Multi-Stage Robust Mixed-Integer Programming

Krzysztof Postek[1], Ward Romeijnders[2], and Wolfram Wiesemann[3]

[1] *Delft University of Technology, The Netherlands,* k.s.postek@tudelft.nl

[2] *University of Groningen, The Netherlands,* w.romeijnders@rug.nl

[3] *Imperial College London, United Kingdom,* ww@imperial.ac.uk

March 31, 2023

Abstract

Multi-stage robust optimization, in which decisions are taken sequentially as new information becomes available about the uncertain problem parameters, is a very versatile yet computationally challenging paradigm for decision-making under uncertainty. In this paper, we propose a new model and solution approach for multi-stage robust mixed-integer programs, which may contain both continuous and discrete decisions in any time stage. Our model builds upon the finite adaptability scheme developed for two-stage robust optimization problems, and it allows us to decompose the multi-stage problem into a large number of much simpler two-stage problems. We discuss how these two-stage problems can be solved both exactly and approximately, and we report numerical results for route planning and location-transportation problems.

Keywords: Robust optimization; multi-stage problems; mixed-integer optimization.

1 Introduction

Real-life decisions almost inevitably need to be taken under considerable uncertainty about key problem parameters, such as future customer demands, raw material prices, exchange rates, equipment outages and traffic conditions. Among the many paradigms that have been developed to safeguard decisions against uncertainty, such as simulation-based optimization, stochastic programming and Markov decision processes, the relatively young field of robust optimization stands out due to its promise to scale to the large problem sizes commonly encountered in application areas. This is achieved by replacing probabilistic descriptions of the uncertainty with a set-based characterization under which the uncertain problem parameters can attain any value from a pre-specified *uncertainty set*, and a decision is sought that performs best in view of the worst anticipated parameter realizations (Ben-Tal et al., 2009; Bertsimas et al., 2011a; Bertsimas and den Hertog, 2022).

While robust optimization has initially been developed for single-stage problems where all decisions are taken here-and-now, subsequent research has studied multi-stage formulations where the uncertain problem parameters are revealed gradually over time, and future recourse decisions can depend on the parameters that have already been observed. In contrast to single-stage robust optimization problems, which can often be solved in polynomial time, multi-stage robust optimization problems are NP-hard even if only two stages are considered, all decisions are continuous and the objective function as well as all constraints are linear (Guslitser, 2002). Two-stage robust optimization problems with continuous recourse can be solved exactly via Benders' decomposition (Jiang et al., 2010; Thiele et al., 2010; Bertsimas et al., 2013; Zhao et al., 2013), column-and-constraint generation (Zeng and Zhao, 2013; Ayoub and Poss, 2016), Fourier-Motzkin elimination (Zhen et al., 2018), copositive programming (Xu and Burer, 2018; Hanasusanto and Kuhn, 2018) or iteratively lifting the uncertainty set (Georghiou et al., 2020). None of these techniques scale well beyond two stages, however, and multi-stage robust optimization problems with more than two stages are typically conservatively approximated by restricting the recourse decisions to affine (Guslitser, 2002; Kuhn et al., 2011), piecewise affine (Chen et al., 2008; Chen and Zhang, 2009; Goh and Sim, 2010; Georghiou et al., 2015), polynomial or trigonometric (Bertsimas et al., 2011b) functions of the observed parameters (the so-called *decision rules*). More recently, multi-stage robust optimization problems have also been solved exactly by an adaptation of stochastic dual dynamic programming (Georghiou et al., 2019). We refer to Ben-Tal et al. (2009), Delage and Iancu (2015), Yanıkoğlu et al. (2019) and Bertsimas and den Hertog (2022) for surveys of the literature.

All of the aforementioned approaches have in common that they require the recourse decisions to be continuous. The ubiquitous presence of integer recourse decisions in practical applications, such as whether or not to place an order, build a facility or change the operation of a plant, has sparked significant interest in the development of solution schemes for multi-stage robust mixed-integer problems. However, the presence of integer recourse decisions prohibits a straightforward application of the aforementioned decision rules, and it precludes the use of strong convex duality results upon which most of the tractable reformulations from the continuous robust optimization literature rely. Two-stage robust optimization problems with mixed-integer recourse decisions have been solved exactly by semi-infinite programming techniques (Zhao and Zeng, 2012) as well as approximately by K-adaptability schemes (Bertsimas and Caramanis, 2010; Hanasusanto et al., 2015; Subramanyam et al., 2020) that restrict the choice of the second-stage decisions to one out of K candidate solutions which are optimized over in the first stage. We are not aware of any successful attempts to generalize either of these techniques to multi-stage robust mixed-integer problems with more than two time stages, however.

In this paper, we propose to approximate multi-stage robust mixed-integer programs by a finite adaptability formulation. Our formulation selects in each time stage the best mixed-integer state decision from a finite set of pre-selected candidate decisions. Continuous control decisions that only affect the feasibility and optimality within a stage, on the other hand, can be selected optimally from pre-defined polyhedral regions. We show that in contrast to the original multi-stage robust mixed-integer program, the finite adaptability approximation admits an equivalent nested formulation that can be solved via backward recursion. This allows us to reduce the monolithic multi-stage problem to a number of much simpler two-stage problems, many of which can be solved in parallel. We show how the arising two-stage problems can be solved exactly or approximately. We also discuss various heuristic strategies to select sets of candidate state decisions for each stage, and we show how to deterministically bound the suboptimality incurred by the current choice of candidate state decision sets. The contributions of this paper can be summarized as follows:

(i) We conservatively approximate multi-stage robust mixed-integer programs via a finite adaptability approximation that admits an equivalent decomposition into two-stage subproblems. To our best knowledge, this is the first approach that decomposes multi-stage robust mixed-integer programs into smaller subproblems that can be solved independently.

(ii) The subproblems of our decomposition scheme constitute two-stage robust optimization problems with a particularly benign structure. We show how these problems can be solved exactly through a disjunctive programming reformulation, as well as approximately by extending recent results from the literature on two-stage robust optimization.

(iii) We demonstrate the promise of our framework in the context of two numerical experiments: a route planning problem involving a graph with 100 nodes and 40 time stages as well as a location-transportation problem involving up to 20 facilities, 40 customer sites and 10 time stages. Our source codes as well as all data sets are released open-source to facilitate comparison with alternative approaches as well as reuse in applications.[1]

Our solution approach builds upon the multi-stage robust mixed-integer programming literature, which can be broadly categorized into two streams: *(i)* generalizations of the decision rule schemes developed for continuous problems and *(ii)* uncertainty set partitioning schemes. In the following, we summarize the most prominent approaches of both streams, and we subsequently explain how our proposed method differs from the literature.

The first decision rule architecture for multi-stage robust mixed-integer programs has been proposed by Bertsimas and Caramanis (2007). The authors model the recourse decisions as affine functions of features formed from the observed parameters. To ensure integrality of the discrete recourse decisions, the parameter realizations are rounded up to the closest integers, and the intercepts and slopes of the corresponding affine functions are restricted to integer numbers. The resulting semi-infinite mixed-integer linear program (MILP) is solved by constraint sampling, which results in probabilistic feasibility and optimality guarantees. Bertsimas and Georghiou (2015), on the other hand, model the continuous recourse decisions as piecewise affine functions of the observed parameters, whereas the discrete decisions (which are assumed to be binary) attain the value 1 precisely when certain piecewise affine functions of the parameters are non-negative. The problem can be formulated as a semi-infinite MILP that is solved by an iterative procedure which alternates between determining the optimal decision rules for a fixed set of parameter realizations and identifying new worst-case parameter realizations for the updated decision rules. For computational reasons, the algorithm is typically terminated prematurely, which implies that the solution may violate the constraints for some scenarios, and the true worst-case costs may exceed the worst-case costs estimated by the procedure. For the same problem class, Bertsimas and Georghiou (2018) model the binary recourse decisions as affine functions of non-anticipative binary features formed from the observed parameters. Binarity of the recourse decisions is achieved by requiring the intercepts and slopes of these affine functions to be integer and by restricting the image of the decision rules to be

[1] Website: `http://www.doc.ic.ac.uk/~wwiesema/multi-stage-ro.zip`.

binary. Using convex duality arguments, the authors obtain a finite-dimensional MILP whose size grows exponentially in general but remains polynomial if the uncertainty set is a hyperrectangle and each feature mapping involves a single parameter. The approach has recently been extended to endogenous (decision-dependent) uncertainty by Feng et al. (2021).

The first uncertainty set partitioning scheme for multi-stage robust mixed-integer programs has been proposed by Vayanos et al. (2011), who pre-partition the uncertainty set in each time stage into hyperrectangles and select affine decision rules (for the continuous recourse decisions) as well as constant decision rules (for the integer recourse decisions) over each hyperrectangle. By invoking convex duality arguments, the authors solve the corresponding conservative approximation exactly as a finite-dimensional MILP. The authors also discuss how their approach generalizes to the presence of endogenous (decision-dependent) uncertainty. Subsequent extensions of the uncertainty set partitioning paradigm have continued to rely on affine and constant decision rules for the continuous and discrete recourse decisions, respectively, but they seek to refine the partitions adaptively in view of the incumbent solutions. Postek and den Hertog (2016) alternate between determining the best decision rules for a fixed partition of the uncertainty set and identifying subsets of the partition where multiple parameter realizations result in binding constraints, thus indicating the need to further subdivide these subsets to obtain better decisions. In the multi-stage case, the stage-wise partitions form a tree structure that is similar to the scenario trees in stochastic programming. This partitioning scheme is very effective when the adaptive problem at hand has a small number of time stages, but it exhibits an exponential growth if many time stages are involved. Romeijnders and Postek (2021) refine the splitting technique of Postek and den Hertog (2016) in the presence of integer recourse decisions, where subsets of the partition may need to be split even if they do not generate binding constraints. The authors show that critical parameter realizations can be identified from the LP relaxations in the branch-and-bound tree that determines the decision rules. Bertsimas and Dunning (2015) propose an alternative scheme where the uncertainty set in each time stage is partitioned by a Voronoi diagram that is constructed from the binding scenarios, thus eliminating the need to choose splits manually. Since their approach splits each subset of the partition in every iteration, however, the size of the partition grows quickly in the number of algorithm iterations even for two-stage problems.

In contrast to the aforementioned decision rule architectures and uncertainty set partitioning schemes, Goerigk and Hartisch (2021) model multi-stage robust pure integer programs with pure integer uncertainties as quantified integer linear programs, which constitute two-person zero-sum games between an existential player (the 'minimizer') and a universal player (the 'maximizer'). The resulting problems can be solved with an open-source quantified integer programming solver.

Multi-stage mixed integer problems have also been studied in the related literature on stochastic programming, where the uncertain problem parameters are assumed to be governed by a known (typically discrete) probability distribution. Most approaches in this domain relax some of the constraints in the scenario tree representation of the problem, which leads to scenario, component and nodal decomposition approaches; see Klein Haneveld and van der Vlerk (1999), Römisch and Schultz (2001), Schultz (2003), Sen (2005) and Boland et al. (2016). Under the assumption of stage-wise independent problem parameters, multi-stage mixed integer stochastic programs are also amenable to extensions of the stochastic dual dynamic programming scheme developed for convex problems (Pereira and Pinto, 1991; Shapiro, 2011). Since the cost to-go functions are no longer convex, the affine cutting planes from stochastic dual dynamic programming are no longer applicable, and they are replaced with step functions (Philpott et al., 2020), nonlinear Lipschitz cuts (Ahmed et al., 2020) or generalized conjugacy cuts (Zhang and Sun, 2022). Alternatively, Zou et al. (2019) show that the convex lower envelope of the cost to-go function remains piecewise affine and convex if the problem has a complete recourse and only contains binary state variables; in this case, cuts generated from Lagrangian relaxations can also be applied. The aforementioned stochastic programming approaches crucially rely on the ability to enumerate all uncertainty realizations explicitly or implicitly as a scenario tree, however, and it is unclear how they would extend to the set-based descriptions of the uncertainty that are employed in robust optimization.

The method developed in this paper complements the existing solution schemes for multi-stage robust mixed-integer programs. The uncertainty set partitioning approaches from the literature are essentially free of hyperparameters and thus do not require any *a priori* knowledge about the problem. On the flip side, they construct set-based analogues of scenario trees that exhibit an exponential growth in the number of time stages as well as, typically, the number of uncertain problem parameters per stage. They are thus ideally suited for smaller problems where the decision maker has little *a priori* insight into the structure of well-performing solutions. Most of the decision rule architectures, on the other hand, require the functional form of the recourse decisions to be selected upfront and thus rely on domain knowledge of the decision maker. In contrast to the uncertainty set partitioning approaches, however, decision rule architectures tend to scale polynomially in the number of decision variables and time stages, which renders them particularly promising for larger problems. Similar to the decision rule architectures, our method requires *a priori* knowledge about the problem to select suitable candidate decisions for each time stage. On the other hand, since our approach decomposes the overall problem into smaller two-stage subproblems, our method scales particularly well in the number of time stages. Moreover, and in contrast to all of the existing approaches, our method is ideally suited for parallelization since many of the subproblems can be solved in parallel. This is attractive in view of the recent growth in cloud computing services, which enable users to rent vast amounts of parallel computing resources at an hourly billing.

The remainder of this paper proceeds as follows. Section 2 formulates the multi-stage robust mixed-integer program of interest, it presents our finite adaptability approximation, and it elucidates how this approximation admits a decomposition into two-stage subproblems. Sections 3 and 4 discuss exact and approximate solution approaches for the two-stage subproblems, respectively. Section 5 presents a progressive bound to estimate the suboptimality of our approximation, and it develops heuristic strategies to select candidate state decision sets for each stage. We conclude with numerical experiments in Section 6. All proofs are relegated to the appendix.

Notation. For a vector $\boldsymbol{\xi} = (\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_T)$ constructed from T subvectors $\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_T$, we denote by $\boldsymbol{\xi}_t$ the t-th subvector, whereas $\boldsymbol{\xi}^t = (\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_t)$ stacks all subvectors up to and including $\boldsymbol{\xi}_t$. The vectors $\mathbf{e}$ and $\mathbf{e}_i$ refer to the all-ones and the i-th basis vector, respectively; their dimension will be clear from the context. Finally, we denote by $\mathbf{1}[\cdot]$ the indicator function that attains the value 1 if the logical expression $\cdot$ is satisfied and 0 otherwise.

2 Problem Formulation

We are interested in *multi-stage robust mixed-integer optimization problems* of the form

$$
\begin{array}{ll}
\text{minimize} & \displaystyle\max_{\boldsymbol{\xi} \in \Xi} \sum_{t=1}^{T} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t(\boldsymbol{\xi}^t) \\
\text{subject to} & \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_{t-1}(\boldsymbol{\xi}^{t-1}) + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t(\boldsymbol{\xi}^t) + \boldsymbol{V}_t\, \boldsymbol{y}_t(\boldsymbol{\xi}^t) \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t) \qquad \forall \boldsymbol{\xi} \in \Xi,\ \forall t = 1, \ldots, T \\
& \boldsymbol{x}_t(\boldsymbol{\xi}^t) \in \mathcal{X}_t \text{ and } \boldsymbol{y}_t(\boldsymbol{\xi}^t) \in \mathbb{R}^{n_t^2} \text{ for all } \boldsymbol{\xi} \in \Xi \text{ and } t = 1, \ldots, T,
\end{array}
\tag{1}
$$

where the uncertain problem parameters $\boldsymbol{\xi}_t$ are revealed at the beginning of each time stage t, $t = 1, \ldots, T$, and the decisions $\boldsymbol{x}_t$ and $\boldsymbol{y}_t$ are taken afterwards under the full knowledge of $\boldsymbol{\xi}^t = (\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_t)$. A solution to problem (1) is immunized against all parameter realizations $\boldsymbol{\xi} = \boldsymbol{\xi}^T$ in the uncertainty set Ξ. The cost vectors $\boldsymbol{q}_t$, the technology matrices $\boldsymbol{T}_t$, the recourse matrices $\boldsymbol{W}_t$ as well as the right-hand side vectors $\boldsymbol{h}_t$ may all depend affinely on $\boldsymbol{\xi}_t$. We stipulate that $\boldsymbol{T}_1(\boldsymbol{\xi}_0) \equiv \mathbf{0}$ and $\boldsymbol{x}_0(\boldsymbol{\xi}^0) \equiv \mathbf{0}$, that is, the first-stage constraints only involve the first-stage decisions $\boldsymbol{x}_1$ and $\boldsymbol{y}_1$. We assume that $\boldsymbol{\xi}_1$ is deterministic, and hence $\boldsymbol{x}_1$ and $\boldsymbol{y}_1$ are here-and-now decisions. We can interpret $\boldsymbol{x}_t$ as *state variables* since they couple constraints of consecutive stages, whereas $\boldsymbol{y}_t$ are *control variables* that only affect the costs and feasibility within stage t. Each set $\mathcal{X}_t \subseteq \mathbb{R}^{n_t^1}$ can be non-convex and comprise both continuous and discrete decision variables. In contrast, the linear constraints of problem (1) fully characterize the set of admissible $\boldsymbol{y}_t$. To ease the exposition, we assume that the set of admissible $\boldsymbol{y}_t$ is bounded; this avoids tedious but otherwise straightforward case distinctions where infinite cost savings in one stage need to be weighed against infeasibility (which can be regarded as infinite costs) in another stage. We do *not* assume that problem (1) has a relatively complete recourse, that is, there may be partial solutions $(\boldsymbol{x}_\tau, \boldsymbol{y}_\tau)_{\tau=1}^t$ satisfying the constraints up to time stage t that cannot be extended to complete solutions $(\boldsymbol{x}_\tau, \boldsymbol{y}_\tau)_{\tau=1}^T$ satisfying all constraints up to the final time stage T. The presence of potentially non-convex stage-wise feasible regions $\mathcal{X}_t$ renders problem (1) strongly NP-hard even in the absence of uncertainty. Even if the stage-wise feasible regions $\mathcal{X}_t$ are polyhedra, however, problem (1) remains strongly NP-hard due to the presence of adaptive decisions, see Guslitser (2002, Theorem 3.5) and Subramanyam et al. (2020, Proposition A.3).

The uncertainty set Ξ comprises discrete *inter-stage uncertainties* $\boldsymbol{\phi}_t \in \mathbb{R}^{k_t^1}$ that can be coupled over consecutive time stages as well as continuous *intra-stage uncertainties* $\boldsymbol{\psi}_t \in \mathbb{R}^{k_t^2}$ that are stage-wise rectangular, apart from their potential dependence on $\boldsymbol{\phi}_t$. More precisely, Ξ emerges from stage-wise uncertainty sets Ξ_t as follows. The parameters of the first stage are deterministic and are thus the only element of the singleton set $\Xi_1 = \{\boldsymbol{\xi}_1 = (\boldsymbol{\phi}_1, \boldsymbol{\psi}_1)\}$. For $t = 2, \ldots, T$, the stage-wise uncertainty sets satisfy

$$\Xi_t(\boldsymbol{\phi}_{t-1}) \;=\; \{\boldsymbol{\xi}_t = (\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \,:\, \boldsymbol{\phi}_t \in \Phi_t(\boldsymbol{\phi}_{t-1}),\;\; \boldsymbol{U}_t(\boldsymbol{\phi}_t)\, \boldsymbol{\psi}_t \leqslant \boldsymbol{b}_t(\boldsymbol{\phi}_t)\},$$

where $\boldsymbol{U}_t : \mathbb{R}^{k_t^1} \to \mathbb{R}^{l_t \times k_t^2}$ and $\boldsymbol{b}_t : \mathbb{R}^{k_t^1} \to \mathbb{R}^{l_t}$ can be arbitrary functions of $\boldsymbol{\phi}_t$. To keep the notation consistent, we stipulate that $\boldsymbol{\phi}_0 \equiv \mathbf{0}$, $\Phi_1(\boldsymbol{\phi}_0) \equiv \Phi_1 = \{\boldsymbol{\phi}_1\}$ and $\Xi_1(\boldsymbol{\phi}_0) \equiv \Xi_1$. We assume that the sets $\Phi_t(\boldsymbol{\phi}_{t-1})$ are finite and their unions $\Phi_t = \bigcup_{\boldsymbol{\phi}_{t-1} \in \Phi_{t-1}} \Phi_t(\boldsymbol{\phi}_{t-1})$, $t = 2, \ldots, T$, are 'not too large', in a sense that we will discuss in more detail later on. We can then define the overall uncertainty set Ξ from the stage-wise uncertainty sets Ξ_t as

$$\Xi \;=\; \big\{\boldsymbol{\xi} = (\underbrace{(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)}_{=\boldsymbol{\xi}_t})_{t=1}^T \;:\; \boldsymbol{\xi}_1 \in \Xi_1,\;\; \boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1}) \;\; \forall t = 2, \ldots, T\big\}. \tag{2}$$

Note that Ξ may fail to be stage-wise rectangular, that is, $\Xi \subsetneq \times_{t=1}^T \{\boldsymbol{\xi}_t \,:\, \boldsymbol{\xi} \in \Xi\}$. To avoid technical but otherwise straightforward case distinctions, we assume in the following that Ξ is bounded.

