcal{H}_t$ where $\mathbb{E}_{\mathbb{P}''}[\boldsymbol{g}(\boldsymbol{w})] = \frac{-\alpha \boldsymbol{m}'}{1-\alpha}$ for some small $\alpha \in (0,1)$. Then the distribution $\mathbb{P} := \alpha \mathbb{P}' + (1-\alpha)\mathbb{P}''$ is the parent distribution we are looking for, namely that $\mathbb{E}_{\mathbb{P}}[\boldsymbol{g}(\boldsymbol{w})] = \boldsymbol{0}$ and $\mathbb{P}[\boldsymbol{w}_{t+1} = \hat{\boldsymbol{w}}_{t+1} \mid \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}] = \mathbb{P}'[\boldsymbol{w}_{t+1} = \hat{\boldsymbol{w}}_{t+1} \mid \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}] = 1$, since $\mathbb{P}''$ is not supported on $\hat{\boldsymbol{w}}_{\leqslant t}$.

To construct $\mathbb{P}''$, consider the sequence of lemmas:

LEMMA 2. *Let $\mathcal{S}$ be a convex set and let $\{\mathcal{G}_m\}_{m=1}^M$ be a finite collection of hyperplanes where $\mathcal{S} \cap \bigcap_{m=1}^M \mathcal{G}_m \neq \varnothing$ or $\mathcal{S}$. Then*

$$\mathcal{S} \cap \bigcap_{m=1}^M \mathcal{G}_m \subseteq \mathrm{cl}\left(\mathcal{S} \Big\backslash \bigcap_{m=1}^M \mathcal{G}_m\right).$$

Proof of Lemma 2. Let $\boldsymbol{a} \in \mathcal{S} \cap \bigcap_{m=1}^M \mathcal{G}_m$ be any point in the set. Take any $\boldsymbol{b} \in \mathcal{S} \backslash \bigcap_{m=1}^M \mathcal{G}_m$. Consider the line segment $\mathcal{L}$ adjoining $\boldsymbol{a}$ and $\boldsymbol{b}$. Now, since each $\mathcal{G}_m$ is a hyperplane, either $\mathcal{L} \cap \mathcal{G}_m = \{\boldsymbol{a}\}$ or $\mathcal{L} \subseteq \mathcal{G}_m$. Hence, either $\mathcal{L} \cap \bigcap_{m=1}^M \mathcal{G}_m = \{\boldsymbol{a}\}$ or $\mathcal{L} \subseteq \bigcap_{m=1}^M \mathcal{G}_m$. However, the latter contradicts the fact that $\boldsymbol{b} \notin \bigcap_{m=1}^M \mathcal{G}_m$. Moreover, since $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{S}$ is convex, $\mathcal{L} \subseteq \mathcal{S}$. Hence $\mathcal{L} \backslash \{\boldsymbol{a}\} \subseteq \mathcal{S} \backslash \bigcap_{m=1}^M \mathcal{G}_m$.

By choosing points on $\mathcal{L}$, there is a sequence $\boldsymbol{b}_n \in \mathcal{L}\backslash\{\boldsymbol{a}\}$, therefore $\boldsymbol{b}_n \in \mathcal{S}\backslash\bigcap_{m=1}^{M}\mathcal{G}_m$, such that $\boldsymbol{b}_n \to \boldsymbol{a}$. Thus, $\boldsymbol{a} \in \mathrm{cl}\left(\mathcal{S}\backslash\bigcap_{m=1}^{M}\mathcal{G}_m\right)$.

□

LEMMA 3. *For a set $\mathcal{S}$, we have* $\mathrm{cl}(\mathrm{conv}\,\{\mathcal{S}\}) = \mathrm{cl}(\mathrm{conv}\,\{\mathrm{cl}(\mathcal{S})\})$.

Proof of Lemma 3. On the one hand, since every set is contained in its closure, and by taking convex hull, $\mathrm{conv}\,\{\mathcal{S}\} \subseteq \mathrm{conv}\,\{\mathrm{cl}(\mathcal{S})\}$. Taking closure again,

$$\mathrm{cl}(\mathrm{conv}\,\{\mathcal{S}\}) \subseteq \mathrm{cl}(\mathrm{conv}\,\{\mathrm{cl}(\mathcal{S})\}). \tag{43}$$

On the other hand, $\mathcal{S} \subseteq \mathrm{conv}\,\{\mathcal{S}\}$, thus $\mathrm{cl}(\mathcal{S}) \subseteq \mathrm{cl}(\mathrm{conv}\,\{\mathcal{S}\})$. Moreover, we claim that if set $\mathcal{R}$ is a convex set and $\mathcal{A} \subseteq \mathrm{cl}(\mathcal{R})$, then $\mathrm{conv}\,\{\mathcal{A}\} \subseteq \mathrm{cl}(\mathcal{R})$. Indeed, let $\boldsymbol{a} \in \mathrm{conv}\,\{\mathcal{A}\}$, then there exists $\boldsymbol{a}^i \in \mathcal{A}$ and $\lambda_i \in [0,1]$, $\sum_i \lambda_i = 1$, such that $\boldsymbol{a} = \sum_i \lambda_i \boldsymbol{a}^i$. Since $\boldsymbol{a}^i \in \mathcal{A} \subseteq \mathrm{cl}(\mathcal{R})$, for each i, there exists sequences $\boldsymbol{r}_n^{(i)} \in \mathcal{R}$ such that $\boldsymbol{r}_n^{(i)} \to \boldsymbol{a}^i$ as $n \to \infty$. Let $\boldsymbol{r}_n = \sum_i \lambda_i \boldsymbol{r}_n^{(i)}$. Then by the convexity of $\mathcal{R}$, $\boldsymbol{r}_n \in \mathcal{R}$, for all n. Hence, $\boldsymbol{r}_n \to \boldsymbol{a}$ as $n \to \infty$ and so $\boldsymbol{a} \in \mathrm{cl}(\mathcal{R})$.

Letting $\mathcal{A} = \mathrm{cl}(\mathcal{S})$ and $\mathcal{R} = \mathrm{conv}\,\{\mathcal{S}\}$, we get that $\mathrm{conv}\,\{\mathrm{cl}(\mathcal{S})\} \subseteq \mathrm{cl}(\mathrm{conv}\,\{\mathcal{S}\})$, again taking closure

$$\mathrm{cl}(\mathrm{conv}\,\{\mathrm{cl}(\mathcal{S})\}) \subseteq \mathrm{cl}(\mathrm{conv}\,\{\mathcal{S}\}) \tag{44}$$

Putting (43) and (44) together yields the result. □

LEMMA 4. *There exists some small radius $\epsilon > 0$ such that $\mathcal{B}_\epsilon(\mathbf{0}) \subseteq \mathrm{conv}\,\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_{t+1}\}$.*

Proof of Lemma 4. By containment property of $\mathcal{U}$, $\mathbf{0}$ lies in the strict interior of $\mathrm{conv}\,\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\}$. In other words, there exists some ball $\mathcal{B}_\epsilon(\mathbf{0}) \subseteq \mathrm{conv}\,\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\}$. We argue that the same ball must lie in $\mathrm{conv}\,\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_{t+1}\}$.