The uncertainty set (2) is quite versatile, and it complements the relatively scarce literature on multi-stage uncertainty sets (Miao et al., 2007; Delage and Iancu, 2015; Lorca and Sun, 2015).

Example 1 (Uncertainty Set Ξ)**.** *In the multi-stage robust optimization literature,* stage-wise rectangular uncertainty sets *of the form $\Xi = \Xi_1 \times \ldots \times \Xi_T$, where $\Xi_1 = \{\boldsymbol{\xi}_1\}$ and each set Ξ_t is a polyhedron, constitute a common choice of uncertainty sets. Stage-wise rectangular uncertainty sets are readily recognized as a special case of (2) if we disregard the inter-stage uncertainties $\boldsymbol{\phi}_t$ and choose $\boldsymbol{U}_t$ and $\boldsymbol{b}_t$ such that $\{\boldsymbol{\psi}_t \,:\, \boldsymbol{U}\boldsymbol{\psi}_t \leqslant \boldsymbol{b}_t\} = \Xi_t$. Despite their popularity, we argue that stage-wise rectangular uncertainty sets can lead to overly conservative solutions as they allow for the worst parameter values to be realized in every single time stage.*

To alleviate the conservatism of stage-wise rectangular uncertainty sets, Bandi and Bertsimas (2012) propose uncertainty sets inspired by central limit theorems*, where it is assumed that the cumulative absolute deviations of the uncertain parameters from some nominal values are bounded. A univariate discrete version of their uncertainty set can be modeled as*

$$\mathcal{Z} \;=\; \left\{\boldsymbol{\zeta} \in \mathbb{R}^T \,:\, \zeta_1 = \zeta_1^0,\;\; \zeta_t \in \left\{\underline{\zeta}_t, \ldots, \overline{\zeta}_t\right\} \;\; \forall t = 2, \ldots T,\;\; \underline{\Gamma} \leqslant \sum_{t=1}^T \zeta_t \leqslant \overline{\Gamma}\right\},$$

where ζ_t is the uncertain parameter in stage t, $(\underline{\zeta}_t, \overline{\zeta}_t)$ are the stage-wise bounds and $(\underline{\Gamma}, \overline{\Gamma})$ characterize the maximum absolute deviations from some nominal values. Note that this uncertainty set includes the well-known discrete budget uncertainty sets *(Bertsimas and Sim, 2004) as a special case. Central limit theorem uncertainty sets emerge if we disregard the intra-stage uncertainties $\boldsymbol{\psi}_t$ in (2) and choose $\boldsymbol{\phi}_t = (\phi'_t, \phi^\Sigma_t)$ with $\Phi_1 = \{(\zeta_1^0, \zeta_1^0)\}$ as well as*

$$\Phi_t(\boldsymbol{\phi}_{t-1}) = \left\{\boldsymbol{\phi}_t = (\phi'_t, \phi^\Sigma_t) \,:\, \exists \boldsymbol{\zeta} \in \mathcal{Z} \text{ such that } \sum_{\tau=1}^{t-1} \zeta_\tau = \phi^\Sigma_{t-1},\;\; \zeta_t = \phi'_t \text{ and } \phi^\Sigma_t = \phi^\Sigma_{t-1} + \phi'_t\right\},$$

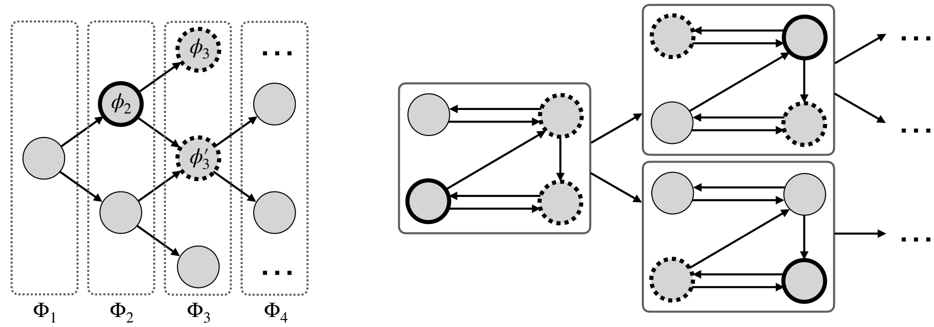

Figure 1. Lattice regime (left) and Markovian regime chain (right) uncertainty sets. In both cases, the regimes are represented as nodes. In the left figure, we have $\Phi_3(\phi_2) = \{\phi_3, \phi_3'\}$. In the right figure, the sets $\Phi_{t+1}(\phi_t)$ corresponding to the nodes ϕ_t highlighted by bold strokes are the nodes with dotted strokes.

$t = 2, \ldots, T$, where ϕ_t' and ϕ_t^{Σ} represent the innovation term at stage t as well as the summation over all previously observed innovations, respectively. The uncertainty set (2) *can readily model multivariate versions of the discrete central limit theorem uncertainty sets as well, albeit at the expense of combinatorially scaling inter-stage uncertainty sets Φ_t.*

Chen et al. (2019), Long et al. (2023) and Cui et al. (2023) propose scenario-based uncertainty sets *for distributionally robust optimization problems, where a random parameter vector can be governed by different moment ambiguity sets with associated (and possibly partially unknown) scenario probabilities. A robust multi-stage version of this uncertainty set,*

$$\mathcal{Z} \;=\; \left\{ \boldsymbol{\zeta} = (\boldsymbol{\zeta}_1, \ldots, \boldsymbol{\zeta}_T) \,:\, \boldsymbol{\zeta}_1 = \boldsymbol{\zeta}_1^0, \quad \begin{bmatrix} \boldsymbol{F}_t^1 \boldsymbol{\zeta}_t \leqslant \boldsymbol{g}_t^1 \\ \vee \qquad \vdots \\ \vee\; \boldsymbol{F}_t^{s_t} \boldsymbol{\zeta}_t \leqslant \boldsymbol{g}_t^{s_t} \end{bmatrix} \quad \forall t = 2, \ldots, T \right\},$$

arises as a special case of (2) *if the scalar inter-stage uncertainty can attain the values $\phi_t \in \{1, \ldots, s_t\}$ and we set $\boldsymbol{U}_t(\phi_t) = \boldsymbol{F}_t^{\phi_t}$ as well as $\boldsymbol{b}_t(\phi_t) = \boldsymbol{g}_t^{\phi_t}$ in every stage $t = 2, \ldots, T$. We can interpret this uncertainty set as a union of scenario polyhedra in every time stage.*

Finally, uncertainty sets of the form (2) *also allow us to model* regime uncertainty sets *that to our best knowledge have not been previously studied in the robust optimization literature. In a* lattice regime uncertainty set*, the discrete regime ϕ_t in stage t is any element of $\Phi_t(\phi_{t-1})$, where $\Phi_t(\phi_{t-1})$ describes the children of the previous regime node ϕ_{t-1} in a lattice structure (*cf. *Figure 1, left). Likewise, in a* Markovian regime chain uncertainty set*, the regime ϕ_t is any of the neighbours $\Phi_t(\phi_{t-1})$ of the previous regime node ϕ_{t-1} in a graph (*cf. *Figure 1, right). In both cases, the discrete* *regimes can be complemented by continuous intra-regime uncertainties.*

Instead of solving the multi-stage robust mixed-integer problem (1) directly, we propose to investigate the following *finite adaptability approximation*,

$$\begin{array}{ll} \text{minimize} & \displaystyle\max_{\boldsymbol{\xi}\in\Xi} \sum_{t=1}^{T} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t(\boldsymbol{\xi}^t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_{t-1}(\boldsymbol{\xi}^{t-1}) + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t(\boldsymbol{\xi}^t) + \boldsymbol{V}_t\,\boldsymbol{y}_t(\boldsymbol{\xi}^t) \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t) \qquad \forall \boldsymbol{\xi}\in\Xi,\ \forall t=1,\ldots,T \\ & \boldsymbol{x}_t(\boldsymbol{\xi}^t) \in \mathcal{X}_t^{\mathrm{C}} \text{ and } \boldsymbol{y}_t(\boldsymbol{\xi}^t) \in \mathbb{R}^{n_t^2} \text{ for all } \boldsymbol{\xi}\in\Xi \text{ and } t=1,\ldots,T, \end{array} \tag{3}$$

where $\mathcal{X}_t^{\mathrm{C}} = \{\boldsymbol{x}_t^1, \ldots, \boldsymbol{x}_t^{p_t}\} \subseteq \mathcal{X}_t$ is a finite set of pre-selected candidate decisions for stage t. The finite adaptability approximation (3) differs from the multi-stage robust mixed-integer problem (1) in that it selects optimal state decisions $\boldsymbol{x}_t$ from the restricted finite sets $\mathcal{X}_t^{\mathrm{C}}$ as opposed to the original sets $\mathcal{X}_t$. This conservative approximation turns out to be crucial for the development of our nested problem formulation and its associated solution method below. Note that in contrast to the K-adaptability schemes for two-stage robust problems (Bertsimas and Caramanis, 2010; Hanasusanto et al., 2015; Subramanyam et al., 2020), the candidate decision sets $\mathcal{X}_t^{\mathrm{C}}$ in problem (3) are fixed. Similar to the choice of decision rule architectures discussed in the introduction, the choice of suitable candidate decision sets $\mathcal{X}_t^{\mathrm{C}}$ requires some *a priori* insight into the structure of well-performing solutions to the multi-stage robust mixed-integer problem (1). We will discuss in Section 5 different heuristic approaches for choosing candidate decision sets $\mathcal{X}_t^{\mathrm{C}}$ if such *a priori* knowledge is not available. We emphasize that problem (3) selects the best decisions from each set $\mathcal{X}_t^{\mathrm{C}}$ adaptively based on the realization of $\boldsymbol{\xi}^t$, and that the problem therefore remains challenging as it generalizes a known NP-hard problem (Subramanyam et al., 2020, Proposition B.3).

We now discuss how the finite adaptability approximation (3) gives rise to an equivalent nested problem formulation that can be solved by a backward recursion. Our result relies on the following interchangeability principle established by Shapiro (2017, Proposition 2.1).

Lemma 1. *For a set $\Omega \subset \mathbb{R}^m$, a set-valued mapping $F : \Omega \rightrightarrows \mathbb{R}^n$, the set of functions $\mathcal{F} = \{[\boldsymbol{x} : \Omega \to \mathbb{R}^n] \,:\, \boldsymbol{x}(\boldsymbol{\omega}) \in F(\boldsymbol{\omega}) \;\; \forall \boldsymbol{\omega} \in \Omega\}$ and a cost function $c : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$, we have*

$$\min_{\boldsymbol{x}\in\mathcal{F}} \max_{\boldsymbol{\omega}\in\Omega} \; c(\boldsymbol{x}(\boldsymbol{\omega}), \boldsymbol{\omega}) \quad = \quad \max_{\boldsymbol{\omega}\in\Omega} \; \min_{\boldsymbol{x}\in F(\boldsymbol{\omega})} \; c(\boldsymbol{x}, \boldsymbol{\omega}) \tag{4}$$

as long as c attains its minimum over $\boldsymbol{x} \in F(\boldsymbol{\omega})$ for all $\boldsymbol{\omega} \in \Omega$.

Note that Lemma 1 does not require any convexity assumptions. Intuitively, any policy $\boldsymbol{x} \in \mathcal{F}$ in the min-max problem on the left-hand side of (4) induces a feasible (but possibly suboptimal) response $\boldsymbol{x}(\boldsymbol{\omega})$ in the max-min problem on the right-hand side of (4), and conversely a collection of responses $\{\boldsymbol{x}(\boldsymbol{\omega})\}_{\boldsymbol{\omega} \in \Omega}$ to the max-min problem in (4) gives rise to a feasible (but possibly suboptimal) policy $\boldsymbol{x} \in \mathcal{F}$ in the min-max problem in (4). Lemma 1 extends to the case where $F(\boldsymbol{\omega}) = \emptyset$ for some $\boldsymbol{\omega} \in \Omega$ if we stipulate that both sides of the identity (4) evaluate to $+\infty$ in that case.

Our nested formulation is defined by the *stage-t worst-case cost to-go problem*

$$\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1}) \;=\; \max_{\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})} \; Q_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t), \tag{5a}$$

$t = 1, \ldots, T$, where the *stage-t nominal cost to-go problem* satisfies

$$Q_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t) \;=\; \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \mathcal{Q}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t) \\ & \boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}}, \; \boldsymbol{y}_t \in \mathbb{R}^{n_t^2} \end{array}\right]. \tag{5b}$$

In problem (5b), we stipulate that $Q_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t) = +\infty$ if the minimization is infeasible, and we set the boundary condition $\mathcal{Q}_{T+1}(\boldsymbol{x}_T; \boldsymbol{\phi}_T) \equiv 0$ for all $\boldsymbol{x}_T \in \mathcal{X}_T^{\mathrm{C}}$ and $\boldsymbol{\phi}_T \in \Phi_T$. For $t = 1$, we also set $\hat{\boldsymbol{x}}_0 \equiv \boldsymbol{0}$ and abbreviate $\mathcal{Q}_1(\hat{\boldsymbol{x}}_0, \boldsymbol{\phi}_0)$ by $\mathcal{Q}_1$.

Proposition 1. *The finite adaptability approximation* (3) *and the stage-*1 *worst-case cost to-go problem $\mathcal{Q}_1$ share the same optimal value. In particular,* (3) *is infeasible if and only if $\mathcal{Q}_1 = \infty$.*

The equivalence of multi-stage problems and their nested formulations is well understood in the stochastic programming community when the worst-case approach is replaced with the expected value. In that case, the pendant of Lemma 1 has been established by Rockafellar and Wets (1998, Theorem 14.60), and the equivalence of both formulations is typically taken as given. Shapiro (2017) studies the more general case where the expected value is replaced with a generic risk measure. While this result includes our worst-case approach, the nested formulation of Shapiro (2017) does not account for any structure in the support of the random vectors, and his stage-wise cost to-go problems therefore depend on the entire parameter sequence $\boldsymbol{\xi}^T$. In contrast, the stage-t worst-case cost to-go problems in our derivations only depend on the realization $\boldsymbol{\phi}_{t-1}$ of the preceding inter-stage uncertainty. This parsimonious dependence of the worst-case cost to-go problems on the past uncertainty realizations lies at the heart of our approach; it is facilitated by the design of our generalized rectangularity of Ξ, and it is key for the tractability of our approach.

In the following, we first discuss how the stage-t worst-case cost to-go problems $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1};\boldsymbol{\phi}_{t-1})$ can be solved by a backward recursion (Algorithm 1). Here, the key step repeatedly solves two-stage subproblems $\mathcal{Q}_t$, given the (previously computed) values of the worst-case cost to-go problems $\mathcal{Q}_{t+1}$. Sections 3 and 4 are devoted to the exact and approximate solution of these subproblems.

Algorithm 1: Evaluation of the stage-t worst-case cost to-go problems $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1};\boldsymbol{\phi}_{t-1})$.

1. **Initialization.** Set $t = T$ as well as $\mathcal{Q}_{T+1}(\boldsymbol{x}_T;\boldsymbol{\phi}_T) = 0$ for all $\boldsymbol{x}_T \in \mathcal{X}_T^{\mathrm{C}}$ and $\boldsymbol{\phi}_T \in \Phi_T$.

2. **Iteration.** Compute $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1};\boldsymbol{\phi}_{t-1})$ for all $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ and $\boldsymbol{\phi}_{t-1} \in \Phi_{t-1}$.

3. **Termination.** If $t > 1$, update $t \leftarrow t-1$ and repeat Step 2. Otherwise, terminate.

Note that the subproblems $\mathcal{Q}_t$ occurring in the same time stage t are independent of another and can thus be solved in parallel, leading to a theoretical maximum speedup of $(1/T)\cdot\sum_{t=1}^{T}|\mathcal{X}_t^{\mathrm{C}}|\cdot|\Phi_{t-1}|$ over the sequential solution of all subproblems. Once Algorithm 1 has been executed, the optimal first-stage decisions $(\boldsymbol{x}_1^\star,\boldsymbol{y}_1^\star)$ can be obtained by solving a single instance of the nominal problem Q_1. In principle, Algorithm 1 is sufficient for solving the finite adaptability approximation (3) in a receding-horizon fashion, where upon observing $\boldsymbol{\xi}_t$ we construct and solve a new $(T-t+1)$-stage problem. However, since Algorithm 1 solves *all* worst-case cost to-go problems $\mathcal{Q}_t$, they do not need to be recomputed at later time stages. The following algorithm exploits this property.

Algorithm 2: Optimal policy for the finite adaptability approximation (3).

1. **Initialization.** Solve all stage-t worst-case cost to-go problems via Algorithm 1.

2. **First Stage.** Implement $(\boldsymbol{x}_1^\star,\boldsymbol{y}_1^\star)$ as determined by the problem $\mathcal{Q}_1$. Set $t = 2$.

3. **Iteration.** Upon observation of the uncertain parameters $\boldsymbol{\xi}_t$, solve the stage-t nominal cost to-go problem $Q_t(\boldsymbol{x}_{t-1}^\star;\boldsymbol{\xi}_t)$ and implement the corresponding stage-t decisions $(\boldsymbol{x}_t^\star,\boldsymbol{y}_t^\star)$.

4. **Termination.** If $t < T$, update $t \leftarrow t+1$ and repeat Step 3. Otherwise, terminate.

We next show that Algorithm 2 *implicitly* defines an optimal policy $\{(\boldsymbol{x}_t^\star,\boldsymbol{y}_t^\star)\}_{t=1}^{T}$ for the finite adaptability approximation (3) in response to a stage-wise revealed scenario $\boldsymbol{\xi} \in \Xi$.

Proposition 2. *Assume that problem* (3) *is feasible. Then the policy* $\{(\boldsymbol{x}_t^\star,\boldsymbol{y}_t^\star)\}_{t=1}^{T}$ *that is implicitly defined by Algorithm 2 in response to the parameter realizations* $\boldsymbol{\xi} \in \Xi$ *is optimal in* (3).

Remark 1 (Alternative Formulations). *While in theory a nested formulation akin to* (5) *could be formulated for a finite adaptability approximation that selects the candidate decision sets* $\mathcal{X}_t^{\mathrm{C}}$ *as part of the optimization, or even more generally for the multi-stage robust mixed-integer problem* (1)*, the corresponding backward recursion would have to compute the worst-case cost to-go for all possible decisions* $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}$*, as opposed to the candidate decisions* $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ *only, in order to quantify the future worst-case cost to-go in each two-stage subproblem. This unfavorable scaling would limit the applicability of the resulting solution scheme to very small problem instances.*

Remark 2 (Robust Markov Decision Processes). *When the continuous decisions* $\boldsymbol{y}_t$ *are absent, the multi-stage robust mixed-integer program* (1) *is reminiscent of a robust Markov decision process, that is, a Markov decision process where the transition probabilities are only known to reside in some uncertainty set (Iyengar, 2005; Nilim and Ghaoui, 2005; Wiesemann et al., 2013). A crucial difference between the two modeling paradigms is, however, that the discrete decisions* $\boldsymbol{x}_t$ *in* (1) *can be selected* after *the uncertain parameters* $\boldsymbol{\xi}_t$ *have been observed, whereas the same decisions would have to be taken* before $\boldsymbol{\xi}_t$ *is known in a robust Markov decision process. This seemingly minor difference is crucial as one can readily construct instances of* (1) *that are feasible under the multi-stage robust optimization paradigm but infeasible if modelled as a robust Markov decision process. Likewise, the applications in our numerical experiments would typically reduce to much simpler deterministic problems if they were modelled as robust Markov decision processes.*

3 Exact Solution of the Two-Stage Subproblems $\mathcal{Q}_t$

The key step in Algorithm 1 for our finite adaptibility approximation (3) is the solution of the two-stage subproblems $\mathcal{Q}_t$. The presence of discrete wait-and-see decisions $\boldsymbol{x}_t$ in these problems renders most of the existing exact solution approaches for two-stage robust optimization problems (*cf.* Section 1) inapplicable. Indeed, Fourier-Motzkin elimination and iterative uncertainty set lifting schemes fundamentally require all second-stage decisions to be continuous, while Benders' decomposition and column-and-constraint generation depend on strong convex duality of the second stage to identify worst-case uncertainty realizations for fixed first-stage decisions. To our best knowledge, the only exact solution approach for two-stage robust optimization problems with mixed-integer recourse is the nested column-and-constraint generation scheme of Zhao and Zeng (2012). Applied to our setting, however, this method would require the two-stage subproblems $\mathcal{Q}_t$

to be solved by extreme point uncertainty realizations $\boldsymbol{\xi}_t \in \operatorname{ext} \Xi_t(\boldsymbol{\phi}_{t-1})$, which is not guaranteed to be the case in our context. Indeed, one can readily construct instances of $\mathcal{Q}_t$ without control variables and either objective uncertainty or right-hand side uncertainty whose worst-case uncertainty realizations are not extreme points, and the two-stage subproblems in our two applications of Section 6 are not optimized by extreme point worst-case parameter realizations in general. Our problem $\mathcal{Q}_t$ is reminiscent of the K-adaptability problems studied by Bertsimas and Caramanis (2010), Hanasusanto et al. (2015) and Subramanyam et al. (2020). In contrast to those problems, however, the candidate decision sets $\mathcal{X}_t^{\mathrm{C}}$ in $\mathcal{Q}_t$ have finite cardinality. This allows us to design exact solution schemes despite the presence of continuous second-stage decisions $\boldsymbol{y}_t$.

To solve the two-stage subproblems $\mathcal{Q}_t$ exactly, we first derive a non-convex strong duality result for the stage-t nominal cost to-go problem Q_t that is embedded in the two-stage subproblem $\mathcal{Q}_t$ (Theorem 1), which subsequently gives rise to an equivalent bilinear programming formulation of $\mathcal{Q}_t$ (Proposition 3). We then identify and discuss four special cases of Proposition 3 that lead to LP or MILP reformulations of polynomial size: *(i)* the constraints of Q_t do not depend on the uncertainty realization $\boldsymbol{\xi}_t$; *(ii)* the control variables $\boldsymbol{y}_t$ are absent in Q_t; *(iii)* the recourse matrix $\boldsymbol{V}_t$ in Q_t is invertible; and *(iv)* the recourse matrix $\boldsymbol{V}_t$ in Q_t has a block-diagonal structure.

The benign structure of the stage-t nominal cost to-go problem Q_t allows us to derive a non-convex strong dual that forms the basis of our solution approaches for the two-stage subproblems $\mathcal{Q}_t$.

Theorem 1 (Strong Duality)**.** *Fix $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ and $\boldsymbol{\xi}_t$. The stage-t nominal cost to-go problem $Q_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t)$ has the same optimal value as the semi-infinite disjunctive program*

$$\begin{array}{ll} \text{maximize} & \theta \\ \text{subject to} & \Big[\theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \mathcal{Q}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \;\; \vee \\ & \;\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \not\geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t)\Big] \qquad \forall \boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}},\ \forall \boldsymbol{y}_t \in \mathbb{R}^{n_t^2} \\ & \theta \in \mathbb{R}. \end{array} \tag{6}$$

Problem (6) can be interpreted as follows. For every second-stage decision $(\boldsymbol{x}_t, \boldsymbol{y}_t) \in \mathcal{X}_t^{\mathrm{C}} \times \mathbb{R}^{n_t^2}$, θ either needs to account for the objective value $\boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \mathcal{Q}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t)$ of $(\boldsymbol{x}_t, \boldsymbol{y}_t)$, or $(\boldsymbol{x}_t, \boldsymbol{y}_t)$ has to violate at least one of the constraints $\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t)$.

We next leverage the semi-infinite non-convex dual (6) to derive an equivalent finite-dimensional single-stage reformulation of the two-stage subproblem $\mathcal{Q}_t$ as a bilinear program. To this end, we from now on denote by m_t the number of second-stage constraints $\boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t + \boldsymbol{V}_t\,\boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t)$ in stage t.

Proposition 3. *Fix $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ and $\boldsymbol{\phi}_{t-1} \in \Phi_{t-1}$. If the stage-$t$ nominal cost to-go problem $Q_t(\hat{\boldsymbol{x}}_{t-1};\boldsymbol{\xi}_t)$ is feasible for some $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$, then the two-stage subproblem $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1};\boldsymbol{\phi}_{t-1})$ has the equivalent reformulation*

$$
\begin{array}{ll}
\text{maximize} & \theta \\
\text{subject to} & \theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \left[\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\hat{\boldsymbol{x}}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t^i\right]^\top \boldsymbol{\lambda}_t^i \\
& \qquad\qquad\qquad\qquad + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i;\boldsymbol{\phi}_t) \qquad \forall i = 1,\ldots,p_t \\
& \boldsymbol{V}_t^\top \boldsymbol{\lambda}_t^i = \boldsymbol{r}_t \qquad\qquad\qquad\qquad\qquad\qquad\quad \forall i = 1,\ldots,p_t \\
& \theta \in \mathbb{R},\;\; \boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1}),\;\; \boldsymbol{\lambda}_t^i \in \mathbb{R}_+^{m_t},\; i = 1,\ldots,p_t.
\end{array}
\tag{7}
$$

Problem (7) constitutes a mixed-integer bilinear program and is as such difficult to solve in general. We next study subclasses of problem (7) that admit practically efficient solution schemes.

Corollary 1 (Deterministic Constraints)**.** *Fix $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ and $\boldsymbol{\phi}_{t-1} \in \Phi_{t-1}$. If the constraints of the stage-t nominal cost to-go problem Q_t do not depend on $\boldsymbol{\xi}_t$, then the two-stage subproblem $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1};\boldsymbol{\phi}_{t-1})$ has the equivalent MILP reformulation*

$$
\begin{array}{ll}
\text{maximize} & \theta \\
\text{subject to} & \theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + r_t^i + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i;\boldsymbol{\phi}_t) \qquad \forall i = 1,\ldots p_t \,:\, r_t^i \neq +\infty \\
& \theta \in \mathbb{R},\;\; \boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})
\end{array}
$$

if $r_t^i \neq -\infty$ for all $i = 1,\ldots,p_t$, and $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1};\boldsymbol{\phi}_{t-1}) = -\infty$ otherwise, where $r_t^i = \inf\{\boldsymbol{r}_t{}^\top \boldsymbol{y}_t^i \,:\, \boldsymbol{T}_t\,\hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t\,\boldsymbol{x}_t^i + \boldsymbol{V}_t\,\boldsymbol{y}_t^i \geqslant \boldsymbol{h}_t,\;\; \boldsymbol{y}_t^i \in \mathbb{R}^{n_t^2}\} \in \mathbb{R} \cup \{-\infty,+\infty\}$.

Several combinatorial optimization problems admit formulations where the constraints are not affected by uncertainty and that thus fall under the umbrella of Corollary 1. Examples include the shortest path, minimum cut, minimum spanning tree and matching problem with uncertain arc weights, the knapsack problem with uncertain utilities as well as the capacitated vehicle routing problem with uncertain travel costs.

Corollary 2 (State Variables Only)**.** *Fix $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ and $\boldsymbol{\phi}_{t-1} \in \Phi_{t-1}$. If the control variables $\boldsymbol{y}_t$ are absent in the stage-t nominal cost to-go problem Q_t, then the two-stage subproblem*

$\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1})$ *has the equivalent MILP reformulation*

$$
\begin{array}{lll}
\text{maximize} & \theta & \\
\text{subject to} & \theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t) + \mathrm{M} \cdot z_t^i & \forall i = 1, \ldots p_t \\
& \left[\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t^i\right]_\ell < [\boldsymbol{h}_t(\boldsymbol{\xi}_t)]_\ell + \mathrm{M} \cdot \lambda_t^{i\ell} & \forall i = 1, \ldots p_t,\ \forall \ell = 1, \ldots, m_t \\
& z_t^i + \displaystyle\sum_{\ell=1}^{m_t} \lambda_t^{i\ell} \leqslant m_t & \forall i = 1, \ldots p_t \\
& \theta \in \mathbb{R},\ \ \boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1}),\ \ \boldsymbol{z}_t \in \{0,1\}^{p_t},\ \ \boldsymbol{\lambda}_t \in \{0,1\}^{p_t \times m_t}, &
\end{array}
\tag{8}
$$

where M *is a sufficiently large positive number.*

For practical purposes, the strict inequalities in problem (8) can be relaxed to weak inequalities if we subtract sufficiently small positive quantities from their right-hand sides. An alternative derivation along the lines of Hanasusanto et al. (2015) is possible, but it would result in a formulation whose number of constraints scales exponentially in m_t.

Corollary 2 is applicable, for example, in purely discrete instances of the multi-stage robust mixed-integer problem (1) where $|\mathcal{X}_t|$ is not too large. This is the case for moderate-sized queueing network problems, where the state variables record the numbers of customers at each node, as well as for display ad allocation problems, where campaigns are matched to ad impressions so as to maximize the publisher's revenues subject to uncertain click-throughs (Chen et al., 2011).

Corollary 3 (Invertible Recourse Matrix)**.** *Fix* $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ *and* $\boldsymbol{\phi}_{t-1} \in \Phi_{t-1}$*. If the recourse matrix* $\boldsymbol{V}_t$ *in the stage-*t *nominal cost to-go problem* Q_t *is invertible, then the two-stage subproblem* $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1})$ *is infeasible if* $(\boldsymbol{V}_t^{-1})^\top \boldsymbol{r}_t \not\geqslant \boldsymbol{0}$*; otherwise, it has the equivalent MILP reformulation*

$$
\begin{array}{ll}
\text{maximize} & \theta \\
\text{subject to} & \theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \left[\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t^i\right]^\top (\boldsymbol{V}_t^{-1})^\top \boldsymbol{r}_t \\
& \qquad\qquad + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t) \qquad \forall i = 1, \ldots, p_t \\
& \theta \in \mathbb{R},\ \ \boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1}).
\end{array}
$$

Note that the stage-t worst-case cost to-go problem $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1})$ is infeasible precisely when for every $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$ the stage-t nominal cost to-go problem $Q_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t)$ is unbounded.

Corollary 4 (Decomposability)**.** *Fix* $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ *and* $\boldsymbol{\phi}_{t-1} \in \Phi_{t-1}$*. If the stage-*$t$ *nominal cost to-go problem* $Q_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t)$ *is feasible for some* $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$ *and the recourse matrix* $\boldsymbol{V}_t$ *in* Q_t *has* *a block-diagonal structure* $\boldsymbol{V}_t = \operatorname{diag}(\boldsymbol{V}_t^1, \ldots, \boldsymbol{V}_t^{k_t})$*, then the two-stage subproblem* $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1})$ *has the equivalent MILP reformulation*

$$
\begin{array}{lll}
\text{maximize} & \theta & \\
\text{subject to} & \theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \displaystyle\sum_{j=1}^{k_t} \tau_t^{ij} + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t) & \forall i = 1, \ldots, p_t \\
& \tau_t^{ij} \leqslant \left[\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t^i\right]_j^\top \boldsymbol{\lambda}_t^j + \mathrm{M} \cdot z_t^{ij}(\boldsymbol{\lambda}_t^j) + \mathrm{M} \cdot \displaystyle\sum_{\boldsymbol{\gamma}_t^j \in \Gamma_t^j} \delta_t^{ij}(\boldsymbol{\gamma}_t^j) & \\
& & \forall \boldsymbol{\lambda}_t^j \in \Lambda_t^j,\ \forall i = 1, \ldots, p_t,\ \forall j = 1, \ldots, k_t \\
& \displaystyle\sum_{\boldsymbol{\lambda}_t^j \in \Lambda_t^j} z_t^{ij}(\boldsymbol{\lambda}_t^j) = |\Lambda_t^j| - 1 & \forall i = 1, \ldots, p_t,\ \forall j = 1, \ldots, k_t \\
& \left[\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t^i\right]_j^\top \boldsymbol{\gamma}_t^j > \mathrm{M} \cdot (\delta_t^{ij}(\boldsymbol{\gamma}_t^j) - 1) & \\
& & \forall \boldsymbol{\gamma}_t^j \in \Gamma_t^j,\ \forall i = 1, \ldots, p_t,\ \forall j = 1, \ldots, k_t \\
& \theta \in \mathbb{R},\ \ \boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1}),\ \ \tau_t^{ij} \in \mathbb{R},\ \ z_t^{ij} : \Lambda_t^j \to \{0,1\},\ \ \delta_t^{ij} : \Gamma_t^j \to \{0,1\}, & \\
& & i = 1, \ldots, p_t \text{ and } j = 1, \ldots, k_t,
\end{array}
$$

where $\boldsymbol{r}_t = (\boldsymbol{r}_t^1, \ldots, \boldsymbol{r}_t^{k_t})$ *such that* $\boldsymbol{V}_t^j$ *and* $\boldsymbol{r}_t^j$ *have matching numbers of rows,* $j = 1, \ldots, k_t$*,* Λ_t^j *and* Γ_t^j *contain the extreme points and extreme rays of the polyhedron* $\{\boldsymbol{\lambda} \geqslant \boldsymbol{0} \,:\, [\boldsymbol{V}_t^j]^\top \boldsymbol{\lambda} = \boldsymbol{r}_t^j\}$*, respectively, and* $[\cdot]_j$ *refers to the subsets of rows matching those of* $\boldsymbol{V}_t^j$ *and* $\boldsymbol{r}_t^j$*.*

The equivalence stated in Corollary 4 extends to unbounded instances of the two-stage subproblem $\mathcal{Q}_t$ in the sense that $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1})$ is unbounded if and only if the MILP in Corollary 4 attains an optimal value greater than or equal to M. Corollary 4 applies to multi-stage robust mixed-integer problems that, as far as the continuous decisions $\boldsymbol{y}_t$ are concerned, decompose into separate components. Note that the discrete decisions $\boldsymbol{x}_t$ as well as the uncertain parameters $\boldsymbol{\xi}_t$ may still connect the different components. The study of optimization problems with decomposable structure goes back to Bellman (1957), and much work has since been dedicated to developing efficient solution techniques (most notably the Dantzig-Wolfe and Benders' decomposition) and applying them, among others, to problems in supply chain management, production scheduling as well as the design and operation of energy systems (Martin, 1999). In supply chain management problems, for example, the flows of different products are often restricted by constraints that exhibit a block-diagonal structure. Likewise, in production scheduling problems, the production activities typically decompose across different facilities.

Remark 3. *If the set $\Phi_t(\boldsymbol{\phi}_{t-1})$ of inter-stage uncertainties $\boldsymbol{\phi}_t$ is not too large, then we can alternatively solve the MILPs in Corollaries 1–4 by enumerating all possible values of $\boldsymbol{\phi}_t \in \Phi(\boldsymbol{\phi}_{t-1})$, treating $\boldsymbol{\phi}_t$ as a constant in the MILPs and solving the resulting problems in parallel.*

Remark 4. *Even if a two-stage subproblem $\mathcal{Q}_t$ is not amenable to Corollaries 3 or 4* per se*, it may admit an equivalent reformulation that is once a candidate decision $\hat{\boldsymbol{x}}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ has been fixed and redundant constraints in the stage-t nominal cost to-go problem Q_t have been removed.*

4 Conservative Approximation of the Two-Stage Subproblems $\mathcal{Q}_t$

In cases where the two-stage subproblems $\mathcal{Q}_t$ cannot be solved exactly, we propose to use conservative approximations $\overline{\mathcal{Q}}_t$ that restrict the adaptivity of the state or control variables. The following observation justifies the use of stage-wise conservative approximations $\overline{\mathcal{Q}}_t$ in a multi-stage setting.

Observation 1 (Conservative Approximations: Propagation of Upper Bounds)**.**

(i) *If $\mathcal{Q}_{t+1}$ in each nominal stage-t cost to-go problem Q_t is replaced with an upper bound $\overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \geqslant \mathcal{Q}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t)$, $t = 1, \ldots, T$, $\boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}}$ and $\boldsymbol{\phi}_t \in \Phi_t$, then the resulting approximations $\overline{\mathcal{Q}}_t$ of $\mathcal{Q}_t$ are conservative: $\overline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}; \boldsymbol{\phi}_{t-1}) \geqslant \mathcal{Q}_t(\boldsymbol{x}_{t-1}; \boldsymbol{\phi}_{t-1})$ for all t, $\boldsymbol{x}_{t-1}$ and $\boldsymbol{\phi}_{t-1}$.*

(ii) *If $\overline{\mathcal{Q}}_1 < \infty$, then the policy $\{(\boldsymbol{x}_t^\star, \boldsymbol{y}_t^\star)\}_{t=1}^T$ that is implicitly defined by Algorithm 2 using the upper bounds $\overline{\mathcal{Q}}_t$, $t = 1, \ldots, T$, is feasible in* (3) *and attains worst-case costs of at most $\overline{\mathcal{Q}}_1$.*

Observation 1 shows that $\overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t^\star; \boldsymbol{\phi}_t)$ remains an upper bound on the worst-case cost to-go when decision $\boldsymbol{x}_t^\star \in \mathcal{X}_t^{\mathrm{C}}$ is taken by Algorithm 2, even when the worst-case cost to-go functions in later time stages are also replaced by conservative approximations.

In the remainder of this section, we discuss two different conservative bounds based on approximations of the state and control variables. Since both approximations distinguish between the inter-stage uncertainties $\boldsymbol{\phi}_t$ and the intra-stage uncertainties $\boldsymbol{\psi}_t$ in time stage t, we define $\Psi_t(\boldsymbol{\phi}_t) = \{\boldsymbol{\psi}_t : (\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \in \Xi_t(\boldsymbol{\phi}_{t-1})\}$ as the set of conditional intra-stage uncertainty realizations.

Approximation 1 (State Approximation)**.** *The state variables $\boldsymbol{x}_t$ in each two-stage subproblem $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_t)$ no longer adapt to the exact values of the intra-stage uncertainties $\boldsymbol{\psi}_t$, but only to their set memberships $\{\mathbf{1}[\boldsymbol{\psi}_t \in \Psi_{tj}(\boldsymbol{\phi}_t)]\}_{j=1}^l$ across a pre-selected polyhedral partition $\{\Psi_{tj}(\boldsymbol{\phi}_t)\}_{j=1}^l$ of the conditional intra-stage uncertainty realizations $\Psi_t(\boldsymbol{\phi}_t)$.*

Note that under Approximation 1, the state variables $\boldsymbol{x}_t$ remain fully adaptive in the inter-stage uncertainties $\boldsymbol{\phi}_t$. Approximation 1 then adopts a piecewise constant decision rule approximation in $\boldsymbol{\psi}_t$ for the state variables $\boldsymbol{x}_t$. The number l of subsets in the partition $\{\Psi_{tj}(\boldsymbol{\phi}_t)\}_{j=1}^{l}$ of $\Psi_t(\boldsymbol{\phi}_t)$ can vary with the time stage t as well as the values of $\boldsymbol{\phi}_{t-1}$, $\boldsymbol{\phi}_t$ and $\hat{\boldsymbol{x}}_{t-1}$; for ease of exposition, however, we notationally suppress this potential dependence.