By Lemma 2, using $\mathcal{S} = \mathcal{W}$ and $\mathcal{G}_m = \{\boldsymbol{w} : \boldsymbol{w}_m = \hat{\boldsymbol{w}}_m\}$ for $\forall m \leqslant t$, we get that $\mathcal{H}_t \subseteq \mathrm{cl}(\mathcal{W}\backslash\mathcal{H}_t)$. Continuity of $\boldsymbol{g}$ implies that $\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{H}_t\} \subseteq \mathrm{cl}(\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\})$. Trivially, $\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\} \subseteq \mathrm{cl}(\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\})$, thus $\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\} = \{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{H}_t\} \cup \{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\} \subseteq \mathrm{cl}(\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\})$.

We can also show that $\mathrm{cl}(\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\}) \subseteq \{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\}$. In fact, we first claim that $\mathrm{cl}(\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\}) \subseteq \{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathrm{cl}(\mathcal{W}\backslash\mathcal{H}_t)\}$. Indeed, suppose there is a sequence of $\boldsymbol{w}_n \in \mathcal{W}\backslash\mathcal{H}_t$ such that $\boldsymbol{g}(\boldsymbol{w}_n) \to \boldsymbol{g}^*$ for some limit $\boldsymbol{g}^*$. Since $\mathcal{W}$ is compact, there is a convergent subsequence $\boldsymbol{w}_{n_k} \to \boldsymbol{w}^* \in \mathrm{cl}(\mathcal{W}\backslash\mathcal{H}_t)$. By continuity, $\boldsymbol{g}(\boldsymbol{w}_{n_k}) \to \boldsymbol{g}(\boldsymbol{w}^*)$. Hence, $\boldsymbol{g}(\boldsymbol{w}^*) = \boldsymbol{g}^*$ for some $\boldsymbol{w}^* \in \mathrm{cl}(\mathcal{W}\backslash\mathcal{H}_t)$. This proves the claim. Moreover, since $\mathrm{cl}(\mathcal{W}\backslash\mathcal{H}_t) \subseteq \mathcal{W}$, hence $\mathrm{cl}(\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\}) \subseteq \{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\}$.

Putting both together, we find that $\mathrm{cl}(\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\}) = \{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\}$. Using Lemma 3, we have $\mathrm{cl}(\mathrm{conv}\,\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\}) = \mathrm{conv}\,\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\}$. Then, $\mathcal{B}_\epsilon(\mathbf{0}) \subseteq \mathrm{cl}(\mathrm{conv}\,\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\})$. As the ball is open, $\mathcal{B}_\epsilon(\mathbf{0}) \subseteq \mathrm{conv}\,\{\boldsymbol{g}(\boldsymbol{w}) : \boldsymbol{w} \in \mathcal{W}\backslash\mathcal{H}_t\}$. □

By Lemma 4, for α small enough, $\frac{-\alpha \boldsymbol{m}'}{1-\alpha} \in \mathcal{B}_\epsilon(\mathbf{0})$. Thus, there is a collection of $\boldsymbol{w}_i \in \mathcal{W} \backslash \mathcal{H}_t$ and $\mu_i \in [0,1]$ where $\sum_i \mu_i = 1$ such that $\frac{-\alpha \boldsymbol{m}'}{1-\alpha} = \sum_i \mu_i \boldsymbol{g}(\boldsymbol{w}_i)$. Let $\mathbb{P}''$ be the distribution which takes value $\boldsymbol{w}_i$ with probability μ_i. Then $\mathbb{E}_{\mathbb{P}''}[\boldsymbol{g}(\boldsymbol{w})] = \frac{-\alpha \boldsymbol{m}'}{1-\alpha}$ by construction. This is the $\mathbb{P}''$ we seek.

A.9. Theorem 4

Given any policy $\boldsymbol{x}$, under time t and history $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_t$, the problems-to-go of (AURA) and (RO) are

$$F_t^{\mathrm{AURA}}(\hat{\boldsymbol{w}}_{\leqslant t}) := c_t(\boldsymbol{x}_t) + \sup_{\mathbb{Q}_{t+1} \in \mathscr{J}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t})} \mathbb{E}_{\mathbb{Q}_{t+1}} \left[F_{t+1}^{\mathrm{AURA}}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) \right];$$
$$F_t^{\mathrm{RO}}(\hat{\boldsymbol{w}}_{\leqslant t}) := c_t(\boldsymbol{x}_t) + \sup_{\boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1}} F_{t+1}^{\mathrm{RO}}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}),$$

where $\sup_{\boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1}} F_{t+1}^{\mathrm{RO}}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) = \sup_{\mathbb{Q}_{t+1} \in \mathcal{P}(\hat{\mathcal{W}}_{t+1})} \mathbb{E}_{\mathbb{Q}_{t+1}} \left[F_{t+1}^{\mathrm{RO}}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) \right]$.

Since $\tau \in \mathscr{C}(\mathcal{U})$, for all history $\hat{\boldsymbol{w}}_{\leqslant \tau} \in \hat{\mathcal{W}}_\tau$, the ambiguity sets $\mathscr{J}_{\tau+1}(\hat{\boldsymbol{w}}_{\leqslant \tau})$ contains all distributions, *i.e.*, $\mathscr{J}_{\tau+1}(\hat{\boldsymbol{w}}_{\leqslant \tau}) = \mathcal{P}(\hat{\mathcal{W}}_{\tau+1})$. This holds for all times $\tau \geqslant t$. Thus, by backward induction, $F_\tau^{\mathrm{AURA}}(\hat{\boldsymbol{w}}_{\leqslant \tau}) = F_\tau^{\mathrm{RO}}(\hat{\boldsymbol{w}}_{\leqslant \tau})$ for all times $\tau \geqslant t$ and histories $\hat{\boldsymbol{w}}_{\leqslant \tau} \in \hat{\mathcal{W}}_{\leqslant \tau}$.

Additionally, when $\mathcal{U}$ is perfectly collapsible, $F_1^{\mathrm{AURA}}(\hat{\boldsymbol{w}}_1) = F_1^{\mathrm{RO}}(\hat{\boldsymbol{w}}_1)$ for histories $\hat{\boldsymbol{w}}_1 \in \hat{\mathcal{W}}_1 := \mathcal{W}_1$. Thus, for the initial problems of (AURA) and (RO),

$$F_0^{\mathrm{AURA}} := c_0(\boldsymbol{x}_0) + \sup_{\mathbb{Q}_1 \in \mathscr{J}_1} \mathbb{E}_{\mathbb{Q}_1} \left[F_1^{\mathrm{AURA}}(\hat{\boldsymbol{w}}_1) \right]$$
$$F_0^{\mathrm{RO}} := c_0(\boldsymbol{x}_0) + \sup_{\boldsymbol{w}_1 \in \mathcal{W}_1} F_1^{\mathrm{RO}}(\hat{\boldsymbol{w}}_1)$$

where $\mathscr{J}_1 = \{\mathbb{P}_{\boldsymbol{w}_1} : \mathbb{P} \in \mathcal{U}\}$ is the set of all marginals of parent ambiguity set $\mathcal{U}$. Moreover, since $\mathcal{U}$ has recoverable property, because of (9), for every $\hat{\boldsymbol{w}}_1 \in \hat{\mathcal{W}}_1$, there exists some distribution $\mathbb{Q}_{>1}^{\hat{\boldsymbol{w}}_1}$ such that $\mathbb{E}_{\mathbb{Q}_{>1}^{\hat{\boldsymbol{w}}_1}}[\boldsymbol{g}(\hat{\boldsymbol{w}}_1, \boldsymbol{w}_{>1})] = \mathbf{0}$. Then we can construct a joint distribution by

$$\tilde{\mathbb{P}} := \delta_{\hat{\boldsymbol{w}}_1} \otimes \mathbb{Q}_{>1}^{\hat{\boldsymbol{w}}_1}.$$

Then, $\mathbb{E}_{\tilde{\mathbb{P}}}[\boldsymbol{g}(\boldsymbol{w})] = \mathbf{0}$, thus $\tilde{\mathbb{P}} \in \mathcal{U}$ and $\tilde{\mathbb{P}}_1 = \delta_{\hat{\boldsymbol{w}}_1} \in \mathscr{J}_1$. This holds for any $\hat{\boldsymbol{w}}_1$, thus $\mathscr{J}_1$ contains all single-point marginal distributions supported on $\mathcal{W}_1$. Thus, on the one hand, $F_0^{\mathrm{AURA}} \leqslant F_0^{\mathrm{RO}}$; on the other hand, since F_0^{RO} attains its maximum at some point, we have

$$\sup_{\mathbb{Q}_1 \in \mathscr{J}_1} \mathbb{E}_{\mathbb{Q}_1} \left[F_1^{\mathrm{AURA}}(\boldsymbol{w}_1) \right] = \sup_{\boldsymbol{w}_1 \in \mathcal{W}_1} F_1^{\mathrm{RO}}(\boldsymbol{w}_1).$$

Consequently, $F_0^{\mathrm{AURA}} = F_0^{\mathrm{RO}}$. Finally, since this holds for any policy $\boldsymbol{x}$, we have the conclusion, (AURA) = (RO).

A.10. Theorem 5i.

Clearly, $L_t(\varepsilon_t,\eta_t)\geqslant 0$. We argue by contradiction. Suppose $L_t(\varepsilon_t,\eta_t)=0$. Let $\ell(\mathbb{P}):=\sup_{\boldsymbol{w}\neq\boldsymbol{w}'\in\mathcal{W}}\frac{\left|p_{\mathbb{P}}(\boldsymbol{w})-p_{\mathbb{P}}(\boldsymbol{w}')\right|}{\|\boldsymbol{w}-\boldsymbol{w}'\|}$. Then there exists a sequence $\{\mathbb{P}_n\in\mathcal{D}_t\}_{n=1}^{\infty}$ such that $\ell(\mathbb{P}_n)\to 0$. Let p_n denote the density of $\mathbb{P}_n$.

Since $\mathcal{W}$ is compact, it is bounded and $\sup_{\boldsymbol{w}\neq\boldsymbol{w}'}\|\boldsymbol{w}-\boldsymbol{w}'\|\leqslant M<\infty$ for some M. Hence, $\ell(\mathbb{P}_n)\to 0$ implies that $\sup_{\boldsymbol{w}\neq\boldsymbol{w}'}|p_n(\boldsymbol{w})-p_n(\boldsymbol{w}')|\to 0$. Thus,

$$0\leqslant\max_{\boldsymbol{w}}p_n(\boldsymbol{w})-\min_{\boldsymbol{w}}p_n(\boldsymbol{w})\leqslant\sup_{\boldsymbol{w},\boldsymbol{w}'}|p_n(\boldsymbol{w})-p_n(\boldsymbol{w}')|\to 0.$$

Thus, $\max_{\boldsymbol{w}}p_n(\boldsymbol{w})-\min_{\boldsymbol{w}}p_n(\boldsymbol{w})\to 0$. Now, $\min_{\boldsymbol{w}}p_n(\boldsymbol{w})\leqslant p_n(\boldsymbol{w})\leqslant\max_{\boldsymbol{w}}p_n(\boldsymbol{w})$ for all $\boldsymbol{w}\in\mathcal{W}$. Hence, integrating over $\boldsymbol{w}\in\mathcal{W}$, we get for all n, $\min_{\boldsymbol{w}}p_n(\boldsymbol{w})\leqslant\frac{1}{\mathrm{vol}(\mathcal{W})}\leqslant\max_{\boldsymbol{w}}p_n(\boldsymbol{w})$. (Note: $0<\mathrm{vol}\,(\mathcal{W})<\infty$.) This implies that

$$\sup_{\boldsymbol{w}\in\mathcal{W}}\left|p_n^t(\boldsymbol{w})-\frac{1}{\mathrm{vol}\,(\mathcal{W})}\right|\leqslant\max_{\boldsymbol{w}\in\mathcal{W}}p_n^t(\boldsymbol{w})-\min_{\boldsymbol{w}\in\mathcal{W}}p_n^t(\boldsymbol{w})\to 0.$$

This is to say that for all $\xi>0$, there exists some $N:=N(\xi)$, such that for all $n\geqslant N$:

$$\frac{1}{\mathrm{vol}\,(\mathcal{W})}-\xi\leqslant p_n^t(\boldsymbol{w})\leqslant\frac{1}{\mathrm{vol}\,(\mathcal{W})}+\xi,\qquad\forall\boldsymbol{w}\in\mathcal{W}.\tag{45}$$

As $\mathbb{P}_n\in\mathcal{D}_t$, there exists some $\hat{\boldsymbol{w}}_{\leqslant t}^n$ and some $\boldsymbol{w}_{>t+1}^n$, corresponding to each $\mathbb{P}_n$, such that

$$\mathbb{P}_n[\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1}^n)|\boldsymbol{w}_{\leqslant t}=\hat{\boldsymbol{w}}_{\leqslant t}^n]=\frac{\int_{\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1}^n)\cap\hat{\mathcal{W}}_{t+1}}p_n(\hat{\boldsymbol{w}}_{\leqslant t}^n,\boldsymbol{w}_{>t})d\boldsymbol{w}_{>t}}{\int_{\hat{\mathcal{W}}_{t+1}}p_n(\hat{\boldsymbol{w}}_{\leqslant t}^n,\boldsymbol{w}_{>t})d\boldsymbol{w}_{>t}}\geqslant 1-\eta_t.$$

By (45), we have for all $n\geqslant N(\xi)$,

$$1-\eta_t\leqslant\mathbb{P}_n[\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1}^n)|\boldsymbol{w}_{\leqslant t}=\hat{\boldsymbol{w}}_{\leqslant t}^n]\leqslant\frac{\frac{1}{\mathrm{vol}(\mathcal{W})}+\xi}{\frac{1}{\mathrm{vol}(\mathcal{W})}-\xi}\cdot\frac{\mathrm{vol}\left(\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1}^n)\cap\hat{\mathcal{W}}_{t+1}\right)}{\mathrm{vol}\left(\hat{\mathcal{W}}_{t+1}\right)}.$$