Observation 2. *Denote by $\eta^\star(\boldsymbol{\phi}_t, j, \boldsymbol{x}_t)$, $\boldsymbol{\phi}_t \in \Phi_t(\boldsymbol{\phi}_{t-1})$, $j = 1, \ldots, l$ and $\boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}}$, the optimal value of the following worst-case optimization problem with continuous recourse.*

$$
\begin{array}{ll}
\text{maximize} & \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)^\top \boldsymbol{x}_t + {\boldsymbol{r}_t}^\top \boldsymbol{y}_t + \overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \\ & \boldsymbol{y}_t \in \mathbb{R}^{n_t^2} \end{array} \right] \\
\text{subject to} & \boldsymbol{\psi}_t \in \Psi_{tj}(\boldsymbol{\phi}_t)
\end{array}
$$

Under Approximation 1, the optimal value of the two-stage subproblem $\overline{\mathcal{Q}}_t$ coincides with

$$
\max \Big\{ \min \left\{ \eta^\star(\boldsymbol{\phi}_t, j, \boldsymbol{x}_t) \,:\, \boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}} \right\} \,:\, \boldsymbol{\phi}_t \in \Phi_t(\boldsymbol{\phi}_{t-1}) \text{ and } j = 1, \ldots, l \Big\}.
$$

The $|\Phi_t(\boldsymbol{\phi}_{t-1})| \cdot l \cdot p_t$ worst-case optimization problems of Observation 2 can be solved in parallel. Together with our earlier observation of the parallel solution of the subproblems in each time stage, there are thus two levels of parallelization: We can solve $|\Phi_{t-1}| \cdot |\Phi_t| \cdot l \cdot p_{t-1} \cdot p_t$ worst-case optimization problems in parallel in each time stage t. In contrast to the two-stage subproblems $\mathcal{Q}_t$, the worst-case optimization problems in Observation 2 contain no first-stage decisions, discrete uncertainties or discrete recourse decisions, and they can be solved using column-and-constraint generation (Zeng and Zhao, 2013) or iterative liftings of the uncertainty set (Georghiou et al., 2020).

Instead of approximating the state variables, we can also approximate the control variables.

Approximation 2 (Control Approximation)**.** *The control variables $\boldsymbol{y}_t$ in each two-stage subproblem $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_t)$ no longer adapt to the exact values of the intra-stage uncertainties $\boldsymbol{\psi}_t$; instead, they are optimally chosen from an optimally selected set of affine decision rules $\{\boldsymbol{y}_{tj} : \Psi_t(\boldsymbol{\phi}_t) \to \mathbb{R}^{n_t^2}\}_{j=1}^{l}$.*

Under Approximation 2, the control variables $\boldsymbol{y}_t$ remain fully adaptive in the inter-stage uncertainties $\boldsymbol{\phi}_t$. Instead of selecting $\boldsymbol{y}_t$ optimally in response to the observation of $\boldsymbol{\psi}_t$, Approximation 2 selects the l best affine decision rules $\{\boldsymbol{y}_{tj} : \Psi_t(\boldsymbol{\phi}_t) \to \mathbb{R}^{n_t^2}\}_{j=1}^{l}$ here-and-now (that is, upon observation of $\boldsymbol{\phi}_t$ but before observing $\boldsymbol{\psi}_t$) and subsequently computes $\boldsymbol{y}_t$ from the best of these l decision rules wait-and-see (that is, upon observation of $\boldsymbol{\psi}_t$). Similar to the number l in Approximation 1, the number l of decision rules in Approximation 2 can vary with the values of t, $\boldsymbol{\phi}_{t-1}$, $\boldsymbol{\phi}_t$ and $\hat{\boldsymbol{x}}_{t-1}$.

Observation 3. *Denote by $\eta^\star(\boldsymbol{\phi}_t)$, $\boldsymbol{\phi}_t \in \Phi_t(\boldsymbol{\phi}_{t-1})$, the optimal value of the following (p_t, l)-adaptable two-stage robust MILP that optimizes over affine decision rules $\{\boldsymbol{y}_{tj}\}_{j=1}^l$ here-and-now and over combinations $(k, j) \in \{1, \ldots, p_t\} \times \{1, \ldots, l\}$ wait-and-see.*

$$\begin{array}{ll} \text{minimize} & \left(\max\limits_{\boldsymbol{\psi}_t \in \Psi_t(\boldsymbol{\phi}_t)} \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)^\top \boldsymbol{x}_t^k + \boldsymbol{r}_t{}^\top \boldsymbol{y}_{tj}(\boldsymbol{\psi}_t) + \overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t^k; \boldsymbol{\phi}_t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \boldsymbol{x}_t^k + \boldsymbol{V}_t\, \boldsymbol{y}_{tj}(\boldsymbol{\psi}_t) \geqslant \boldsymbol{h}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \\ & (k, j) \in \{1, \ldots, p_t\} \times \{1, \ldots, l\} \end{array} \right] \right) \\ \text{subject to} & \boldsymbol{y}_{tj} : \Xi_t(\boldsymbol{\phi}_{t-1}) \rightarrow \mathbb{R}^{n_t^2},\ j = 1, \ldots, l. \end{array}$$

Under Approximation 2, the optimal value of the two-stage subproblem $\overline{\mathcal{Q}}_t$ coincides with

$$\max\{\eta^\star(\boldsymbol{\phi}_t) \,:\, \boldsymbol{\phi}_t \in \Phi_t(\boldsymbol{\phi}_{t-1})\}.$$

The $|\Phi_t(\boldsymbol{\phi}_{t-1})|$ (p_t, l)-adaptable two-stage robust MILPs of Observation 3 can be solved in parallel. Together with our earlier observation of the parallel solution of the subproblems in each time stage, we can thus solve $|\Phi_{t-1}| \cdot |\Phi_t| \cdot p_{t-1}$ (p_t, l)-adaptable two-stage robust optimization problems in parallel in each time stage t. In contrast to the two-stage subproblems $\mathcal{Q}_t$, the (p_t, l)-adaptable two-stage robust MILPs in Observation 3 contain neither discrete uncertainties nor continuous recourse decisions, and they are thus amenable to standard solution schemes for K-adaptable two-stage robust optimization problems (Hanasusanto et al., 2015; Subramanyam et al., 2020).

5 Lower Bounds and Selection of Candidate State Decision Sets

So far we assumed that the sets $\mathcal{X}_t^{\mathrm{C}}$ of candidate state decisions in the finite adaptability approximation (3) are pre-specified and fixed. This is reasonable if domain expertise can be leveraged to design promising candidate state decisions, but it may place an undue burden on the decision maker when such information is not available. To address such situations, we develop in this section a greedy heuristic that iteratively expands the sets $\mathcal{X}_t^{\mathrm{C}}$ based on optimal solutions to lower bounds of the multi-stage robust mixed-integer optimization problem (1). The lower bounds are useful in their own right to estimate the suboptimality of a solution to (3), regardless of whether the sets $\mathcal{X}_t^{\mathrm{C}}$ are being expanded or kept fixed.

Our expansion heuristic is motivated by a necessary and a sufficient condition for the improvement of the optimal value of the finite adaptability approximation (3). To ease the exposition, we denote by $\mathcal{X}^{\mathrm{C}} = \{\mathcal{X}_1^{\mathrm{C}}, \ldots, \mathcal{X}_T^{\mathrm{C}}\}$ the collection of all candidate state decision sets, and for a fixed choice of $\mathcal{X}^{\mathrm{C}}$ and Ξ we let $P(\mathcal{X}^{\mathrm{C}}, \Xi)$ denote both the formulation (3) and its optimal value.

Proposition 4 (Necessary and Sufficient Conditions for Improvement)**.** *Consider two collections of candidate state decision sets* $\mathcal{X}^{\mathrm{C}} = \{\mathcal{X}_t^{\mathrm{C}}\}_{t=1}^T$ *and* $\hat{\mathcal{X}}^{\mathrm{C}} = \{\hat{\mathcal{X}}_t^{\mathrm{C}}\}_{t=1}^T$ *with* $\hat{\mathcal{X}}_\tau^{\mathrm{C}} = \mathcal{X}_\tau^{\mathrm{C}} \cup \{\boldsymbol{x}'_\tau\}$ *for some* $\tau = 1, \ldots, T$ *and* $\boldsymbol{x}'_\tau \in \mathcal{X}_t$*, and* $\hat{\mathcal{X}}_t^{\mathrm{C}} = \mathcal{X}_t^{\mathrm{C}}$ *for all* $t \neq \tau$*. A* necessary condition *for* $P(\hat{\mathcal{X}}^{\mathrm{C}}, \Xi) < P(\mathcal{X}^{\mathrm{C}}, \Xi)$ *is the existence of* $\boldsymbol{\phi}_{\tau-1} \in \Phi_{\tau-1}$*,* $\hat{\boldsymbol{x}}_{\tau-1} \in \mathcal{X}_{\tau-1}^{\mathrm{C}}$*,* $\boldsymbol{\xi}_\tau \in \Xi_\tau(\boldsymbol{\phi}_{\tau-1})$ *and* $\boldsymbol{y}'_\tau \in \mathbb{R}^{n_\tau^2}$ *such that*

(1) $\boldsymbol{T}_\tau(\boldsymbol{\xi}_\tau)\, \hat{\boldsymbol{x}}_{\tau-1} + \boldsymbol{W}_\tau(\boldsymbol{\xi}_\tau)\, \boldsymbol{x}'_\tau + \boldsymbol{V}_\tau\, \boldsymbol{y}'_\tau \geqslant \boldsymbol{h}_\tau(\boldsymbol{\xi}_\tau)$*, and*

(2) $\boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}'_\tau + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}'_\tau + \mathcal{Q}_{\tau+1}(\boldsymbol{x}'_\tau; \boldsymbol{\phi}_\tau) < Q_\tau(\hat{\boldsymbol{x}}_{\tau-1}; \boldsymbol{\xi}_\tau)$.

A sufficient condition *for* $P(\hat{\mathcal{X}}^{\mathrm{C}}, \Xi) < P(\mathcal{X}^{\mathrm{C}}, \Xi)$ *is the existence of an optimal solution* $(\boldsymbol{x}^\star, \boldsymbol{y}^\star)(\boldsymbol{\xi})$ *to* $P(\mathcal{X}^{\mathrm{C}}, \Xi)$ *such that for all corresponding worst-case uncertainties* $\boldsymbol{\xi}^\star \in \Xi$*, conditions (1) and (2) hold for* $\hat{\boldsymbol{x}}_{\tau-1} = \boldsymbol{x}^\star_{\tau-1}([\boldsymbol{\xi}^\star]^{\tau-1})$ *and* $\boldsymbol{\xi}_\tau = \boldsymbol{\xi}^\star_\tau$*.*

In Proposition 4, $\boldsymbol{\xi}^\star \in \Xi$ is a worst-case uncertainty if $\sum_{t=1}^T \boldsymbol{q}_t(\boldsymbol{\xi}_t^\star)^\top \boldsymbol{x}_t^\star([\boldsymbol{\xi}^\star]^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t^\star([\boldsymbol{\xi}^\star]^t) = P(\mathcal{X}^{\mathrm{C}}, \Xi)$. The conditions (1) and (2) correspond to the feasibility and objective improvement of $(\hat{\boldsymbol{x}}_t, \hat{\boldsymbol{y}}_t)$ in the stage-t nominal cost to-go problem Q_t, respectively. The necessary condition in Proposition 4 is not sufficient since the candidate state solution $\boldsymbol{x}'_\tau$ may only improve nominal stage-t problems $Q_t(\boldsymbol{x}_{\tau-1}; \boldsymbol{\xi}_\tau)$ that do not involve worst-stage uncertainties $\boldsymbol{\xi}_\tau$ or optimal decisions $\boldsymbol{x}_{\tau-1}$ in $P(\mathcal{X}^{\mathrm{C}}, \Xi)$. Likewise, the sufficient condition in Proposition 4 is not necessary since it is possible that by adding $\boldsymbol{x}'_\tau$ to $\mathcal{X}_\tau^{\mathrm{C}}$ the worst-case cost to-go $\mathcal{Q}_t$ and nominal cost to-go Q_t decrease in early stages $t = 1, \ldots, \tau - 1$ for previously sub-optimal decisions, thus leading to an improvement without the sufficient condition being satisfied.

Motivated by Proposition 4, we limit our attention to worst-case uncertainties $\boldsymbol{\xi}^\star$ when selecting new candidate state decisions to add to $\mathcal{X}^{\mathrm{C}}$. This idea is formalized in the following algorithm.

Algorithm 3: Iterative expansion of the candidate state decision sets.

1. **Initialization.** Initialize the collection of candidate state decision sets $\mathcal{X}^{\mathrm{C}}$.

2. **Conservative Bound.** Solve the current finite adaptability problem $P(\mathcal{X}^{\mathrm{C}}, \Xi)$.

3. **Progressive Bound.** Identify a set $\hat{\Xi} \subseteq \Xi$ of worst-case uncertainties and solve $P(\mathcal{X}, \hat{\Xi})$, where $\mathcal{X} = \{\mathcal{X}_t\}_{t=1}^T$. Identify the candidate state decision(s) to add to $\mathcal{X}^{\mathrm{C}}$.

4. **Termination.** Stop if $\mathcal{X}^{\mathrm{C}}$ has not been updated. Otherwise, go back to Step 2.

Algorithm 3 produces a sequence of upper and lower bounds to the multi-stage robust mixed-integer optimization problem (1) by solving finite adaptability problems $P(\mathcal{X}^{\mathrm{C}}, \Xi)$ with updated candidate state decision sets $\mathcal{X}^{\mathrm{C}}$ and relaxations of (1) that only involve subsets $\hat{\Xi} \subseteq \Xi$ of the uncertain parameter realizations, respectively. Algorithm 3 constitutes a template that allows for several ways of *(i)* initializing and updating the collection $\mathcal{X}^{\mathrm{C}}$, *(ii)* constructing reduced uncertainty sets $\hat{\Xi}$ and *(iii)* solving the upper and lower bounds $P(\mathcal{X}^{\mathrm{C}}, \Xi)$ and $P(\mathcal{X}, \hat{\Xi})$.

We initialize the collection $\mathcal{X}^{\mathrm{C}}$ of candidate state decision sets by solving a static version of the multi-stage robust mixed-integer optimization problem (1) and identifying $\mathcal{X}_t^{\mathrm{C}}$ with the decision $\boldsymbol{x}_t^\star \in \mathcal{X}_t$ taken by the optimal solution to the static problem. Likewise, we update $\mathcal{X}^{\mathrm{C}}$ by adding to each set $\mathcal{X}_t^{\mathrm{C}}$ the decision $\boldsymbol{x}_t^\star \in \mathcal{X}_t$ taken by an optimal solution to the lower bound $P(\mathcal{X}, \hat{\Xi})$. This approach is heuristic as it uses the sufficient (but not necessary) condition of Proposition 4.

We consider two alternative constructions of the reduced uncertainty set $\hat{\Xi}$. In the first, we identify $\hat{\Xi}$ with all worst-case uncertainties $\boldsymbol{\xi}^\star$ determined during the solution of the upper bound problem $P(\mathcal{X}^{\mathrm{C}}, \Xi)$, which results in a set $\hat{\Xi}$ of finite cardinality. In the second approach, we set $\hat{\Xi} = \{\boldsymbol{\xi} = (\boldsymbol{\phi}^\star, \boldsymbol{\psi}) \in \Xi \,:\, \boldsymbol{\phi}^\star \in \Phi^\star\}$, where $\Phi^\star$ denotes the set of worst-case inter-stage uncertainties in problem $P(\mathcal{X}^{\mathrm{C}}, \Xi)$. In other words, the second approach considers all uncertainty realizations that emerge from combinations of the worst-case inter-stage uncertainties with any admissible intra-stage uncertainties. If necessary, the reduced uncertainty set $\hat{\Xi}$ resulting from either approach can be restricted further by selecting subsets of the worst-case parameter realizations.

We solve the upper bound problems $P(\mathcal{X}^{\mathrm{C}}, \Xi)$ using Algorithm 1. As part of this algorithm, we solve the emerging two-stage subproblems $\mathcal{Q}_t$ either exactly (using the techniques of Section 3) or approximately (using the methods of Section 4). The lower bound problems $P(\mathcal{X}, \hat{\Xi})$, on the other hand, can be solved as large-scale MILPs as long as the reduced uncertainty set $\hat{\Xi}$ is finite (*i.e.*, it takes the form of a scenario tree or a scenario fan). If $\hat{\Xi}$ has infinite cardinality, on the other hand, then we solve a single-stage robust optimization problem in which the decisions $(\boldsymbol{x}_t, \boldsymbol{y}_t)$ only adapt to the inter-stage uncertainties $\boldsymbol{\phi}_t$ but not to the intra-stage uncertainties $\boldsymbol{\psi}_t$. Note that this approach no longer provides a lower bound on the multi-stage robust mixed-integer optimization problem (1), but it remains a valid heuristic to choose candidate state decisions to add to $\mathcal{X}^{\mathrm{C}}$.

6 Numerical Experiments

We compare our iterative solution scheme for the finite adaptability approximation (3) with different benchmark approaches for solving the multi-stage robust mixed-integer problem (1) on two case studies. The first case study concerns a route planning problem in which the stage-wise feasible regions $\mathcal{X}_t$ are sufficiently benign for our approach to solve the problem exactly within seconds. The second case study concerns a transportation-location planning problem where the stage-wise feasible regions $\mathcal{X}_t$ grow exponentially in the problem description; here, we use our iterative state variable selection procedure from Section 5 in combination with the piecewise constant state decision approximation from Section 4 to approximately solve the problem to a high accuracy.

6.1 Route Planning

We consider a dynamic robust route planning problem on a directed graph $G = (N, A)$ with nodes $N = \{1, \ldots, n\}$ and arcs $A \subseteq V \times V$. The goal is to determine an adaptive shortest path from the start node 1 to the terminal node n that minimizes the worst-case route length when the arc lengths are uncertain and non-stationary over time. The problem can be formulated as the following instance of our multi-stage robust mixed-integer optimization problem (1):

$$
\begin{array}{lll}
\text{minimize} & \displaystyle\max_{\boldsymbol{\xi} \in \Xi} \sum_{t=1}^{T} \sum_{(i,j) \in A} y_{tij}(\boldsymbol{\xi}^t) & \\
\text{subject to} & y_{tij}(\boldsymbol{\xi}^t) \geqslant r_{tij}(\boldsymbol{\xi}_t)\Big(x_{tj}(\boldsymbol{\xi}^t) + x_{t-1,i}(\boldsymbol{\xi}^{t-1}) - 1\Big) & \forall \boldsymbol{\xi} \in \Xi,\ \forall (i,j) \in A,\ \forall t = 1, \ldots, T \\
& x_{tj}(\boldsymbol{\xi}^t) + x_{t-1,i}(\boldsymbol{\xi}^{t-1}) \leqslant 1 & \forall \boldsymbol{\xi} \in \Xi,\ \forall (i,j) \notin A,\ \forall t = 1, \ldots, T \\
& \boldsymbol{x}_T(\boldsymbol{\xi}^T) = \mathbf{e}_n & \forall \boldsymbol{\xi} \in \Xi \\
& \boldsymbol{x}_t(\boldsymbol{\xi}^t) \in \{\mathbf{e}_i \,:\, i = 1, \ldots, n\},\ \ \boldsymbol{y}_t(\boldsymbol{\xi}^t) \geqslant \mathbf{0} & \forall \boldsymbol{\xi} \in \Xi,\ \forall t = 1, \ldots, T.
\end{array}
$$

Here, $r_{tij}(\boldsymbol{\xi}_t)$ represents the uncertain length of arc $(i,j) \in A$ in time stage $t = 1, \ldots, T$. The state variable $\boldsymbol{x}_t$ records the node entered at time stage t, with $\boldsymbol{x}_0(\boldsymbol{\xi}^0) \equiv \mathbf{e}_1$, and the components of the auxiliary control variable $\boldsymbol{y}_t$ evaluate (at optimality) to $y_{tij}(\boldsymbol{\xi}^t) = r_{tij}(\boldsymbol{\xi}_t)$ if arc (i,j) is traversed in time stage t and $y_{tij}(\boldsymbol{\xi}^t) = 0$ otherwise. The objective function evaluates the worst-case cumulative length of the arcs traversed over time. The first constraint ensures that $y_{tij}(\boldsymbol{\xi}^t) \geqslant r_{tij}(\boldsymbol{\xi}_t)$ precisely when we traverse from node i to node j in time stage t, the second constraint ensures that this is possible only when $(i, j) \in A$, and the third constraint ensures that we reach node n by time T.