Hence,

$$(1-\eta_t)\frac{\frac{1}{\mathrm{vol}(\mathcal{W})}-\xi}{\frac{1}{\mathrm{vol}(\mathcal{W})}+\xi}\leqslant\frac{\mathrm{vol}\left(\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1}^n)\cap\hat{\mathcal{W}}_{t+1}\right)}{\mathrm{vol}\left(\hat{\mathcal{W}}_{t+1}\right)}\leqslant\sup_{\substack{\hat{\boldsymbol{w}}_{\leqslant t}\in\hat{\mathcal{W}}_{\leqslant t},\\ \boldsymbol{w}_{t+1}\in\hat{\mathcal{W}}_{t+1}}}\frac{\mathrm{vol}\left(\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1})\cap\hat{\mathcal{W}}_{t+1}\right)}{\mathrm{vol}\left(\hat{\mathcal{W}}_{t+1}\right)}<1-\eta_t,$$

where the last inequality is by assumption of (12). However, $\frac{\frac{1}{\mathrm{vol}(\mathcal{W})}-\xi}{\frac{1}{\mathrm{vol}(\mathcal{W})}+\xi}\to 1$, as $\xi\downarrow 0$, thus

$$(1-\eta_t)\leqslant\sup_{\substack{\hat{\boldsymbol{w}}_{\leqslant t}\in\hat{\mathcal{W}}_{\leqslant t},\\ \boldsymbol{w}_{t+1}\in\hat{\mathcal{W}}_{t+1}}}\frac{\mathrm{vol}\left(\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1})\cap\hat{\mathcal{W}}_{t+1}\right)}{\mathrm{vol}\left(\hat{\mathcal{W}}_{t+1}\right)}<1-\eta_t,$$

which is a contradiction.

A.11. Theorem 5ii.

Fix history $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$. Hereon, we shorthand $J_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}; \mathcal{U}^*(L))$ as $\mathscr{J}_{t+1}$. Since

$$\inf_{\mathbb{P}\in\mathcal{D}_t(\varepsilon_t,\eta_t)} \sup_{\substack{\boldsymbol{w},\boldsymbol{w}'\in\mathcal{W},\\ \boldsymbol{w}\neq\boldsymbol{w}'}} \frac{|p_{\mathbb{P}}(\boldsymbol{w}) - p_{\mathbb{P}}(\boldsymbol{w}')|}{\|\boldsymbol{w}-\boldsymbol{w}'\|} := L_t > L,$$

it follows that for all $\mathbb{P} \in \mathcal{D}_t$, there are some $\boldsymbol{w}, \boldsymbol{w}' \in \mathcal{W}$, $\boldsymbol{w} \neq \boldsymbol{w}'$, where for the corresponding density p, $|p(\boldsymbol{w}) - p(\boldsymbol{w}')| > L\|\boldsymbol{w}-\boldsymbol{w}'\|$. Thus, $\mathbb{P} \notin \mathcal{U}^*(L)$, *i.e.* $\mathcal{D}_t \cap \mathcal{U}^*(L) = \varnothing$.

Claim: For all $\mathbb{Q} \in \mathscr{J}_{t+1}$ and for all $\boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1}$, $\mathbb{Q}[\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1})] < 1-\eta_t$.
Proof: Suppose there is some $\mathbb{Q} \in \mathscr{J}_{t+1}$ and some $\boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1}$, such that $\mathbb{Q}[\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1})] \geqslant 1-\eta_t$. Then there exists some $\tilde{\mathbb{P}} \in \mathcal{U}^*(L)$ such that $\mathbb{Q}[\cdot] = \tilde{\mathbb{P}}[\cdot|\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}]$. Hence, $\tilde{\mathbb{P}}[\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1}) \mid \boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}] \geqslant 1-\eta_t$, and thus $\tilde{\mathbb{P}} \in \mathcal{D}_t$, contradicting $\mathcal{D}_t \cap \mathcal{U}^*(L) = \varnothing$. □

We now prove the respective statements.

a) Now, we seek $r := r(\hat{\boldsymbol{w}}_{\leqslant t}) > 0$ such that $B_r(\delta_{\boldsymbol{w}_{t+1}}) \cap \mathscr{J}_{t+1} = \varnothing, \forall \boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1}$. Define

$$r := \inf\{\xi : \exists \boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1} \text{ s.t. } B_\xi(\delta_{\boldsymbol{w}_{t+1}}) \cap \mathscr{J}_{t+1} \neq \varnothing\}.$$

If $r > 0$, then this is the r we seek and we are done.

Instead, assume that $r = 0$. Hence, for $\xi_n = \frac{1}{n}$, $n = 1, \ldots, \infty$, there exists some $\boldsymbol{w}^n_{t+1} \in \hat{\mathcal{W}}_{t+1}$ such that $B_{1/n}(\delta_{\boldsymbol{w}^n_{t+1}}) \cap \mathscr{J}_{t+1} \neq \varnothing$. Note that if there exists some n where for all $\boldsymbol{w}^n_{t+1} \in \hat{\mathcal{W}}_{t+1}$ such that $B_{1/n}(\delta_{\boldsymbol{w}^n_{t+1}}) \cap \mathscr{J}_{t+1} = \varnothing$, then because $B_\xi(\delta_{\boldsymbol{w}^n_{t+1}}) \subset B_{1/n}(\delta_{\boldsymbol{w}^n_{t+1}})$ if $\xi \leqslant \frac{1}{n}$, therefore $r \geqslant \frac{1}{n}$, contradicting the supposition that $r = 0$. Since $\hat{\mathcal{W}}_{t+1}$ is compact, there is a subsequence of $(\boldsymbol{w}^n_{t+1})_n$ that is convergent. Hence, w.l.o.g., for notational clarity, we assume that $\boldsymbol{w}^n_{t+1} \to \tilde{\boldsymbol{w}}_{t+1} \in \hat{\mathcal{W}}_{t+1}$. Moreover, since $B_{1/n}(\delta_{\boldsymbol{w}^n_{t+1}}) \cap \mathscr{J}_{t+1} \neq \varnothing$, hence there exists a sequence of distributions $(\mathbb{P}^n)_n$ such that $\mathbb{P}^n \in B_{1/n}(\delta_{\boldsymbol{w}^n_{t+1}}) \cap \mathscr{J}_{t+1}$.

We argue that for all $\theta > 0$, $B_\theta(\delta_{\tilde{\boldsymbol{w}}_{t+1}}) \cap \mathscr{J}_{t+1} \neq \varnothing$. Note that

$$d(\delta_{\tilde{\boldsymbol{w}}_{t+1}}, \mathbb{P}^n) \leqslant d(\delta_{\tilde{\boldsymbol{w}}_{t+1}}, \delta_{\boldsymbol{w}^n_{t+1}}) + d(\delta_{\boldsymbol{w}^n_{t+1}}, \mathbb{P}^n) \leqslant D\|\boldsymbol{w}^n_{t+1} - \tilde{\boldsymbol{w}}_{t+1}\| + \frac{1}{n}, \tag{46}$$

where the first inequality holds due to the triangle inequality; and in the second inequality, the first term holds as the metric $d(\cdot,\cdot)$ is PDL and the second term holds by definition of $\mathbb{P}^n \in B_{1/n}(\delta_{\boldsymbol{w}^n_{t+1}})$. Consequently, for any $\theta > 0$, since $\boldsymbol{w}^n_{t+1} \to \tilde{\boldsymbol{w}}_{t+1}$, we can pick some m_1 such that $\|\boldsymbol{w}^n_{t+1} - \tilde{\boldsymbol{w}}_{t+1}\| < \frac{\theta}{2D}$ for all $n \geqslant m_1$, and pick m_2 such that $\frac{1}{n} < \frac{\theta}{2}$ for all $n \geqslant m_2$. Take $m = \max\{m_1, m_2\}$. Thus, for all $n \geqslant m$, $d(\delta_{\tilde{\boldsymbol{w}}_{t+1}}, \mathbb{P}^n) < \theta$, *i.e.* $\mathbb{P}^n \in B_\theta(\delta_{\tilde{\boldsymbol{w}}_{t+1}}) \cap \mathscr{J}_{t+1}$ for all $n \geqslant m$, in particular, $B_\theta(\delta_{\tilde{\boldsymbol{w}}_{t+1}}) \cap \mathscr{J}_{t+1} \neq \varnothing$. This proves our claim.

Thus, there exists a sequence of distributions $\mathbb{Q}^m \in \mathscr{J}_{t+1}$, such that $\mathbb{Q}^m \to \delta_{\tilde{\boldsymbol{w}}_{t+1}}$ weakly under $d(\cdot,\cdot)$. Consider the function $f(\boldsymbol{w}) = \max\left\{1 - \frac{\|\boldsymbol{w}-\tilde{\boldsymbol{w}}_{t+1}\|}{\varepsilon_t}, 0\right\}$, which is bounded and continuous. Thus,

$$\int f(\boldsymbol{w}) d\mathbb{Q}_n(\boldsymbol{w}) \to \int f(\boldsymbol{w}) d\delta_{\tilde{\boldsymbol{w}}_{t+1}}(\boldsymbol{w}) = 1. \tag{47}$$

However, since $\frac{\|\boldsymbol{w}-\tilde{\boldsymbol{w}}_{t+1}\|}{\varepsilon_t} \leqslant 1$ is equivalent to $\boldsymbol{w} \in \mathcal{B}_{\varepsilon_t}(\tilde{\boldsymbol{w}}_{t+1})$, according to the Claim,

$$\int f(\boldsymbol{w}) d\mathbb{Q}_n(\boldsymbol{w}) \leqslant \tilde{\mathbb{Q}}_n[\mathcal{B}_{\varepsilon_t}(\tilde{\boldsymbol{w}}_{t+1})] < 1-\eta_t,$$

which contradicts (47). Thus, $r \neq 0$.

b) For any $\mathbb{Q} \in \mathscr{J}_{t+1}$, the Claim gives that $\mathbb{Q}[\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}_{t+1})] < 1-\eta_t, \forall \boldsymbol{w}_{t+1} \in \hat{\mathcal{W}}_{t+1}$, in particular, for worst-case $\boldsymbol{w}^*_{t+1}$, we have $\mathbb{Q}[\hat{\mathcal{W}}_{t+1} \backslash \mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})] \geqslant \eta_t$. Moreover,

$$\begin{aligned}
&\mathbb{E}_{\mathbb{Q}}[F^{\mathrm{AURA}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})] \leqslant \mathbb{E}_{\mathbb{Q}}[F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})] && (48)\\
\leqslant &\mathbb{E}_{\mathbb{Q}}[F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) | \mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})]\, \mathbb{Q}[\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})] \\
&\quad + \mathbb{E}_{\mathbb{Q}}[F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) | \hat{\mathcal{W}}_{t+1} \backslash \mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})]\, \mathbb{Q}[\hat{\mathcal{W}}_{t+1} \backslash \mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})] \\
\leqslant &F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}^*_{t+1}) \mathbb{Q}[\mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})] + (F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}^*_{t+1}) - \Delta_t(\hat{\boldsymbol{w}}_{\leqslant t}; \varepsilon_t)) \mathbb{Q}[\hat{\mathcal{W}}_{t+1} \backslash \mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})] && (49)\\
= &F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}^*_{t+1}) - \Delta(\hat{\boldsymbol{w}}_{\leqslant t}; \varepsilon_t) \tilde{\mathbb{Q}}[\hat{\mathcal{W}}_{t+1} \backslash \mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})] \\
\leqslant &F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}^*_{t+1}) - \Delta(\hat{\boldsymbol{w}}_{\leqslant t}; \varepsilon_t)\, \eta_t, && (50)
\end{aligned}$$

where (48) holds due to Theorem 1i., (49) holds by definition of $\Delta_t(\hat{\boldsymbol{w}}_{\leqslant t}; \varepsilon_t)$ and (50) holds as $\Delta_t(\hat{\boldsymbol{w}}_{\leqslant t}; \varepsilon_t) \geqslant 0$ by definition. Taking supremum yields the desired bound:

$$F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}^*_{t+1}) - \sup_{\mathbb{Q} \in \mathscr{J}_{t+1}} \mathbb{E}_{\mathbb{Q}}[F^{\mathrm{AURA}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})] \geqslant \eta_t \Delta(\hat{\boldsymbol{w}}_{\leqslant t}; \varepsilon_t).$$

Moreover, since $\hat{\mathcal{W}}_{t+1} \backslash \mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})$ is compact and the RO problem-to-go $F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \cdot)$ is continuous, the latter attains its maximum on $\hat{\mathcal{W}}_{t+1} \backslash \mathcal{B}_{\varepsilon_t}(\boldsymbol{w}^*_{t+1})$. Under Assumption 6, this objective value is strictly smaller than $F^{\mathrm{RO}}_{t+1}(\hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}^*_{t+1})$, hence $\Delta(\hat{\boldsymbol{w}}_{\leqslant t}; \varepsilon_t) > 0$.

A.12. Proof that $D(\cdot,\cdot)$ defines a pseudo-metric

PROPOSITION 5. *The distance $D(\cdot,\cdot)$ is a pseudo-metric on bounded sets. Moreover, $D(B, B') = 0$ if and only if the sets B and B' differ on a set of measure 0.*

Proof of Propostion 5 We verify the defining properties of a pseudo-metric. First, for any bounded sets B_1 and B_2,

$$D(B_1, B_2) = W_1(U(B_1), U(B_2)) \geqslant 0,$$

because W_1 is nonnegative. Moreover,

$$D(B_1, B_1) = W_1(U(B_1), U(B_1)) = 0.$$

Second, for any bounded sets B_1 and B_2,

$$D(B_1, B_2) = W_1(U(B_1), U(B_2)) = W_1(U(B_2), U(B_1)) = D(B_2, B_1),$$

where the second equality follows from the symmetry of W_1.

Third, for any bounded sets B_1, B_2, B_3,

$$\begin{aligned} D(B_1, B_3) &= W_1(U(B_1), U(B_3)) \\ &\leqslant W_1(U(B_1), U(B_2)) + W_1(U(B_2), U(B_3)) \\ &= D(B_1, B_2) + D(B_2, B_3), \end{aligned}$$

where the inequality follows from the triangle inequality of W_1.