For our numerical experiments, we generate random graphs with nodes $N = \{1, \ldots, 100\}$ that are located uniformly at random on the 'N'-shaped subset $[0, 10]^2 \backslash ([2, 4] \times [0, 8] \cup [6, 8] \times [2, 10])$ of the two-dimensional square $[0, 10]^2$; the exceptions are the start node 1 and the terminal node $n = 100$, which are located at the bottom-left corner $(0, 0)$ and top-right corner $(10, 10)$, respectively. To construct the arc set, we start with $A = N \times N$ (including all self-loops) and remove all arcs that cross the boundaries of our 'N'-shaped set. We subsequently remove arcs in order of decreasing arc lengths until either $|A| = 6n$ or the removal of any further arcs would make the graph disconnected. This eliminates trivial problem instances in which the shortest path contains few arcs. Note that the inclusion of self-loops allows the decision maker to reside at the current node in any time stage; in particular, the decision maker can arrive and reside at the terminal node n prior to time stage T. To model the uncertain arc lengths $r_{tij}(\boldsymbol{\xi}_t)$, we construct a budget uncertainty set of the form

$$\Xi_t(\phi_{t-1}) := \left\{ \boldsymbol{\xi}_t = (\phi_t, \boldsymbol{\psi}_t) : \phi_t \in \Phi_t(\phi_{t-1}),\ \boldsymbol{\psi}_t \in [0, 1]^{|A|},\ \mathbf{e}^\top \boldsymbol{\psi}_t \leqslant \phi_{t-1} - \phi_t,\ \psi_{tii} = 0 \ \ \forall i \in N \right\}.$$

Here, the intra-stage uncertainties $\boldsymbol{\psi}_t$ determine the arc length excesses via $r_{tij}(\boldsymbol{\xi}_t) := (1 + \psi_{tij}) r^0_{ij}$, where the nominal length r^0_{ij} of arc $(i, j) \in A$ is set to the Euclidean distance between nodes i and j, and the inter-stage uncertainty budget ϕ_t evolves according to the set $\Phi_t(\phi_{t-1}) := \{\phi_t \in \mathbb{N}_0 : \phi_t \leqslant \phi_{t-1}\}$ for some initial budget ϕ_0 (selected below). We set the time horizon to $T = 40$.

The static problem, where the uncertain arc lengths are stationary and observed after choosing the shortest path, can be reduced to the solution of multiple deterministic shortest path problems (Bertsimas and Sim, 2003). In our finite adaptability approximation of the dynamic problem, we employ full adaptivity, that is, we set $\mathcal{X}^{\mathrm{C}}_t = \mathcal{X}_t$ for all t; this is possible since $|\mathcal{X}_t| = n$ in this case study. We solve the resulting problem with Algorithm 1 (*cf.* Section 2), where the subproblems $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1})$ are solved exactly using Corollary 3 and Remarks 3 and 4 from Section 3. In fact, each subproblem $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1})$ of our finite adaptability approximation can be solved to any fixed precision ϵ in time $\mathcal{O}(\log \epsilon^{-1})$ via a binary search. All instances of the static as well as the adaptive route planning formulation were solved within 10 seconds each, which is why we omit detailed runtime comparisons in this section.

Figure 2 compares the average worst-case route lengths over 50 randomly generated instances and initial uncertainty budgets $\phi_0 \in \{0, \ldots, T\}$ of three alternative methods: *(i)* the static robust

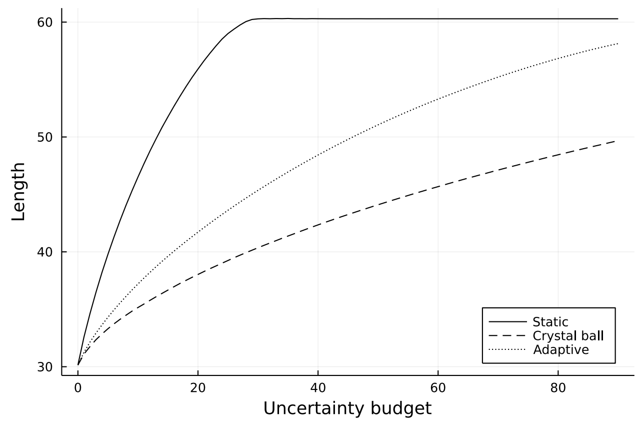

Figure 2. Worst-case route lengths of the static, the crystal ball and our dynamic solution across 50 randomly generated instances and varying uncertainty budgets ϕ_0.

route ('static') that solves a single-stage robust optimization problem to determine a fixed here-and-now route to minimize the worst case route length before any arc lengths are observed; *(ii)* the worst-case crystal ball route ('crystal ball') over all possible uncertainty realizations that benefits from observing all arc lengths ahead of time; and *(iii)* our adaptive robust route ('adaptive'). As expected, the figure confirms that the worst-case route length of our adaptive solution is bounded from above and below by the static and the crystal ball solution, respectively. Moreover, we can see that the static solution 'saturates' quickly even for relatively small uncertainty budgets, in which case the adversarial nature maximally perturbs all arcs lengths along the worst-case optimal path. This illustrates the value of adaptivity in the context of robust route planning.

To further study the differences between static and adaptive routes, we run a simulation study where in each time stage t, the adversarial nature randomly selects each arc $(i, j) \in A$ emanating from the decision maker's present location $x_{t-1,i}(\boldsymbol{\xi}^{t-1}) = 1$ with probability p, and the lengths of the selected arc are maximally increased (subject to the remaining uncertainty budget). Figure 3 compares the static robust route (in blue) with different adaptive robust routes (in red) that respond in a worst-case optimal way to randomly drawn arc length sequences $\boldsymbol{r}_1(\boldsymbol{\xi}_1), \ldots, \boldsymbol{r}_T(\boldsymbol{\xi}_T)$ for the initial uncertainty budgets $\phi_0 \in \{0, 15, 30, 45\}$. In the figure, the optimal static route changes as the initial budget increases from 0 to 15, but it subsequently returns to its original form when the budget increases further. Indeed, an initial budget of 30 is sufficient to maximally perturb all arcs on the worst-case optimal path, in which case the static robust problem becomes equivalent to

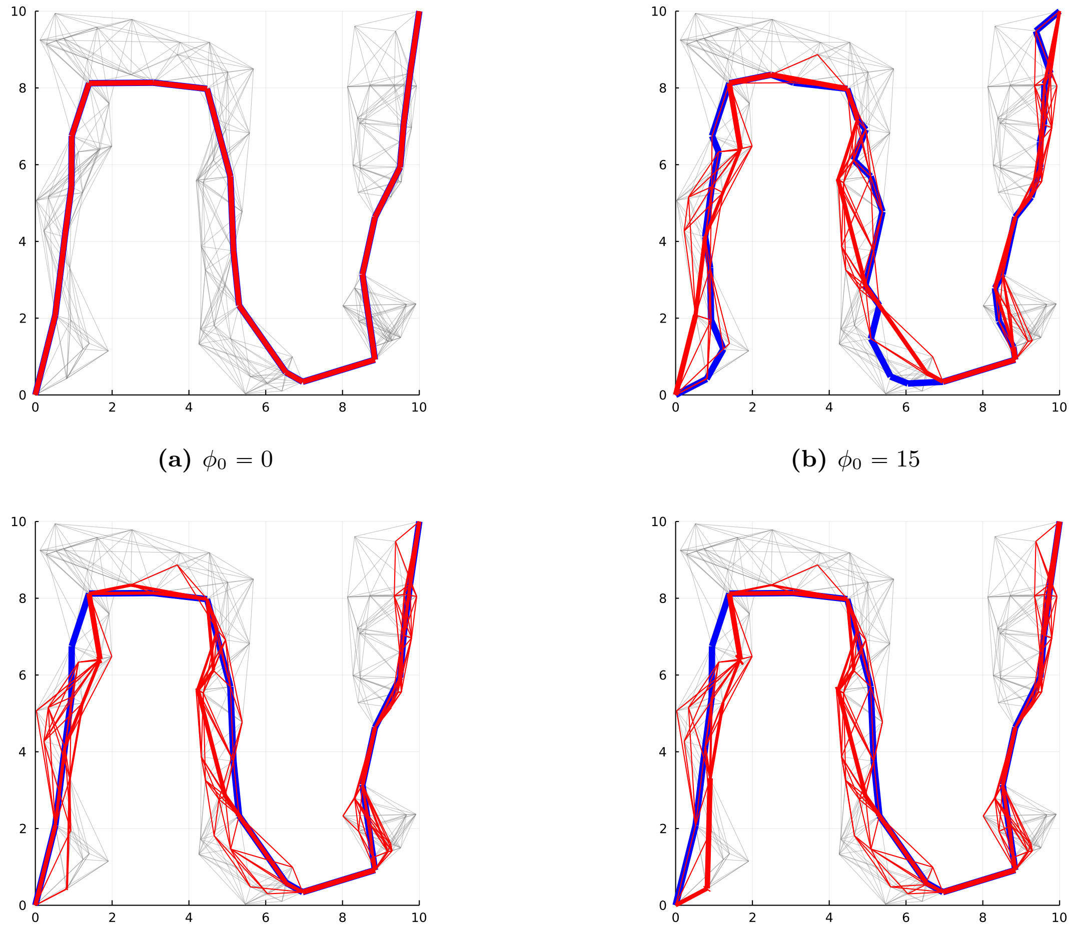

(a) $\phi_0 = 0$

(b) $\phi_0 = 15$

(c) $\phi_0 = 30$

(d) $\phi_0 = 45$

Figure 3. Comparison of the optimal static route (blue) with 1,000 simulated paths of the optimal adaptive route (red; thickness corresponds to traversal frequency) on a randomly generated instance with 4 different initial uncertainty budgets.

a nominal shortest path problem. For an initial uncertainty budget of 0, the simulated paths of our adaptive route coincide with the optimal static route. For all other uncertainty budgets, however, our adaptive route reacts to the revealed arc lengths; this is illustrated by the red arcs, whose thickness is proportional to the traversal frequency across 1,000 simulated arc length sequences. It is worth noting that the initial arc selected by our adaptive route often differs from the one chosen by the optimal static route. This shows that future adaptivity needs to be accounted for by the decision taken here-and-now, and it demonstrates that a receding horizon implementation of a static problem may result in suboptimal solutions.

6.2 Location-Transportation Planning

We study a dynamic location-transportation problem where facilities can be built at locations $i = 1, \ldots, m$ to serve uncertain customer demands at locations $j = 1, \ldots, n$ over a time span of T stages. The problem can be formulated as the following instance of our multi-stage robust mixed-integer optimization problem (1):

$$
\begin{array}{lll}
\text{minimize} & \displaystyle\max_{\boldsymbol{\xi}\in\Xi} \sum_{t=1}^{T} \boldsymbol{p}_t{}^\top \boldsymbol{y}_t(\boldsymbol{\xi}^t) - \boldsymbol{c}_t{}^\top \boldsymbol{z}_t(\boldsymbol{\xi}^t) - \boldsymbol{b}_t{}^\top \boldsymbol{x}_t(\boldsymbol{\xi}^t) & \\
\text{subject to} & \displaystyle\sum_{i=1}^{m} y_{tij}(\boldsymbol{\xi}^t) + z_{tj}(\boldsymbol{\xi}^t) = d_{tj}(\boldsymbol{\xi}_t) & \forall \boldsymbol{\xi} \in \Xi,\ \forall j = 1, \ldots, n,\ \forall t = 1, \ldots, T \\
& \displaystyle\sum_{j=1}^{n} y_{tij}(\boldsymbol{\xi}^t) \leqslant C_i \cdot x_{ti}(\boldsymbol{\xi}^t) & \forall \boldsymbol{\xi} \in \Xi,\ \forall i = 1, \ldots, m,\ \forall t = 1, \ldots, T \\
& \boldsymbol{x}_t(\boldsymbol{\xi}^t) \geqslant \boldsymbol{x}_{t-1}(\boldsymbol{\xi}^{t-1}) & \forall \boldsymbol{\xi} \in \Xi,\ \forall t = 1, \ldots, T \\
& \boldsymbol{x}_t(\boldsymbol{\xi}^t) \in \{0,1\}^m,\ \boldsymbol{y}_t(\boldsymbol{\xi}^t) \in \mathbb{R}^{mn}_+,\ \boldsymbol{z}_t(\boldsymbol{\xi}^t) \in \mathbb{R}^n_+ & \forall \boldsymbol{\xi} \in \Xi,\ \forall t = 1, \ldots, T.
\end{array}
$$

Here, p_{tij} and c_{tj} denote the profit and penalty for serving or disregarding one unit of demand at location j from location i in time stage t, respectively, b_{ti} represents the cost of operating a facility at location i in time stage t, d_{tj} are the uncertain customer demands at location j in time stage t, and C_i is the capacity of the candidate facility at location i. The state variables x_{ti} record whether a facility exists at location i in time stage t, whereas the control variables y_{tij} and z_{tj} record the demand at location j that is served from location i or disregarded in time stage t, respectively. The objective function computes the cumulative worst-case difference of sales profits and penalty as well as facility operating costs. The first constraint stipulates that all customer demands are either served or disregarded, the second constraint ensures that the facility capacities are obeyed, and the third constraint prohibits the closure of previously opened facilities.

For our numerical experiments, we generate random problem instances with $m \in \{5, 10, 15, 20\}$ facility locations and $n = 2m$ customer locations that are located uniformly at random on the two-dimensional square $[0, 10]^2$. We set the per-unit sales profits p_{tij} to 80 minus d times the Euclidean distance between the locations of facility i and customer j, where the per-unit transportation costs are $d \in \{1, 5, 10\}$. The per-unit penalty $c_{tj} \in \{10, 25, 50\}$ for unserved demand is the same across all customer locations j and time stages t. We set the facility capacities to $C_i = 90$ for all locations

i, whereas the operating costs are set to $b_{ti} = 80 \cdot \gamma \cdot C_i$, where $\gamma \in \{0.6, 0.7, 0.8\}$ is the facility cost factor. With this choice, the per-period operating cost b_{ti} of facility i is amortized by its sales revenues if it utilizes $100\% \cdot \gamma$ of its production capacity C_i to serve customer demands. We employ a lattice regime chain uncertainty set for the customer demands (*cf.* Example 1), where the demand regime ϕ_t satisfies $\phi_1 = 100$ and $\phi_t \in \Phi_t(\phi_{t-1}) = \{\phi_{t-1}, \phi_{t-1} + \Delta\phi\}$ for $t = 2, \ldots, T$, where the demand increments $\Delta\phi$ are selected from the set $\{10, 20, 30\}$. The stage-wise uncertainty sets are

$$\Xi_t(\phi_{t-1}) = \left\{ \boldsymbol{\xi}_t = (\phi_t, \boldsymbol{\psi}_t) \in \Phi_t(\phi_{t-1}) \times \mathbb{R}^4_+ \;:\; \mathbf{e}^\top \boldsymbol{d}_t(\boldsymbol{\xi}_t) = \phi_t,\; 0.8 \cdot \frac{\phi_t}{4}\mathbf{e} \leqslant \boldsymbol{\psi}_t \leqslant 1.2 \cdot \frac{\phi_t}{4}\mathbf{e} \right\},$$

where the customer demands satisfy $d_{tj}(\boldsymbol{\xi}_t) = \sum_{f=1}^{4} \left[\frac{1}{\mathrm{dist}(j,f)} / \sum_{k=1}^{n} \frac{1}{\mathrm{dist}(k,f)} \right] \psi_{tf}$. Here, $\boldsymbol{\psi}_t \in \mathbb{R}^4_+$ is a vector of risk factors and $\mathrm{dist}(j, f)$ is the Euclidean distance between the customer location j and the f-th corner of the square $[0, 10]^2$. The uncertainty set reflects the idea that the overall demand in time stage t is known to be ϕ_t, but its geographic breakdown is governed by uncertain risk factors that are only known to reside in a hypercube. In summary, our problem instances are characterized by the five parameters m, d, c_{tj}, γ and $\Delta\phi$. We generate 50 random problem instances for each of these $4 \cdot 3^4 = 324$ parameter combinations and set the time horizon to $T = 10$.

Since the stage-wise feasible regions $\mathcal{X}_t$ of the state decisions scale exponentially in the number of facility locations m, we solve the finite adaptability approximation of the dynamic transportation-location problem with our iterative state variable selection procedure outlined in Algorithm 3 (*cf.* Section 5). In each iteration, we solve the finite adaptability problem $P(\mathcal{X}^{\mathrm{C}}, \Xi)$ using Approximation 1 with piecewise constant state decisions based on a single partition (*cf.* Section 4), and we identify the set $\hat{\Xi}$ of worst-case uncertainties with the 2 worst scenarios in this problem. We compare our method with two benchmark models. The first model ('intra-stage uncertainty') restricts the state decisions $\boldsymbol{x}_t$ to adapt only to the inter-stage uncertainties ϕ_t, but not to the intra-stage uncertainties $\boldsymbol{\psi}_t$. The resulting problem can be reformulated as a large MILP that combines a scenario tree for the demand regimes ϕ_t with robust constraints for the intra-stage demand realizations $\boldsymbol{d}_t(\boldsymbol{\xi}_t)$. The second model ('nominal intra-stage') additionally replaces the uncertain intra-stage demand realizations $\boldsymbol{d}_t(\boldsymbol{\xi}_t)$ with their nominal values $d^0_{tj} = \frac{1}{4} \sum_{f=1}^{4} \left[\frac{1}{\mathrm{dist}(j,f)} / \sum_{k=1}^{n} \frac{1}{\mathrm{dist}(k,f)} \right] \cdot \phi_t$. The resulting model is smaller and thus easier to solve, but it constitutes a weaker approximation of the original problem. In our experiments, we use the incumbent solution of 'nominal intra-stage' that has been determined after 10 minutes runtime to warm-start our iterative state variable selection procedure. We set the runtime limit of each approach to 12 hours per problem instance.