Therefore, $D(\cdot,\cdot)$ is a pseudo-metric on bounded sets. Moreover, B_1 and B_2 differ on a set of measure 0 if and only if $U(B_1)$ and $U(B_2)$ are the same distribution if and only if $W_1(U(B_1), U(B_2)) = 0$ since W_1 is a metric on the space of distributions. □

A.13. Theorem 6

ASSUMPTION 7 **(Slater's condition).** *i. there exists some* $\tilde{\mathbb{P}} \in \mathcal{P}(\mathcal{W})$ *satisfying* $\mathbb{E}_{\tilde{\mathbb{P}}}[\boldsymbol{g}(\boldsymbol{w})] = 0$ *and* $\left|\tilde{\mathbb{P}}(\bar{B}_k) - \tilde{\mathbb{P}}(\underline{B}_k)\right| < LD(\bar{B}_k, \underline{B}_k), k = 1, \ldots, K.$

ii. For every stage t*, history* $\hat{\boldsymbol{w}}_{\leqslant t-1}$*, and sufficiently small perturbation vector* α*, define* $\mathscr{J}_t^{\alpha}(\hat{\boldsymbol{w}}_{\leqslant t-1})$ *as the collection of conditional distributions satisfying the perturbed moment constraints as in Assumption 4, together with the* (CVD) *condition* $|p(\boldsymbol{w}) - p(\boldsymbol{w}')| \leqslant L\|\boldsymbol{w} - \boldsymbol{w}'\|, \forall \boldsymbol{w}, \boldsymbol{w}'$*, where* p *is the density of the corresponding parent distribution. Assume that there exists a tight family* $\hat{\mathscr{J}}_t(\hat{\boldsymbol{w}}_{\leqslant t-1})$ *such that* $\mathscr{J}_t^{\alpha}(\hat{\boldsymbol{w}}_{\leqslant t-1}) \subseteq \hat{\mathscr{J}}_t(\hat{\boldsymbol{w}}_{\leqslant t-1})$ *for all sufficiently small* α*. Moreover, assume that there exists at least one strictly feasible density* p *satisfying the perturbed moment constraints and the above density regularity conditions.*

The one-step reformulation for the problems-to-go is presented below. Both the proof of the Lemma and the Theorem are omitted as they are identical to Lemma 1 and Theorem 2 respectively.

LEMMA 5. *For all times* $t = 0, ..., T-1$ *and histories* $\hat{\boldsymbol{w}}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$*, if* $\mathscr{J}_{t+1}^{CVD}(\hat{\boldsymbol{w}}_{\leqslant t}) \neq \varnothing$*, the worst-case distribution problem over the ambiguity set* (CVD) *has upper bound,*

$$\sup_{\mathbb{Q}_{t+1} \in \mathscr{J}_{t+1}^{CVD}(\hat{\boldsymbol{w}}_{\leqslant t})} \mathbb{E}_{\mathbb{Q}_{t+1}}\left[F_{t+1}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1})\right] \leqslant D_t^{CVD}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}),$$

where $D_t^{CVD}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t})$ *is defined as the problem:*

$$\begin{aligned} \min_{\substack{\boldsymbol{\lambda} \geqslant 0, \bar{\boldsymbol{\rho}}, \underline{\boldsymbol{\rho}} \geqslant 0 \\ \boldsymbol{\theta}, \eta \text{ free}}} \quad & \theta^{\hat{\boldsymbol{w}}_{\leqslant 0}} \qquad (51) \\ s.t. \quad & {\boldsymbol{\lambda}^{\hat{\boldsymbol{w}}_{\leqslant t}}}^{\top} \boldsymbol{g}(\boldsymbol{w}) + \left(\eta^{\hat{\boldsymbol{w}}_{\leqslant t}} + \sum_{k=1}^{K} \left(\mathbf{1}\left[\boldsymbol{w} \in \bar{B}_k\right] - \mathbf{1}\left[\boldsymbol{w} \in \underline{B}_k\right]\right)\left(\bar{\rho}_k^{\hat{\boldsymbol{w}}_{\leqslant t}} - \underline{\rho}_k^{\hat{\boldsymbol{w}}_{\leqslant t}}\right)\right) \mathbf{1}[t \geqslant 1] \\ & \quad + \theta^{\hat{\boldsymbol{w}}_{\leqslant t}} \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right] \geqslant F_{t+1}(\boldsymbol{x}_{\geqslant t+1}; \hat{\boldsymbol{w}}_{\leqslant t}, \boldsymbol{w}_{t+1}) \mathbf{1}\left[\boldsymbol{w}_{\leqslant t} = \hat{\boldsymbol{w}}_{\leqslant t}\right], \qquad \forall \boldsymbol{w} \in \mathcal{W}, \\ & \quad - \eta^{\hat{\boldsymbol{w}}_{\leqslant t}} - \sum_{k=1}^{K} (\bar{\rho}_k^{\hat{\boldsymbol{w}}_{\leqslant t}} + \underline{\rho}_k^{\hat{\boldsymbol{w}}_{\leqslant t}}) \cdot L \cdot D(\bar{B}_k, \underline{B}_k) \geqslant 0. \end{aligned}$$

B. Further details on the water management problem

In this Appendix, we present further details on the problem we considered for better context. For reference, our notation used in the main text is provided in the table below.

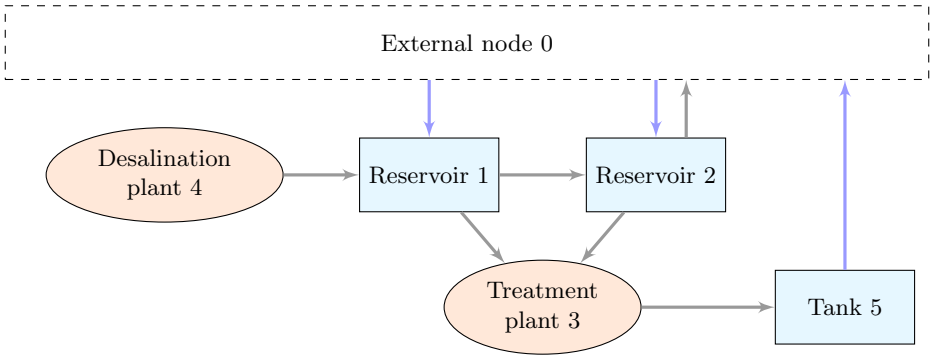
Table 2 List of Parameters and Variables

		Sets and indices
$t \in \mathcal{T}$	:	Time periods
$r \in \mathcal{R}$	:	Reservoirs holding rainwater
$h \in \mathcal{H}$	:	Tanks holding potable water
$p \in \mathcal{P}$	:	Treatment plants converting rainwater to potable water
$d \in \mathcal{D}$	:	Desalination plants converting sea water to rainwater
$\mathcal{E}^{\leftharpoonup}, \mathcal{E}^{\rightharpoonup}$	:	Edge sets for flows to be decided as here-and-now and recourse decisions respectively
		Parameters
κ_v	:	Upper capacity of water body $v \in \mathcal{R} \cup \mathcal{C}$
ι_v	:	Lower capacity of water body $v \in \mathcal{R} \cup \mathcal{C}$
τ_v	:	Maximum production capacity of production facility $v \in \mathcal{P} \cup \mathcal{D}$
$K^{\leftharpoonup}_{m,n}, K^{\rightharpoonup}_{m,n}$	:	Maximum flow capacity on edge (m,n) in $\mathcal{E}^{\leftharpoonup}$ and $\mathcal{E}^{\rightharpoonup}$ respectively
$c^{t,\leftharpoonup}_{m,n}, c^{t,\rightharpoonup}_{m,n}$	:	Unit cost of flow in period t on edge (m,n) in $\mathcal{E}^{\leftharpoonup}$ and $\mathcal{E}^{\rightharpoonup}$ respectively
		Uncertainties
w^t_r	:	Uncertain volume of rain that falls into reservoir $r \in \mathcal{R}$ at time t
w^t_h	:	Uncertain volume of water consumption demanded from tank $c \in \mathcal{H}$ at time t
		State variables
s^t_v	:	Initial volume of water at water body $v \in \mathcal{R} \cup \mathcal{C}$ at the end of time period t
		Decision variables
$x^{t,\leftharpoonup}_{m,n}(\boldsymbol{w}_{<t})$	:	Here-and-now decision given historical realization $\boldsymbol{w}_{<t} \in \hat{\mathcal{W}}_{<t}$ in period t on edge $(m,n) \in \mathcal{E}^{\leftharpoonup}$
$x^{t,\rightharpoonup}_{m,n}(\boldsymbol{w}_{\leqslant t})$	:	Recourse decision given historical realization $\boldsymbol{w}_{\leqslant t} \in \hat{\mathcal{W}}_{\leqslant t}$ in period t on edge $(m,n) \in \mathcal{E}^{\rightharpoonup}$

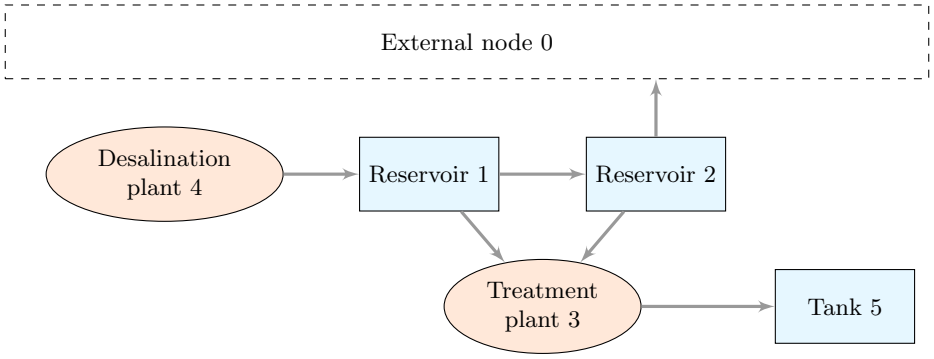
B.1. Details on the problem setting and experimental set-up

The specific network considered in the simulation is presented in Figure 9 below.

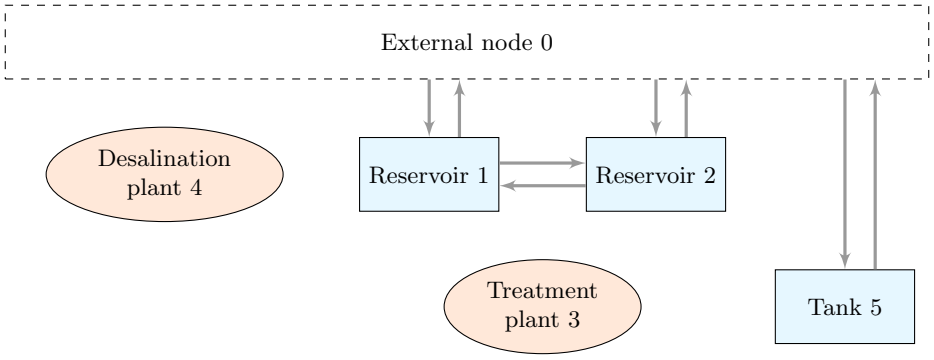
Figure 9 Water network considered in this numerical illustration, showing the full network (topmost), here-and-now decision structure (middle) and recourse decision structure (bottom-most).



(a) Full graph of the network



(b) Permitted here-and-now decisions before uncertainties materialize, $\mathcal{E}^{\curvearrowleft}$



(c) Permitted recourse decisions after uncertainties materialize, $\mathcal{E}^{\curvearrowright}$

B.2. Optimization formulations

Denote the feasible set of the water management problem under uncertainty as

$$\mathcal{X} := \left\{ \boldsymbol{x} \;\middle|\; \begin{array}{ll} \sum_{(v,n)\in\mathcal{E}^{\leftarrow}} x^{t,\leftarrow}_{v,n}(\boldsymbol{w}_{<t}) \leqslant \tau_v, & \forall v \in \mathcal{P}\cup\mathcal{D}, t\in\mathcal{T}, \forall \boldsymbol{w}_{<t}\in\mathcal{W}_{<t} \\ \sum_{(m,p)\in\mathcal{E}^{\leftarrow}} x^{t,\leftarrow}_{m,p}(\boldsymbol{w}_{<t}) = \sum_{(p,n)\in\mathcal{E}^{\leftarrow}} x^{t,\leftarrow}_{p,n}(\boldsymbol{w}_{<t}), & \forall p \in \mathcal{P}, t\in\mathcal{T}, \forall \boldsymbol{w}_{<t}\in\mathcal{W}_{<t} \\ \iota_v \leqslant s^t_v + \sum_{(m,v)\in\mathcal{E}^{\leftarrow}} x^{t,\leftarrow}_{m,v}(\boldsymbol{w}_{<t}) - \sum_{(v,m)\in\mathcal{E}^{\leftarrow}} x^{t,\leftarrow}_{v,m}(\boldsymbol{w}_{<t}) & \\ \quad + \sum_{(i,v)\in\mathcal{E}^{\rightarrow}} x^{t,\rightarrow}_{i,v}(\boldsymbol{w}_{\leqslant t}) - \sum_{(v,j)\in\mathcal{E}^{\rightarrow}} x^{t,\rightarrow}_{v,j}(\boldsymbol{w}_{\leqslant t}) & \\ \quad + w^t_v 1_{\{v\in\mathcal{R}\}} - w^t_v 1_{\{v\in\mathcal{H}\}} \leqslant \kappa_v, & \forall v \in \mathcal{R}\cup\mathcal{H}, t\in\mathcal{T}, \forall \boldsymbol{w}_{\leqslant t}\in\mathcal{W}_{\leqslant t} \\ 0 \leqslant x^{t,\leftarrow}_{m,n}(\boldsymbol{w}_{<t}) \leqslant K^{\leftarrow}_{m,n}, & \forall (m,n)\in\mathcal{E}^{\leftarrow}, t\in\mathcal{T}, \forall \boldsymbol{w}_{<t}\in\mathcal{W}_{<t} \\ 0 \leqslant x^{t,\rightarrow}_{m,n}(\boldsymbol{w}_{\leqslant t}) \leqslant K^{\rightarrow}_{m,n}, & \forall (m,n)\in\mathcal{E}^{\rightarrow}, t\in\mathcal{T}, \forall \boldsymbol{w}_{\leqslant t}\in\mathcal{W}_{\leqslant t} \end{array} \right\}$$

The parent DRO version of (WMP) is

$$\begin{aligned} \min_{\boldsymbol{x}^{\leftarrow}(\cdot),\boldsymbol{x}^{\rightarrow}(\cdot)} \sup_{\mathbb{Q}\in\mathcal{U}} \quad & \mathbb{E}_{\mathbb{Q}}\left[\sum_{t=1}^{T}\left(\boldsymbol{c}^{t,\leftarrow\top}\boldsymbol{x}^{t,\leftarrow}(\boldsymbol{w}_{<t}) + \boldsymbol{c}^{t,\rightarrow\top}\boldsymbol{x}^{t,\rightarrow}(\boldsymbol{w}_{\leqslant t})\right)\right] \\ \text{s.t.} \quad & (\boldsymbol{x}^{\leftarrow},\boldsymbol{x}^{\rightarrow})\in\mathcal{X}. \end{aligned} \tag{WMP-DRO}$$

where the parent ambiguity set $\mathcal{U}$ is defined in (20). As the adaptive policies $\boldsymbol{x}(\cdot)$ are constrained by the uncertainties $\boldsymbol{w}_{\leqslant t}$, each constraint in (WMP-DRO) is individually reformulated into its robust counterpart, ensuring feasibility for all admissible uncertainty realizations.