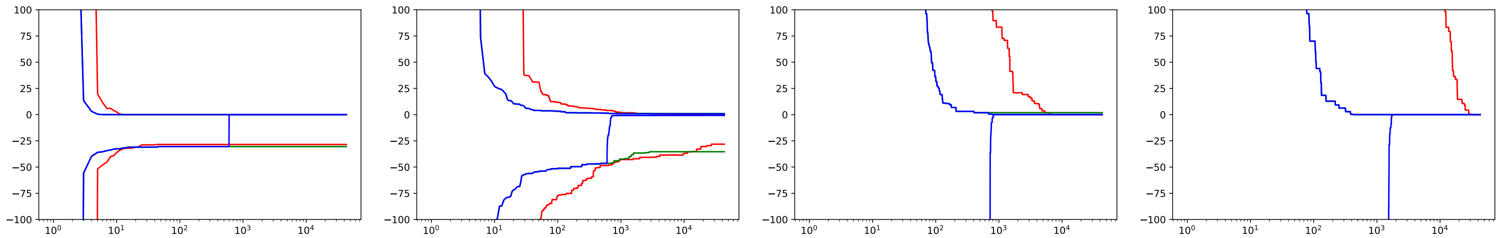

Figure 4. Average (bold lines) and 25%-75% quartile ranges (shaded areas) of the upper and lower bounds generated from our approach (blue), the 'intra-stage uncertainty' model (red) and the 'nominal intra-stage' model (green). Shown are the relative optimality gaps (ordinates) as functions of the runtime in seconds (abscissae) for $m = 5, 10, 15, 20$ (from left to right), averaged over all other $3^4 = 81$ parameter choices. Some of the lower bounds are absent since their optimality gaps lie outside of the displayed region.

Figure 4 and Table 1 present numerical results for different parameter settings. Note that each of the three methods provides both upper (conservative) and lower (progressive) bounds on the optimal value of the dynamic location-transportation problem. In particular, upper bounds can be obtained via Algorithm 1, whereas the lower bounds result from computing the crystal ball solution $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ over a reduced uncertainty set $\hat{\Xi}$ that is extracted from the upper bound problem. Our results show that the 'nominal intra-stage' model can obtain feasible solutions quickly, but the lower bounds tend to provide poor estimates of the optimal value. The 'intra-stage uncertainty' model, on the other hand, can result in good upper and lower bounds, but its runtime grows excessively as a function of the problem size. Our iterative approach, finally, provides tighter bounds than both benchmark methods under almost all parameter settings. The kinks in our lower bounds for $m = 5$ and $m = 10$ in Figure 4 are due to the transition from the incumbent solution to the 'nominal intra-stage' model after 10 minutes to the lower bounds produced by Algorithm 3. We note that while the final gap between our lower and upper bounds is typically small, in most problem instances the bounds do not collapse, see also Table 1.

Acknowledgments The research of Ward Romeijnders has been supported by the Netherlands Organisation for Scientific Research [Grant VI.Vidi.211.066]. Wolfram Wiesemann gratefully acknowledges funding from the EPSRC grants EP/R045518/1 and EP/W003317/1.

		$\Delta\phi = 10$			$\Delta\phi = 20$			$\Delta\phi = 30$		
		$\gamma = 0.6$	$\gamma = 0.7$	$\gamma = 0.8$	$\gamma = 0.6$	$\gamma = 0.7$	$\gamma = 0.8$	$\gamma = 0.6$	$\gamma = 0.7$	$\gamma = 0.8$
$d = 1$	$c_{tj} = 10$	**0.03%**	**0.03%**	**0.20%**	**0.08%**	**0.00%**	**0.07%**	**0.00%**	**0.00%**	**0.00%**
		13.80%	16.78%	23.85%	13.47%	19.83%	31.68%	46.94%	105.41%	160.32%
		14.08%	16.55%	22.35%	14.01%	19.83%	20.75%	45.37%	95.32%	161.55%
	$c_{tj} = 25$	**0.04%**	**0.03%**	**0.00%**	**0.04%**	**0.22%**	**0.00%**	**0.01%**	**0.00%**	**1.17%**
		18.97%	22.35%	52.22%	19.22%	21.37%	89.57%	46.48%	102.69%	1,860.14%
		19.49%	22.78%	51.34%	19.51%	22.66%	89.70%	46.04%	101.82%	1,721.78%
	$c_{tj} = 50$	**0.02%**	**0.06%**	**0.15%**	**0.06%**	**0.13%**	**0.04%**	**0.01%**	**0.00%**	**1.49%**
		29.62%	37.92%	59.82%	29.61%	38.69%	90.96%	46.68%	104.93%	1,809.24%
		29.22%	37.92%	59.91%	30.09%	38.65%	89.87%	46.01%	104.98%	1,712.70%
$d = 5$	$c_{tj} = 10$	**0.12%**	**0.77%**	**1.88%**	**1.31%**	**1.95%**	**6.65%**	**0.00%**	**0.00%**	**0.00%**
		11.51%	8.35%	13.46%	9.41%	10.81%	16.21%	83.24%	118.32%	191.45%
		13.23%	16.43%	24.54%	13.93%	17.12%	39.76%	69.30%	149.66%	226.17%
	$c_{tj} = 25$	**0.19%**	**0.96%**	**4.67%**	**0.71%**	**1.75%**	**1.18%**	**0.00%**	**0.00%**	**0.00%**
		17.61%	28.99%	125.98%	17.28%	41.45%	572.63%	82.72%	724.77%	115.39%
		20.83%	25.58%	122.20%	21.57%	36.38%	354.85%	78.52%	844.00%	130.67%
	$c_{tj} = 50$	**0.18%**	**0.40%**	**7.69%**	**0.59%**	**4.55%**	**0.90%**	**0.11%**	**0.00%**	**0.59%**
		31.08%	50.68%	342.05%	34.52%	64.79%	224.49%	78.75%	778.90%	120.36%
		34.18%	52.08%	290.22%	38.18%	66.19%	212.99%	76.84%	755.01%	109.46%
$d = 10$	$c_{tj} = 10$	**0.15%**	**1.30%**	3.26%	**0.48%**	**8.32%**	4.73%	**3.69%**	**1.71%**	3.90%
		9.73%	8.78%	**3.01%**	21.38%	30.85%	**3.12%**	203.35%	17.59%	**2.76%**
		21.47%	54.77%	35.37%	37.74%	151.21%	17.80%	216.81%	40.20%	17.24%
	$c_{tj} = 25$	**2.39%**	**9.87%**	**6.47%**	**8.28%**	**13.16%**	**3.12%**	**0.00%**	**2.57%**	**4.33%**
		18.17%	50.84%	10.75%	23.85%	79.68%	7.68%	301.31%	44.58%	6.89%
		27.19%	90.59%	33.02%	46.08%	106.60%	23.96%	320.86%	53.14%	17.79%
	$c_{tj} = 50$	**2.42%**	**9.27%**	**4.06%**	**12.70%**	**9.55%**	**1.58%**	**4.52%**	**1.39%**	**1.83%**
		51.57%	134.73%	38.27%	78.79%	67.62%	33.11%	205.81%	61.46%	30.49%
		60.06%	180.83%	38.38%	112.34%	72.59%	28.88%	213.31%	60.69%	34.48%

Table 1. Optimality gaps for instances with $m = 10$ facility locations. Within each cell, the numbers correspond to our iterative approach, followed by the benchmark models with and without intra-stage uncertainty. The best method is highlighted in bold print, and shaded cells correspond to instances with large optimality gaps.

References

S. Ahmed, F. G. Cabral, and B. F. P. da Costa. Stochastic Lipschitz dynamic programming. *Mathematical Programming*, 2020.

J. Ayoub and M. Poss. Decomposition for adjustable robust linear optimization subject to uncertainty polytope. *Computational Management Science*, 13(2):219–239, 2016.

C. Bandi and D. Bertsimas. Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical Programming B*, 134(1):23–70, 2012.

R. Bellman. On the computational solution of linear programming problems involving almost-block-diagonal matrices. *Management Science*, 3(4):403–406, 1957.

A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009.

D. Bertsimas and C. Caramanis. Adaptability via sampling, 2007.

D. Bertsimas and C. Caramanis. Finite adaptibility in multistage linear optimization. *IEEE Transactions on Automatic Control*, 55(12):2751–2766, 2010.

D. Bertsimas and D. den Hertog. *Robust and Adaptive Optimization*. Dynamic Ideas, 2022.

D. Bertsimas and I. Dunning. Multistage robust mixed-integer optimization with adaptive partitions. *Operations Research*, 64(4):980–998, 2015.

D. Bertsimas and A. Georghiou. Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Operations Research*, 63(3):610–627, 2015.

D. Bertsimas and A. Georghiou. Binary decision rules for multistage adaptive mixed-integer optimization. *Mathematical Programming*, 167(2):395–433, 2018.

D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Mathematical Programming*, 98 (1-3):49–71, 2003.

D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.

D. Bertsimas, D. B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011a.

D. Bertsimas, D. A. Iancu, and P. A. Parrilo. A hierarchy of near-optimal policies for multistage adaptive optimization. *IEEE Transactions on Automatic Control*, 56(12):2809–2824, 2011b.

D. Bertsimas, E. Litvinov, X. A. Sun, J. Zhao, and T. Zheng. Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Transactions on Power Systems*, 28(1):52–63, 2013.

N. Boland, I. Dumitrescu, G. Froyland, and T. Kalinowski. Minimum cardinality non-anticipativity constraints sets for multistage stochastic programming. *Mathematical Programming*, 157(2):69–93, 2016.

X. Chen and Y. Zhang. Uncertain linear programs: Extended affinely adjustable robust counterparts. *Operations Research*, 57(6):1469–1482, 2009.

X. Chen, M. Sim, P. Sun, and J. Zhang. A linear decision-based approximation approach to stochastic programming. *Operations Research*, 56(2):344–357, 2008.

Y. Chen, P. Berkhin, B. Anderson, and N. R. Devanur. Real-time bidding algorithms for performance-based display ad allocation. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1307–1315. 2011.

Z. Chen, M. Sim, and H. Xu. Distributionally robust optimization with infinitely constrained ambiguity sets. *Operations Research*, 67(5):1328–1344, 2019.

Z. Cui, D. Z. Long, J. Qi, and L. Zhang. The inventory routing problem under uncertainty. *Operations Research*, 71(1):378–395, 2023.

E. Delage and D. A. Iancu. Robust multistage decision making. *INFORMS Tutorials in Operations Research*, pages 20–46, 2015.

W. Feng, Y. Feng, and Q. Zhang. Multistage robust mixed-integer optimization under endogenous uncertainty. *European Journal of Operational Research*, 294(2):460–475, 2021.

A. Georghiou, W. Wiesemann, and D. Kuhn. Generalized decision rule approximations for stochastic programming via liftings. *Mathematical Programming*, 152(1):301–338, 2015.

A. Georghiou, A. Tsoukalas, and W. Wiesemann. Robust dual dynamic programming. *Operations Research*, 67(3):813–830, 2019.

A. Georghiou, A. Tsoukalas, and W. Wiesemann. A primal-dual lifting scheme for two-stage robust optimization. *Operations Research*, 68(2):572–590, 2020.

M. Goerigk and M. Hartisch. Multistage robust discrete optimization via quantified integer programming. *Computers & Operations Research*, 135:105434, 2021.

J. Goh and M. Sim. Distributionally robust optimization and its tractable approximations. *Operations Research*, 58(4):902–917, 2010.

E. Guslitser. Uncertainty-immunized solutions in linear programming. Master's thesis, Technion, 2002.

G. A. Hanasusanto and D. Kuhn. Conic programming reformulations of two-stage distributionally robust linear programs over Wasserstein balls. *Operations Research*, 66(3):849–869, 2018.

G. A. Hanasusanto, D. Kuhn, and W. Wiesemann. K-adaptability in two-stage robust binary programming. *Operations Research*, 63(4):877–891, 2015.

G. N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.

R. Jiang, M. Zhang, G. Li, and Y. Guan. Two-stage robust power grid optimization problem. Available on Optimization Online, 2010.

W. K. Klein Haneveld and M. H. van der Vlerk. Stochastic integer programming: General models and algorithms. *Annals of Operations Research*, 85:39–57, 1999.

D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2011.

D. Z. Long, J. Qi, and A. Zhang. Supermodularity in two-stage distributionally robust optimization. Forthcoming in *Management Science*, 2023.

A. Lorca and X. A. Sun. Adaptive robust optimization with dynamic uncertainty sets for multi-period economic dispatch under significant wind. *IEEE Transactions on Power Systems*, 30(4):1702–1713, 2015.

A. Martin. Integer programs with block structure. Technical Report SC-99-03, Zuse Institute Berlin, 1999.

W. Miao, H. Yin, D. Du, and J. Han. Robust optimization model for a class of uncertain linear programs. In B. Chen, M. Paterson, and G. Zhang, editors, *Combinatorics, Algorithms, Probabilistic and Experimental Methodologies*, volume 4614 of *Lecture Notes in Computer Science*, pages 384–395. Springer, 2007.

A. Nilim and L. El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

M. V. F. Pereira and L. M. V. G. Pinto. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52(1):359–375, 1991.

A. B. Philpott, F. Wahid, and J. F. Bonnans. Midas: A mixed integer dynamic approximation scheme. *Mathematical Programming*, (1):19–50, 2020.

K. Postek and D. den Hertog. Multistage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set. *INFORMS Journal on Computing*, 28(3):553–574, 2016.

R. T. Rockafellar and R. J.-B. Wets. *Variational Analysis*. Springer, 1998.

W. Romeijnders and K. Postek. Piecewise constant decision rules via branch-and-bound based scenario detection for integer adjustable robust optimization. *INFORMS Journal on Computing*, 33(1):390–400, 2021.

W. Römisch and R. Schultz. Multistage stochastic integer programs: An introduction. In M. Grötschel, S. O. Krumke, and J. Rambau, editors, *Online Optimization of Large Scale Systems*, pages 581–600. Springer, 2001.

R. Schultz. Stochastic programming with integer variables. *Mathematical Programming B*, 97(1):285–309, 2003.

Suvrajeet Sen. Algorithms for stochastic mixed-integer programming models. In K. Aardal, G.L. Nemhauser, and R. Weismantel, editors, *Discrete Optimization*, volume 12 of *Handbooks in Operations Research and Management Science*, pages 515–558. Elsevier, 2005.

A. Shapiro. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209(1):63–72, 2011.

A. Shapiro. Interchangeability principle and dynamic equations in risk averse stochastic programming. *Operations Research Letters*, 45(4):377–381, 2017.

A. Subramanyam, C. E. Gounaris, and W. Wiesemann. K-adaptability in two-stage mixed-integer robust optimization. *Mathematical Programming Computation*, 12(2):193–224, 2020.

A. Thiele, T. Terry, and M. Epelman. Robust linear optimization with recourse. Technical report, Lehigh University and University of Michigan, 2010.

P. Vayanos, D. Kuhn, and B. Rustem. Decision rules for information discovery in multi-stage stochastic programming, 2011.

W. Wiesemann, D. Kuhn, and B. Rustem. Robust Markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013.

G. Xu and S. Burer. A copositive approach for two-stage adjustable robust optimization with uncertain right-hand sides. *Computational Optimization and Applications*, 70(1):33–59, 2018.

I. Yanıkoğlu, B. L. Gorissen, and D. den Hertog. A survey of adjustable robust optimization. *European Journal of Operational Research*, 277(3):799–813, 2019.

B. Zeng and L. Zhao. Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–561, 2013.

S. Zhang and X. A. Sun. Stochastic dual dynamic programming for multistage stochastic mixed-integer nonlinear optimization. *Mathematical Programming*, 196(1–2):935–985, 2022.

C. Zhao, J. Wang, J.-P. Watson, and Y. Guan. Multi-stage robust unit commitment considering wind and demand response uncertainties. *IEEE Transactions on Power Systems*, 28(3):2708–2717, 2013.

L. Zhao and B. Zeng. An exact algorithm for two-stage robust optimization with mixed integer recourse problems. Technical report, University of South Florida, 2012.

J. Zhen, D. den Hertog, and M. Sim. Adjustable robust optimization via Fourier-Motzkin elimination. *Operations Research*, 66(4):1086–1100, 2018.

J. Zou, S. Ahmed, and X. A. Sun. Stochastic dual dynamic integer programming. *Mathematical Programming*, 175(1):461–502, 2019.

Appendix: Proofs

Proof of Lemma 1. This follows immediately from Proposition 2.1 and Example 2 of Shapiro (2017). □

Proof of Proposition 1. We show that the finite adaptability approximation (3) shares the same optimal value as the nested optimization problem

$$\max_{\boldsymbol{\xi}_1 \in \Xi_1} \min_{(\boldsymbol{x}_1, \boldsymbol{y}_1) \in F_1(\boldsymbol{\xi}_1)} \Big[\boldsymbol{q}_1(\boldsymbol{\xi}_1)^\top \boldsymbol{x}_1 + \boldsymbol{r}_1{}^\top \boldsymbol{y}_1 \;+ \max_{\boldsymbol{\xi}_2 \in \Xi_2(\boldsymbol{\phi}_1)} \min_{(\boldsymbol{x}_2, \boldsymbol{y}_2) \in F_2(\boldsymbol{x}_1, \boldsymbol{\xi}_2)} \Big[\boldsymbol{q}_2(\boldsymbol{\xi}_2)^\top \boldsymbol{x}_2 + \boldsymbol{r}_2{}^\top \boldsymbol{y}_2 \\ + \cdots + \max_{\boldsymbol{\xi}_T \in \Xi_T(\boldsymbol{\phi}_{T-1})} \min_{(\boldsymbol{x}_T, \boldsymbol{y}_T) \in F_T(\boldsymbol{x}_{T-1}, \boldsymbol{\xi}_T)} \Big[\boldsymbol{q}_T(\boldsymbol{\xi}_T)^\top \boldsymbol{x}_T + \boldsymbol{r}_T{}^\top \boldsymbol{y}_T \Big] \cdots \Big], \tag{9}$$

where for all $t = 1, \ldots, T$, $\boldsymbol{x}_{t-1} \in \mathcal{X}_{t-1}^{\mathrm{C}}$ and $\boldsymbol{\xi}_t = (\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \in \Xi_t$,

$$F_t(\boldsymbol{x}_{t-1}, \boldsymbol{\xi}_t) \;=\; \Big\{ (\boldsymbol{x}_t, \boldsymbol{y}_t) \in \mathcal{X}_t^{\mathrm{C}} \times \mathbb{R}^{n_t^2} \;:\; \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t) \Big\},$$

and where we notationally suppress the dependence on $\boldsymbol{x}_0$ in $F_1(\boldsymbol{\xi}_1)$. The statement of the proposition then follows from the fact that for all $t = 1, \ldots, T$, the value of $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1})$ coincides with the optimal value of the partial problem

$$\max_{\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})} \min_{(\boldsymbol{x}_t, \boldsymbol{y}_t) \in F_t(\hat{\boldsymbol{x}}_{t-1}, \boldsymbol{\xi}_t)} \Big[\boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \cdots \\ + \max_{\boldsymbol{\xi}_T \in \Xi_T(\boldsymbol{\phi}_{T-1})} \min_{(\boldsymbol{x}_T, \boldsymbol{y}_T) \in F_T(\boldsymbol{x}_{T-1}, \boldsymbol{\xi}_T)} \Big[\boldsymbol{q}_T(\boldsymbol{\xi}_T)^\top \boldsymbol{x}_T + \boldsymbol{r}_T{}^\top \boldsymbol{y}_T \Big] \cdots \Big],$$

which can be readily shown by a backward induction on t.