The corresponding RO version is

$$\begin{aligned} \min_{\boldsymbol{x}^{\leftarrow}(\cdot),\boldsymbol{x}^{\rightarrow}(\cdot)} \max_{\boldsymbol{w}\in\mathcal{W}} \quad & \sum_{t=1}^{T}\left(\boldsymbol{c}^{t,\leftarrow\top}\boldsymbol{x}^{t,\leftarrow}(\boldsymbol{w}_{<t}) + \boldsymbol{c}^{t,\rightarrow\top}\boldsymbol{x}^{t,\rightarrow}(\boldsymbol{w}_{\leqslant t})\right) \\ \text{s.t.} \quad & (\boldsymbol{x}^{\leftarrow},\boldsymbol{x}^{\rightarrow})\in\mathcal{X}. \end{aligned} \tag{WMP-RO}$$

where the support set $\mathcal{W}$ is given as discrete in 4.2 or continuous in 4.3.

B.3. Details of parameters and simulations setting

B.3.1. Out-of-sample generation. To test solutions obtained from above models, for each $\boldsymbol{\mu}_{\text{OOS}}$, we generate 20 test distributions over $\mathcal{W}$, with five at each density level: dense, moderately dense, moderately sparse, and sparse. The first three levels require positive probability on all, at least 50%, and at least 10% of the samples, respectively. The sparse level imposes no minimum coverage requirement. All distributions satisfy the prescribed moment constraints. Moreover, for the CVD experiments, we follow the same procedure with the CVD constraints additionally imposed. The pseudo code to generate test distributions is as follow:

Algorithm 1 Generation of OOS test distributions

Require: Support $\mathcal{W}$, test moment bounds $\boldsymbol{\mu}_{\text{oos}}$, threshold $\delta > 0$

Require: Moment functions $\boldsymbol{g}$; CVD inputs L and $(\bar{B}_k, \underline{B}_k)_{k=1}^K$

1: **for** each $\alpha \in \{1, 0.5, 0.1, 0\}$ **do**

2: **for** $i = 1, \ldots, 5$ **do**

3: Draw $\boldsymbol{\zeta} \sim \mathcal{N}(0, I_{|\mathcal{W}|})$

4: Uniformly select $\mathcal{N} \subseteq \{1, \ldots, |\mathcal{W}|\}$ with $|\mathcal{N}| = \text{round}(\alpha|\mathcal{W}|)$

5: Solve the following LP:

$$
\begin{aligned}
\min_{\boldsymbol{p} \in \mathbb{R}^{|\mathcal{W}|}} \quad & \sum_{j=1}^{|\mathcal{W}|} \zeta_j p_j \\
\text{s.t.} \quad & \sum_{j=1}^{|\mathcal{W}|} p_j = 1, \\
& \sum_{j=1}^{|\mathcal{W}|} p_j \boldsymbol{g}(\boldsymbol{w}^j) \leqslant \boldsymbol{\mu}_{\text{oos}}, \\
& p_j \geqslant 0, && j = 1, \ldots, |\mathcal{W}|, \\
& p_j \geqslant \delta, && j \in \mathcal{N}, \\
& LD(\bar{B}_k, \underline{B}_k) \leqslant \sum_{j:\, \boldsymbol{w}^j \in \bar{B}_k} p_j - \sum_{j:\, \boldsymbol{w}^j \in \underline{B}_k} p_j \leqslant LD(\bar{B}_k, \underline{B}_k), && k = 1, ..., K. \text{ (if with CVD)}
\end{aligned}
$$

6: If feasible, store $\boldsymbol{p}^\star$ with labels $(\boldsymbol{\mu}_{\text{oos}}, \alpha, i)$; otherwise report infeasibility

7: **end for**

8: **end for**

B.3.2. Parameters setting for CVD experiments. Here, we discuss two aspects of the CVD-regularization experiments in the continuous case, namely, the construction of the pairs of sets $(\bar{B}, \underline{B})$ and how $(\boldsymbol{\varepsilon}, \boldsymbol{\eta})$ are chosen in the computation of the bounds.

We construct two types of $(\bar{B}, \underline{B})$ pairs in (22) based on the worst-case sample paths of (AURA) with parent (20), where ambiguity collapse occurs. Given the optimal policy $\boldsymbol{x}^*$, these paths are identified as $\mathcal{W}^{\text{wc}} := \arg\max_{\boldsymbol{w} \in \mathcal{W}} \sum_{t=1}^T c_t\big(\boldsymbol{x}_t^*(\boldsymbol{w}_{\leqslant t})\big)$. Let $\mathcal{H}^{\text{wc}} := \{\boldsymbol{w}_{\leqslant T-1} : \boldsymbol{w} \in \mathcal{W}^{\text{wc}}\}$ denote the corresponding worst-case histories. For each $\boldsymbol{h} \in \mathcal{H}^{\text{wc}}$, define

$$
\mathcal{W}_{\boldsymbol{h}}^{\text{wc}} := \{\boldsymbol{w} \in \mathcal{W}^{\text{wc}} : \boldsymbol{w}_{\leqslant T-1} = \boldsymbol{h}\},
$$

$$
\mathcal{W}_{\boldsymbol{h}}^{\text{nwc}} := \{\boldsymbol{w} \in \mathcal{W} \backslash \mathcal{W}^{\text{wc}} : \boldsymbol{w}_{\leqslant T-1} = \boldsymbol{h}\}.
$$

Type I (within-history): For each worst-case history $\boldsymbol{h}$, we construct $(\bar{B}_k, \underline{B}_k)$ from $\mathcal{W}_{\boldsymbol{h}}^{\text{wc}}$ and $\mathcal{W}_{\boldsymbol{h}}^{\text{nwc}}$, respectively. These constraints limit probability-mass concentration on worst-case paths relative to alternative continuations under the same history. A sufficiently restrictive L can exclude the

targeted CC distributions.

Type II (cross-history): For each worst-case history $\boldsymbol{h}$, we construct pairs from $\mathcal{W}_{\boldsymbol{h}}^{\text{wc}}$ and $\left\{\boldsymbol{w}_{\leqslant T-1} \in \hat{\mathcal{W}}_{\leqslant T-1} : \boldsymbol{w}_{\leqslant T-1} \neq \boldsymbol{h}\right\}$. These constraints further restrict probability-mass imbalances across histories, limiting the ability of parent distributions to exploit arbitrarily small marginal probabilities of worst-case histories to induce deterministic conditional distributions.