To show the equivalence between (3) and (9), we denote by $\Xi_{1:t} = \{\boldsymbol{\xi}^t \,:\, \boldsymbol{\xi} \in \Xi\}$ the projection of Ξ onto the first t stages, and we rewrite the finite adaptability problem (3) as

$$\min_{(\boldsymbol{x}_1, \boldsymbol{y}_1) \in \mathcal{F}_1} \min_{(\boldsymbol{x}_2, \boldsymbol{y}_2) \in \mathcal{F}_2(\boldsymbol{x}_1)} \cdots \min_{(\boldsymbol{x}_T, \boldsymbol{y}_T) \in \mathcal{F}_T(\boldsymbol{x}_{T-1})} \max_{\boldsymbol{\xi} \in \Xi} \sum_{t=1}^{T} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t(\boldsymbol{\xi}^t),$$

where for all $t = 1, \ldots, T$,

$$\mathcal{F}_t(\boldsymbol{x}_{t-1}) \;=\; \Big\{ [(\boldsymbol{x}_t, \boldsymbol{y}_t) : \Xi_{1:t} \to \mathcal{X}_t^{\mathrm{C}} \times \mathbb{R}^{n_t^2}] \;:\; (\boldsymbol{x}_t, \boldsymbol{y}_t)(\boldsymbol{\xi}^t) \in F_t(\boldsymbol{x}_{t-1}(\boldsymbol{\xi}^{t-1}), \boldsymbol{\xi}_t) \;\; \forall \boldsymbol{\xi}^t \in \Xi_{1:t} \Big\},$$

and where we notationally suppress the dependence on $\boldsymbol{x}_0$ in $\mathcal{F}_1$. By Lemma 1, we can interchange the innermost minimization and maximization operators, and thus the finite adaptability approximation (3) is equivalent to

$$\min_{(\boldsymbol{x}_1,\boldsymbol{y}_1)\in\mathcal{F}_1} \cdots \min_{(\boldsymbol{x}_{T-1},\boldsymbol{y}_{T-1})\in\mathcal{F}_{T-1}(\boldsymbol{x}_{T-2})}$$

$$\max_{\boldsymbol{\xi}\in\Xi} \min_{(\boldsymbol{x}_T,\boldsymbol{y}_T)\in F_T(\boldsymbol{x}_{T-1}(\boldsymbol{\xi}^{T-1}),\boldsymbol{\xi}_T)} \left[\sum_{t=1}^{T-1} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t(\boldsymbol{\xi}^t)\right] + \boldsymbol{q}_T(\boldsymbol{\xi}_T)^\top \boldsymbol{x}_T + \boldsymbol{r}_T{}^\top \boldsymbol{y}_T.$$

Decomposing the maximization over $\boldsymbol{\xi} \in \Xi$ into one over $\boldsymbol{\xi}^{T-1} \in \Xi_{1:T-1}$ and one over $\boldsymbol{\xi}_T \in \Xi_T(\boldsymbol{\phi}_{T-1})$, which is admissible due to the generalized rectangularity of Ξ, the problem becomes

$$\begin{aligned}
&\min_{(\boldsymbol{x}_1,\boldsymbol{y}_1)\in\mathcal{F}_1} \cdots \min_{(\boldsymbol{x}_{T-1},\boldsymbol{y}_{T-1})\in\mathcal{F}_{T-1}(\boldsymbol{x}_{T-2})} \max_{\boldsymbol{\xi}^{T-1}\in\Xi_{1:T-1}} \\
&\qquad \max_{\boldsymbol{\xi}_T\in\Xi_T(\boldsymbol{\phi}_{T-1})} \min_{(\boldsymbol{x}_T,\boldsymbol{y}_T)\in F_T(\boldsymbol{x}_{T-1}(\boldsymbol{\xi}^{T-1}),\boldsymbol{\xi}_T)} \left[\sum_{t=1}^{T-1} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t(\boldsymbol{\xi}^t)\right] + \boldsymbol{q}_T(\boldsymbol{\xi}_T)^\top \boldsymbol{x}_T + \boldsymbol{r}_T{}^\top \boldsymbol{y}_T \\
= &\min_{(\boldsymbol{x}_1,\boldsymbol{y}_1)\in\mathcal{F}_1} \cdots \min_{(\boldsymbol{x}_{T-1},\boldsymbol{y}_{T-1})\in\mathcal{F}_{T-1}(\boldsymbol{x}_{T-2})} \max_{\boldsymbol{\xi}^{T-1}\in\Xi_{1:T-1}} \left[\sum_{t=1}^{T-1} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t(\boldsymbol{\xi}^t)\right] \\
&\qquad + \max_{\boldsymbol{\xi}_T\in\Xi_T(\boldsymbol{\phi}_{T-1})} \min_{(\boldsymbol{x}_T,\boldsymbol{y}_T)\in F_T(\boldsymbol{x}_{T-1}(\boldsymbol{\xi}^{T-1}),\boldsymbol{\xi}_T)} \boldsymbol{q}_T(\boldsymbol{\xi}_T)^\top \boldsymbol{x}_T + \boldsymbol{r}_T{}^\top \boldsymbol{y}_T.
\end{aligned}$$

We can now invoke Lemma 1 to interchange the minimization over $(\boldsymbol{x}_{T-1}, \boldsymbol{y}_{T-1})$ with the maximization over $\boldsymbol{\xi}^{T-1}$. A backward induction over all time stages then proves that the finite adaptability approximation (3) indeed attains the same optimal value as the nested formulation (9). □

Proof of Proposition 2. To prove that $\{(\boldsymbol{x}_t^\star, \boldsymbol{y}_t^\star)\}_{t=1}^T$ is optimal in (3), it suffices to show that

$$\sum_{t=1}^{T} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^\star(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t^\star(\boldsymbol{\xi}^t) \;\leqslant\; \mathcal{Q}_1 \qquad \forall \boldsymbol{\xi} \in \Xi \tag{10}$$

since $\mathcal{Q}_1$ coincides with the optimal value of (3) by Proposition 1. To this end, we conduct a backward induction on t to show that for all $\boldsymbol{\xi} \in \Xi$, we have

$$\sum_{\tau=t}^{T} \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}_\tau^\star(\boldsymbol{\xi}^\tau) + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}_\tau^\star(\boldsymbol{\xi}^\tau) \;\leqslant\; \mathcal{Q}_t(\boldsymbol{x}_{t-1}^\star(\boldsymbol{\xi}^{t-1}); \boldsymbol{\phi}_{t-1})$$

with our usual convention that $\boldsymbol{x}_0^\star(\boldsymbol{\xi}^0) \equiv \boldsymbol{0}$. Indeed, the base case $t = T$ holds since for any $\boldsymbol{\xi} \in \Xi$,

$$\begin{aligned}
\boldsymbol{q}_T(\boldsymbol{\xi}_T)^\top \boldsymbol{x}_T^\star(\boldsymbol{\xi}^T) + \boldsymbol{r}_T{}^\top \boldsymbol{y}_T^\star(\boldsymbol{\xi}) \;&=\; Q_T(\boldsymbol{x}_{T-1}^\star(\boldsymbol{\xi}^{T-1}); \boldsymbol{\xi}_T) \\
&\leqslant\; \max_{\boldsymbol{\xi}'_T\in\Xi_T(\boldsymbol{\phi}_{T-1})} Q_T(\boldsymbol{x}_{T-1}^\star(\boldsymbol{\xi}^{T-1}); \boldsymbol{\xi}'_T) \\
&=\; \mathcal{Q}_T(\boldsymbol{x}_{T-1}^\star(\boldsymbol{\xi}^{T-1}); \boldsymbol{\phi}_{T-1}),
\end{aligned}$$

where the first equality holds since $(\boldsymbol{x}_T^\star(\boldsymbol{\xi}^T), \boldsymbol{y}_T^\star(\boldsymbol{\xi}^T))$ is an optimal solution to the nominal problem $Q_T(\boldsymbol{x}_{T-1}^\star(\boldsymbol{\xi}^{T-1}); \boldsymbol{\xi}_T)$ by the design of Algorithm 2. Next, assume that the induction hypothesis holds for stage $t+1$. We then have that

$$\begin{aligned}
\sum_{\tau=t}^{T} \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}_\tau^\star(\boldsymbol{\xi}^\tau) + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}_\tau^\star(\boldsymbol{\xi}^\tau) &= \Big[\boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^\star(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t^\star(\boldsymbol{\xi}^t)\Big] + \sum_{\tau=t+1}^{T} \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}_\tau^\star(\boldsymbol{\xi}^\tau) + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}_\tau^\star(\boldsymbol{\xi}^\tau) \\
&\leqslant \Big[\boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^\star(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t^\star(\boldsymbol{\xi}^t)\Big] + \mathcal{Q}_{t+1}(\boldsymbol{x}^\star(\boldsymbol{\xi}^t); \boldsymbol{\phi}_t) \\
&= Q_t(\boldsymbol{x}_{t-1}^\star(\boldsymbol{\xi}^{t-1}); \boldsymbol{\xi}_t) \;\leqslant\; \max_{\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})} Q_t(\boldsymbol{x}_{t-1}^\star(\boldsymbol{\xi}^{t-1}); \boldsymbol{\xi}_t) \\
&= \mathcal{Q}_t(\boldsymbol{x}_{t-1}^\star(\boldsymbol{\xi}^{t-1}); \boldsymbol{\phi}_{t-1}).
\end{aligned}$$

Here, the first inequality holds because of the induction hypothesis, while the second equality is due to the fact that $(\boldsymbol{x}_t^\star(\boldsymbol{\xi}^t), \boldsymbol{y}_t^\star(\boldsymbol{\xi}^t))$ is an optimal solution to the nominal problem $Q_t(\boldsymbol{x}_{t-1}^\star(\boldsymbol{\xi}^{t-1}); \boldsymbol{\xi}_t)$ by the design of Algorithm 2. In particular, the above inequality also holds for $t = 1$, which shows that (10) holds and thus Algorithm 2 *implicitly* defines an optimal policy $\{(\boldsymbol{x}_t^\star, \boldsymbol{y}_t^\star)\}_{t=1}^T$ for (3). $\square$

Proof of Theorem 1. We first prove *weak duality*, that is, the optimal value of Q_t is bounded from below by the optimal value of (6). This is trivially the case if Q_t is infeasible. Assume therefore that Q_t is feasible, and fix any feasible decision $(\boldsymbol{x}_t, \boldsymbol{y}_t) \in \mathcal{X}_t^{\mathrm{C}} \times \mathbb{R}^{n_t^2}$ in Q_t and any feasible decision $\theta \in \mathbb{R}$ in (6). The feasibility of $(\boldsymbol{x}_t, \boldsymbol{y}_t)$ in Q_t implies that $\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t)$, which in turns implies that $\theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \mathcal{Q}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t)$ as θ is feasible in (6).

We now prove *strong duality*, that is, the optimal values of Q_t and (6) coincide. If Q_t is infeasible, that is, if $\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \ngeqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t)$ for all $(\boldsymbol{x}_t, \boldsymbol{y}_t) \in \mathcal{X}_t^{\mathrm{C}} \times \mathbb{R}^{n_t^2}$, then (6) is unbounded as desired. Let us now assume that Q_t is minimized by $(\boldsymbol{x}_t^\star, \boldsymbol{y}_t^\star) \in \mathcal{X}_t \times \mathbb{R}^{n_t^2}$, that is, we have $\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t^\star + \boldsymbol{V}_t\, \boldsymbol{y}_t^\star \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t)$, and for all $(\boldsymbol{x}_t, \boldsymbol{y}_t) \in \mathcal{X}_t^{\mathrm{C}} \times \mathbb{R}^{n_t^2}$ we either have $\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \ngeqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t)$ or $\boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \mathcal{Q}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \geqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^\star + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t^\star + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^\star; \boldsymbol{\phi}_t)$. Then the solution $\theta^\star = \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^\star + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t^\star + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^\star; \boldsymbol{\phi}_t)$ is feasible in (6) by construction. $\square$

Proof of Proposition 3. The semi-infinite disjunction in problem (6) can be expressed as

$$\theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t^i + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t) \qquad \forall i = 1, \ldots, p_t,\ \forall \boldsymbol{y}_t^i \in \mathcal{Y}_t(\hat{\boldsymbol{x}}_{t-1}, \boldsymbol{\xi}_t, \boldsymbol{x}_t^i),$$

where

$$\mathcal{Y}_t(\hat{\boldsymbol{x}}_{t-1}, \boldsymbol{\xi}_t, \boldsymbol{x}_t) \;=\; \left\{ \boldsymbol{y}_t \in \mathbb{R}^{n_t^2} \;:\; \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t + \boldsymbol{V}_t\,\boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t) \right\}.$$

This constraint, in turn, can be reformulated as

$$\theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \inf \left\{ {\boldsymbol{r}_t}^\top \boldsymbol{y}_t^i \;:\; \boldsymbol{y}_t^i \in \mathcal{Y}_t(\hat{\boldsymbol{x}}_{t-1}, \boldsymbol{\xi}_t, \boldsymbol{x}_t^i) \right\} + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t) \qquad \forall i = 1, \ldots, p_t. \tag{11}$$

The minimization problem embedded in this constraint,

$$\begin{array}{ll} \text{minimize} & {\boldsymbol{r}_t}^\top \boldsymbol{y}_t^i \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t^i + \boldsymbol{V}_t\,\boldsymbol{y}_t^i \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t) \\ & \boldsymbol{y}_t^i \in \mathbb{R}^{n_t^2}, \end{array} \tag{12}$$

has the linear programming dual

$$\begin{array}{ll} \text{maximize} & \left[\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\hat{\boldsymbol{x}}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t^i\right]^\top \boldsymbol{\lambda}_t^i \\ \text{subject to} & \boldsymbol{V}_t^\top \boldsymbol{\lambda}_t^i = \boldsymbol{r}_t \\ & \boldsymbol{\lambda}_t^i \in \mathbb{R}_+^{m_t}. \end{array}$$

Strong linear programming duality holds unless both the primal minimization problem and the dual maximization problem are infeasible. Note that the feasibility of the dual maximization problem does not depend on $\hat{\boldsymbol{x}}_{t-1}$, $\boldsymbol{x}_t^i$ or $\boldsymbol{\xi}_t$. In particular, the dual maximization problem is infeasible for *some* $\hat{\boldsymbol{x}}_{t-1}$ and $\boldsymbol{\xi}_t$ if and only if it is infeasible for *all* $\hat{\boldsymbol{x}}_{t-1}$ and $\boldsymbol{\xi}_t$. In the latter case, the primal problem has to be infeasible or unbounded for all $\hat{\boldsymbol{x}}_{t-1}$ and $\boldsymbol{\xi}_t$. By assumption, the set of admissible $\boldsymbol{y}_t^i$ in problem (12) is bounded for all $\hat{\boldsymbol{x}}_{t-1}$ and $\boldsymbol{\xi}_t$, which implies that the primal problem would have to be infeasible for all $\hat{\boldsymbol{x}}_{t-1}$ and $\boldsymbol{\xi}_t$. This, however, contradicts our assumption that the stage-t nominal cost to-go problem Q_t is feasible for some $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$. We thus conclude that strong linear programming duality holds.

The statement of the proposition now follows from substituting the linear programming dual (12) into the constraint (11), replacing the semi-infinite constraint in problem (6) with the revised constraint (11) and finally embedding the updated problem (6) in the overall two-stage subproblem $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1})$ that optimizes over the uncertain parameters $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$. □

Proof of Corollary 1. Using similar arguments as in the proof of Proposition 3, the semi-infinite disjunction in problem (6) can be expressed as

$$\theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \inf \left\{ {\boldsymbol{r}_t}^\top \boldsymbol{y}_t^i \;:\; \boldsymbol{T}_t\,\hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t\,\boldsymbol{x}_t^i + \boldsymbol{V}_t\,\boldsymbol{y}_t^i \geqslant \boldsymbol{h}_t, \;\; \boldsymbol{y}_t^i \in \mathbb{R}^{n_t^2} \right\} + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t)$$

for all $i = 1, \ldots, p_t$, see also constraint (11). The statement then follows from substituting this constraint into the stage-t nominal cost to-go problem (6) and optimizing over the uncertainties $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$. Note that the validity of this reformulation does not depend on the feasibility of Q_t since we do not rely on strong duality in this proof. $\square$

Proof of Corollary 2. Similar arguments as in the proof of Theorem 1 show that when the control variables $\boldsymbol{y}_t$ are absent, then the nominal cost to-go $Q_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t)$ equals the optimal value of the problem

$$
\begin{array}{ll}
\text{maximize} & \theta \\
\text{subject to} & \left[\theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t + \mathcal{Q}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \ \vee\ \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t \ngeqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t)\right] \quad \forall \boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}} \\
& \theta \in \mathbb{R}.
\end{array}
\tag{13}
$$

We claim that for fixed $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$, the problems (8) and (13) share the same optimal value. The statement then follows from the fact that $\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1}) = \max\left\{Q_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t) \,:\, \boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})\right\}$.

We first show that for fixed $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$, the optimal value of (13) is bounded from above by the optimal value of (8). To this end, we show that any $\theta \in \mathbb{R}$ feasible in (13) is also feasible in (8) if we complement it with the variable assignment $(\boldsymbol{z}_t, \boldsymbol{\lambda}_t)$ detailed next. If $\boldsymbol{x}_t^i \in \mathcal{X}_t^{\mathrm{C}}$, $i = 1, \ldots, p_t$, satisfies $\theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t)$, then the variable assignment $z_t^i = 0$ as well as $\lambda_t^{i1} = \ldots = \lambda_t^{im_t} = 1$ satisfies the relevant constraints in problem (8). Likewise, if $\boldsymbol{x}_t^i \in \mathcal{X}_t^{\mathrm{C}}$ satisfies $\left[\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t^i\right]_\ell < \left[\boldsymbol{h}_t(\boldsymbol{\xi}_t)\right]_\ell$ for at least one $\ell = 1, \ldots, m_t$, then the variable assignment $z_t^i = 1$ as well as $\lambda_t^{il} = \mathbf{1}_{[l \neq \ell]}$, $l = 1, \ldots, m_t$, satisfies the relevant constraints in problem (8).

We now prove that for fixed $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$, the optimal value of (8) is also bounded from above by the optimal value of (13). To this end, we show that any solution $(\theta, \boldsymbol{z}_t, \boldsymbol{\lambda}_t)$ feasible in (8) implies that the solution θ satisfies the disjunctive constraint in (13) for each $\boldsymbol{x}_t^i$, $i = 1, \ldots, p_t$. Indeed, if $z_t^i = 0$, then $\theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t)$. If, on the other hand, $z_t^i = 1$, then $\lambda_t^{i\ell} = 0$ for at least one $\ell = 1, \ldots, m_t$ since $z_t^i + \sum_{\ell=1}^{m_t} \lambda_t^{i\ell} \leqslant m_t$. We thus conclude that $\left[\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t^i\right]_\ell < \left[\boldsymbol{h}_t(\boldsymbol{\xi}_t)\right]_\ell$, that is, $\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t^i \ngeqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t)$ in this case, as desired. $\square$

Proof of Corollary 3. If the recourse matrix $\boldsymbol{V}_t$ is invertible, then so is its transpose $\boldsymbol{V}_t^\top$, and we can therefore conduct the variable substitution $\boldsymbol{\lambda}_t^i \leftarrow (\boldsymbol{V}_t^{-1})^\top \boldsymbol{r}_t$ in the stage-t worst-case cost to-go problem (7). If $(\boldsymbol{V}_t^{-1})^\top \boldsymbol{r}_t \ngeqslant \mathbf{0}$, then $\boldsymbol{\lambda}_t^i \ngeqslant \mathbf{0}$ and problem (7) is infeasible; otherwise, the variable substitution $\boldsymbol{\lambda}_t^i \leftarrow (\boldsymbol{V}_t^{-1})^\top \boldsymbol{r}_t$ leads to the maximization problem in the statement of the corollary. $\square$

Proof of Corollary 4. By eliminating the epigraphical variable θ and separating the maximization over $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$ and $\boldsymbol{\lambda}_t^i \in \mathbb{R}_+^{m_t}$, we can equivalently express problem (7) as

$$\begin{aligned} \text{maximize} \quad & \min_{i=1,\ldots,p_t} \max_{\boldsymbol{\lambda}_t^i \in \mathbb{R}_+^{m_t}} \Big\{ \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \big[\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\hat{\boldsymbol{x}}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t^i\big]^\top \boldsymbol{\lambda}_t^i \\ & \qquad\qquad + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t) \;:\; \boldsymbol{V}_t^\top \boldsymbol{\lambda}_t^i = \boldsymbol{r}_t \Big\} \\ \text{subject to} \quad & \boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1}). \end{aligned} \tag{14}$$

Here, we exchanged the order of the maximization over $\boldsymbol{\lambda}_t^i$ and the minimization over i, which is admissible since the problem decomposes in the decisions $\boldsymbol{\lambda}_t^i$, irrespective of the structure of $\boldsymbol{V}_t$.

For any fixed $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$, problem (14) is unbounded if and only if for every $i = 1, \ldots, p_t$, the vector $\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\hat{\boldsymbol{x}}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t^i$ has a strictly positive inner product with one of the extreme rays of the polyhedron $\{\boldsymbol{\lambda} \geqslant \boldsymbol{0} \,:\, \boldsymbol{V}_t^\top \boldsymbol{\lambda} = \boldsymbol{r}_t\}$, that is, if and only if the problem

$$\begin{aligned} \text{maximize} \quad & \theta \\ \text{subject to} \quad & \theta \leqslant \mathrm{M} \cdot \sum_{\boldsymbol{\gamma}_t \in \Gamma_t} \delta_t^i(\boldsymbol{\gamma}_t) && \forall i = 1, \ldots, p_t \\ & \big[\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\hat{\boldsymbol{x}}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t^i\big]^\top \boldsymbol{\gamma}_t > \mathrm{M} \cdot (\delta_t^i(\boldsymbol{\gamma}_t) - 1) && \forall \boldsymbol{\gamma}_t \in \Gamma_t,\ \forall i = 1, \ldots, p_t \\ & \theta \in \mathbb{R},\ \ \delta_t^i : \Gamma_t \to \{0,1\},\ i = 1, \ldots, p_t \end{aligned}$$

has an optimal value greater than or equal to M, where Γ_t is the set of extreme rays of the polyhedron $\{\boldsymbol{\lambda} \geqslant \boldsymbol{0} \,:\, \boldsymbol{V}_t^\top \boldsymbol{\lambda} = \boldsymbol{r}_t\}$. Likewise, if (14) is bounded for a fixed $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$, then it is optimized by extreme points $\boldsymbol{\lambda}_t^i$ of the polyhedron $\{\boldsymbol{\lambda} \geqslant \boldsymbol{0} \,:\, \boldsymbol{V}_t^\top \boldsymbol{\lambda}_t = \boldsymbol{r}_t\}$, $i = 1, \ldots, p_t$, since the inner maximization in (14) is affine in $\boldsymbol{\lambda}_t^i$. In that case, the optimal value of (14) with $\boldsymbol{\xi}_t$ fixed equals

$$\begin{aligned} \text{maximize} \quad & \theta \\ \text{subject to} \quad & \theta \leqslant \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^i + \big[\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\hat{\boldsymbol{x}}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_t^i\big]^\top \boldsymbol{\lambda}_t + \mathcal{Q}_{t+1}(\boldsymbol{x}_t^i; \boldsymbol{\phi}_t) + \mathrm{M} \cdot z_t^i(\boldsymbol{\lambda}_t) \\ & && \forall \boldsymbol{\lambda}_t \in \Lambda_t,\ \forall i = 1, \ldots, p_t \\ & \sum_{\boldsymbol{\lambda}_t \in \Lambda_t} z_t^i(\boldsymbol{\lambda}_t) = |\Lambda_t| - 1 && \forall i = 1, \ldots, p_t \\ & z_t^i : \Lambda_t \to \{0,1\}, \end{aligned}$$

where Λ_t is the set of extreme points of the polyhedron $\{\boldsymbol{\lambda} \geqslant \boldsymbol{0} \,:\, \boldsymbol{V}_t^\top \boldsymbol{\lambda} = \boldsymbol{r}_t\}$.

We obtain the optimization problem in the statement of the corollary, finally, if we combine the unbounded and bounded case from above into a single optimization problem, optimize over $\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})$ and exploit the assumed decomposition $\Lambda_t = \times_{j=1}^{k_t} \Lambda_t^j$ and $\Gamma_t = \times_{j=1}^{k_t} \Gamma_t^j$. □

Proof of Observation 1. In view of the first statement, we observe that

$$\begin{aligned}
\mathcal{Q}_t(\hat{\boldsymbol{x}}_{t-1}; \boldsymbol{\phi}_{t-1}) \;&=\; \max_{\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})} \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \mathcal{Q}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t) \\ & \boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}},\ \boldsymbol{y}_t \in \mathbb{R}^{n_t^2} \end{array} \right] \\
&\leqslant\; \max_{\boldsymbol{\xi}_t \in \Xi_t(\boldsymbol{\phi}_{t-1})} \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\xi}_t) \\ & \boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}},\ \boldsymbol{y}_t \in \mathbb{R}^{n_t^2} \end{array} \right] \\
&=\; \overline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}; \boldsymbol{\phi}_{t-1}),
\end{aligned}$$

where the first and last equality hold by definition, and the inequality holds because $\overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \geqslant \mathcal{Q}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t)$ for all $\boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}}$ and $\boldsymbol{\phi}_t \in \Phi_t$.

As for the second statement, we show that

$$\sum_{t=1}^{T} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^\star(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t^\star(\boldsymbol{\xi}^t) \;\leqslant\; \overline{\mathcal{Q}}_1 \qquad \forall \boldsymbol{\xi} \in \Xi.$$

To this end, we conduct a backward induction on t to show that for all $\boldsymbol{\xi} \in \Xi$, we have

$$\sum_{\tau=t}^{T} \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}_\tau^\star(\boldsymbol{\xi}^\tau) + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}_\tau^\star(\boldsymbol{\xi}^\tau) \;\leqslant\; \overline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}^\star(\boldsymbol{\xi}^{t-1}); \boldsymbol{\phi}_{t-1})$$

with our usual convention that $\boldsymbol{x}_0^\star(\boldsymbol{\xi}^0) \equiv \mathbf{0}$. The backward induction follows the same reasoning as in the proof of Proposition 2, with $\overline{Q}_t$ and $\overline{\mathcal{Q}}_t$ taking the roles of Q_t and $\mathcal{Q}_t$, respectively, $t = 1, \ldots, T$. We omit the details for the sake of brevity. □

Proof of Observation 2. The two-stage subproblem $\overline{\mathcal{Q}}_t$ admits the following representation as a two-stage robust optimization problem:

$$\max_{\substack{\boldsymbol{\phi}_t \in \Phi_t(\boldsymbol{\phi}_{t-1}), \\ j \in \{1,\ldots,l\}}} \; \max_{\boldsymbol{\psi}_t \in \Psi_{tj}(\boldsymbol{\phi}_t)} \; \min_{\boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}}} \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \\ & \boldsymbol{y}_t \in \mathbb{R}^{n_t^2} \end{array} \right]$$

Here, we have replaced the intra-stage uncertainties $\boldsymbol{\psi}_t$ with pairs $(j, \boldsymbol{\psi}_{tj}) \in \{1, \ldots, l\} \times \Psi_{tj}(\boldsymbol{\phi}_t)$, where j denotes the realized set of the polyhedral partition $\{\Psi_{tj}(\boldsymbol{\phi}_t)\}_{j=1}^{l}$ and $\boldsymbol{\psi}_{tj}$ the intra-stage realization within the j-th set, respectively, and we have split the maximization and minimization problems into two separate subproblems each.

Under Approximation 1, the state variables $\boldsymbol{x}_t$ no longer adapt to $(j, \boldsymbol{\psi}_{tj}) \in \{1, \ldots, l\} \times \Psi_{tj}(\boldsymbol{\phi}_t)$, but only to $j \in \{1, \ldots, l\}$. This amounts to exchanging the order of the inner maximization and the outer minimization in the above optimization problem:

$$\max_{\substack{\boldsymbol{\phi}_t \in \Phi_t(\boldsymbol{\phi}_{t-1}),\\ j \in \{1,\ldots,l\}}} \ \min_{\boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}}} \ \max_{\boldsymbol{\psi}_t \in \Psi_{tj}(\boldsymbol{\phi}_t)} \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)^\top \boldsymbol{x}_t + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t; \boldsymbol{\phi}_t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \boldsymbol{x}_t + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \\ & \boldsymbol{y}_t \in \mathbb{R}^{n_t^2} \end{array} \right]$$

One readily recognizes that the expression in the statement of the observation computes the optimal value of this optimization problem. □

Proof of Observation 3. The two-stage subproblem $\overline{\mathcal{Q}}_t$ admits the following representation as a two-stage robust optimization problem:

$$\max_{\boldsymbol{\phi}_t \in \Phi_t(\boldsymbol{\phi}_{t-1})} \ \max_{\boldsymbol{\psi}_t \in \Psi_t(\boldsymbol{\phi}_t)} \ \min_{\boldsymbol{y}_t \in \mathbb{R}^{n_t^2}} \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)^\top \boldsymbol{x}_t^k + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t + \overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t^k; \boldsymbol{\phi}_t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \boldsymbol{x}_t^k + \boldsymbol{V}_t\, \boldsymbol{y}_t \geqslant \boldsymbol{h}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \\ & k \in \{1, \ldots, p_t\} \end{array} \right]$$

Here, we have replaced the minimization over $\boldsymbol{x}_t \in \mathcal{X}_t^{\mathrm{C}}$ with the equivalent minimization over $\boldsymbol{x}_t^k$, $k \in \{1, \ldots, p_t\}$, and we have split the maximization and minimization problems into two separate subproblems each. We can further exchange the order of the inner maximization and the outer minimization operators by optimizing over decisions rules as follows:

$$\max_{\boldsymbol{\phi}_t \in \Phi_t(\boldsymbol{\phi}_{t-1})} \ \min_{\boldsymbol{y}_t : \Psi_t(\boldsymbol{\phi}_t) \to \mathbb{R}^{n_t^2}} \ \max_{\boldsymbol{\psi}_t \in \Psi_t(\boldsymbol{\phi}_t)} \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)^\top \boldsymbol{x}_t^k + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t(\boldsymbol{\psi}_t) + \overline{\mathcal{Q}}_{t+1}(\boldsymbol{x}_t^k; \boldsymbol{\phi}_t) \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \hat{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t)\, \boldsymbol{x}_t^k + \boldsymbol{V}_t\, \boldsymbol{y}_t(\boldsymbol{\psi}_t) \geqslant \boldsymbol{h}_t(\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \\ & k \in \{1, \ldots, p_t\} \end{array} \right]$$

Under Approximation 2, the control variables $\boldsymbol{y}_t$ are no longer fully adaptive in $\boldsymbol{\psi}_t$. Instead, l affine decision rules $\boldsymbol{y}_{tj} : \Psi_t(\boldsymbol{\phi}_t) \to \mathbb{R}^{n_t^2}$ are chosen before $\boldsymbol{\psi}_t$ is observed, and the decision of the best decision rule is implemented upon observation of $\boldsymbol{\psi}_t$. One readily verifies that this restriction leads to the optimization problem in the statement of the observation. □

Proof of Proposition 4. Recall from the proof of Proposition 1 that $P(\mathcal{X}^{\mathrm{C}}, \Xi)$ evaluates to

$$\max_{\boldsymbol{\xi}_1 \in \Xi_1} \min_{(\boldsymbol{x}_1, \boldsymbol{y}_1) \in F_1(\boldsymbol{\xi}_1)} \Big[\boldsymbol{q}_1(\boldsymbol{\xi}_1)^\top \boldsymbol{x}_1 + \boldsymbol{r}_1{}^\top \boldsymbol{y}_1 \; + \max_{\boldsymbol{\xi}_2 \in \Xi_2(\boldsymbol{\phi}_1)} \min_{(\boldsymbol{x}_2, \boldsymbol{y}_2) \in F_2(\boldsymbol{x}_1, \boldsymbol{\xi}_2)} \Big[\boldsymbol{q}_2(\boldsymbol{\xi}_2)^\top \boldsymbol{x}_2 + \boldsymbol{r}_2{}^\top \boldsymbol{y}_2 + \cdots + \mathcal{Q}_\tau(\boldsymbol{x}_{\tau-1}; \boldsymbol{\phi}_{\tau-1}) \Big] \cdots \Big],$$

and thus a necessary condition for $P(\hat{\mathcal{X}}^{\mathrm{C}}, \Xi) < P(\mathcal{X}^{\mathrm{C}}, \Xi)$ is that the inclusion of $\boldsymbol{x}'_\tau$ in $\mathcal{X}^{\mathrm{C}}_t$ strictly decreases the optimal value of at least one of the two-stage subproblems $\mathcal{Q}_\tau(\boldsymbol{x}_{\tau-1}; \boldsymbol{\phi}_{\tau-1})$, $\boldsymbol{x}_{\tau-1} \in \mathcal{X}^{\mathrm{C}}_{\tau-1}$ and $\boldsymbol{\phi}_{\tau-1} \in \boldsymbol{\Phi}_{\tau-1}(\boldsymbol{\phi}_{\tau-2})$. Recall also that for $\hat{\boldsymbol{x}}_{\tau-1} \in \mathcal{X}^{\mathrm{C}}_{\tau-1}$ and $\boldsymbol{\phi}_{\tau-1} \in \boldsymbol{\Phi}_{\tau-1}(\boldsymbol{\phi}_{\tau-2})$, we have

$$\mathcal{Q}_\tau(\hat{\boldsymbol{x}}_{\tau-1}; \boldsymbol{\phi}_{\tau-1}) \;=\; \left[\begin{array}{ll} \text{maximize} & \left[\begin{array}{ll} \text{minimize} & \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}_\tau + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}_\tau + \mathcal{Q}_{\tau+1}(\boldsymbol{x}_\tau; \boldsymbol{\phi}_\tau) \\ \text{subject to} & \boldsymbol{T}_\tau(\boldsymbol{\xi}_\tau)\, \hat{\boldsymbol{x}}_{\tau-1} + \boldsymbol{W}_\tau(\boldsymbol{\xi}_\tau)\, \boldsymbol{x}_\tau + \boldsymbol{V}_\tau\, \boldsymbol{y}_\tau \geqslant \boldsymbol{h}_\tau(\boldsymbol{\xi}_\tau) \\ & \boldsymbol{x}_\tau \in \mathcal{X}^{\mathrm{C}}_\tau, \; \boldsymbol{y}_\tau \in \mathbb{R}^{n^2_\tau} \end{array} \right] \\ \text{subject to} & \boldsymbol{\xi}_\tau = (\boldsymbol{\phi}_\tau, \boldsymbol{\psi}_\tau) \in \Xi_\tau(\boldsymbol{\phi}_{\tau-1}) \end{array} \right].$$

A necessary condition for the optimal value of this problem to decrease when $\boldsymbol{x}'_\tau$ is included in the set $\mathcal{X}^{\mathrm{C}}_\tau$ is that there is $\boldsymbol{\phi}_{\tau-1} \in \Phi_{\tau-1}$, $\boldsymbol{\xi}_\tau \in \Xi_\tau(\boldsymbol{\phi}_{\tau-1})$ and $\boldsymbol{y}'_\tau \in \mathbb{R}^{n^2_\tau}$ such that $(\boldsymbol{x}'_\tau, \boldsymbol{y}'_\tau)$ is feasible in the inner minimization problem and improves its optimal value under the realization $\boldsymbol{\xi}_\tau$, which is exactly what the necessary condition in the statement of the proposition demands.

On the other hand, the sufficient condition in the statement of the proposition guarantees that

$$\begin{aligned} P(\mathcal{X}^{\mathrm{C}}, \Xi) \;&=\; \max_{\boldsymbol{\xi} \in \Xi} \left\{ \sum_{t=1}^{T} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}^\star_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}^\star_t(\boldsymbol{\xi}^t) \right\} \\ &=\; \max \left\{ \max_{\boldsymbol{\xi} \in \Xi^\star} \left\{ \sum_{t=1}^{\tau-1} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}^\star_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}^\star_t(\boldsymbol{\xi}^t) + \right.\right. \\ &\qquad\qquad \left. \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}^\star_\tau(\boldsymbol{\xi}^\tau) + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}^\star_\tau(\boldsymbol{\xi}^\tau) \; + \; \mathcal{Q}_{\tau+1}(\boldsymbol{x}^\star_\tau(\boldsymbol{\xi}^\tau); \boldsymbol{\phi}_\tau) \right\}, \\ &\qquad \max_{\boldsymbol{\xi} \in \Xi \setminus \Xi^\star} \left\{ \sum_{t=1}^{\tau-1} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}^\star_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}^\star_t(\boldsymbol{\xi}^t) + \right. \\ &\qquad\qquad \left.\left. \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}^\star_\tau(\boldsymbol{\xi}^\tau) + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}^\star_\tau(\boldsymbol{\xi}^\tau) \; + \; \mathcal{Q}_{\tau+1}(\boldsymbol{x}^\star_\tau(\boldsymbol{\xi}^\tau); \boldsymbol{\phi}_\tau) \right\} \right\} \\ &>\; \max \left\{ \max_{\boldsymbol{\xi} \in \Xi^\star} \left\{ \sum_{t=1}^{\tau-1} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}^\star_t(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}^\star_t(\boldsymbol{\xi}^t) + \right.\right. \\ &\qquad\qquad \left. \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}'_\tau + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}'_\tau(\boldsymbol{\xi}^t) \; + \; \mathcal{Q}_{\tau+1}(\boldsymbol{x}'_\tau; \boldsymbol{\phi}_\tau) \right\}, \end{aligned}$$

$$
\max_{\boldsymbol{\xi} \in \Xi \setminus \Xi^\star} \left\{ \sum_{t=1}^{\tau-1} \boldsymbol{q}_t(\boldsymbol{\xi}_t)^\top \boldsymbol{x}_t^\star(\boldsymbol{\xi}^t) + \boldsymbol{r}_t{}^\top \boldsymbol{y}_t^\star(\boldsymbol{\xi}^t) + \right.
$$

$$
\left. \left. \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}_\tau^\star(\boldsymbol{\xi}^\tau) + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}_\tau^\star(\boldsymbol{\xi}^\tau) \; + \; \mathcal{Q}_{\tau+1}(\boldsymbol{x}_\tau^\star(\boldsymbol{\xi}^\tau); \boldsymbol{\phi}_\tau) \right\} \right\}
$$

$$
\geqslant \; P(\hat{\mathcal{X}}^{\mathrm{C}}, \Xi),
$$

where $\Xi^\star \subseteq \Xi$ is the set of worst-case uncertainties corresponding to the policy $(\boldsymbol{x}^\star, \boldsymbol{y}^\star)$ in $P(\mathcal{X}^{\mathrm{C}}, \Xi)$. Here, the first two identities hold by definition. Note that second embedded maximization is strictly smaller than the first embedded maximization by definition of the set $\Xi^\star$. The first inequality is due to the sufficient condition in the statement of the proposition, which guarantees that for any $\boldsymbol{\xi} \in \Xi^\star$, there is a feasible decision pair $(\boldsymbol{x}'_\tau, \boldsymbol{y}'_\tau(\boldsymbol{\xi}^t))$ such that

$$
\boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}'_\tau + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}'_\tau(\boldsymbol{\xi}^\tau) + \mathcal{Q}_{\tau+1}(\boldsymbol{x}'_\tau; \boldsymbol{\phi}_\tau) \; < \; \boldsymbol{q}_\tau(\boldsymbol{\xi}_\tau)^\top \boldsymbol{x}_\tau^\star(\boldsymbol{\xi}^\tau) + \boldsymbol{r}_\tau{}^\top \boldsymbol{y}_\tau^\star(\boldsymbol{\xi}^\tau) + \mathcal{Q}_{\tau+1}(\boldsymbol{x}_\tau^\star(\boldsymbol{\xi}^\tau); \boldsymbol{\phi}_\tau),
$$

and thus the value of the first embedded maximization strictly decreases when we replace $\boldsymbol{x}_\tau^\star(\boldsymbol{\xi}^\tau)$ with $\boldsymbol{x}'_\tau$. The second inequality, finally, holds since the policy in the preceding line is feasible but not necessarily optimal in $P(\hat{\mathcal{X}}^{\mathrm{C}}, \Xi)$. $\square$