Electronic Companion

EC.1. Proof of Theorem 1

We first recall the Integer Programming (IP) Feasibility problem:

INTEGER PROGRAMMING FEASIBILITY.
Instance. Given are $\boldsymbol{F} \in \mathbb{Z}^{l \times k}$ and $\boldsymbol{g} \in \mathbb{Z}^{l}$.
Question. Is there a vector $\boldsymbol{z} \in \{0,1\}^k$ such that $\boldsymbol{F}\boldsymbol{z} \leq \boldsymbol{g}$?

The IP Feasibility problem is well-known to be strongly NP-hard (Garey and Johnson 1979).

PROOF OF THEOREM 1. We claim that the following problem can be cast as a stage-t worst-case cost problem $\mathcal{Q}_t(\boldsymbol{x}_{t-1})$ with an uncertain technology matrix $\boldsymbol{T}_t$ or uncertain right-hand sides $\boldsymbol{h}_t$:

$$\begin{array}{ll} \text{maximize} & \displaystyle\min_{\boldsymbol{x} \in \mathbb{R}^k} \left\{ \mathbf{e}^\top \boldsymbol{x} : \boldsymbol{x} \geq \boldsymbol{\xi},\ \boldsymbol{x} \geq \mathbf{e} - \boldsymbol{\xi} \right\} \\ \text{subject to} & \boldsymbol{F}\boldsymbol{\xi} \leq \boldsymbol{g} \\ & \boldsymbol{\xi} \in [0,1]^k \end{array} \tag{EC.1}$$

Indeed, for the stage-t worst-case cost problem $\mathcal{Q}_t(\boldsymbol{x}_{t-1})$ with uncertain technology matrix $\boldsymbol{T}_t$, we set $n_t = k_t = k$, $m_t = 2k$, $\Xi_t = \{\boldsymbol{\xi}_t \in [0,1]^{k_t} : \boldsymbol{F}\boldsymbol{\xi}_t \leq \boldsymbol{g}\}$ and

$$\boldsymbol{q}_t = \mathbf{e} \in \mathbb{R}^k, \quad \mathcal{Q}_{t+1}(\boldsymbol{x}_t) = 0 \ \forall \boldsymbol{x}_t \in \mathbb{R}^{n_t}, \quad \boldsymbol{T}_t(\boldsymbol{\xi}_t) = \begin{bmatrix} \operatorname{diag}(-\boldsymbol{\xi}_t) & \mathbf{0} \\ \mathbf{0} & \operatorname{diag}(\boldsymbol{\xi}_t) \end{bmatrix} \in \mathbb{R}^{2k \times 2k},$$

$$\boldsymbol{x}_{t-1} = \mathbf{e} \in \mathbb{R}^{2k}, \quad \boldsymbol{W}_t = \begin{bmatrix} \boldsymbol{I} \\ \boldsymbol{I} \end{bmatrix} \in \mathbb{R}^{2k \times k} \quad \text{and} \quad \boldsymbol{h}_t(\boldsymbol{\xi}_t) = \begin{bmatrix} \mathbf{0} \\ \mathbf{e} \end{bmatrix} \in \mathbb{R}^{2k},$$

while for the stage-t worst-case cost problem $\mathcal{Q}_t(\boldsymbol{x}_{t-1})$ with uncertain right-hand sides $\boldsymbol{h}_t$, we set $n_t = k_t = k$, $m_t = 2k$, $\Xi_t = \{\boldsymbol{\xi}_t \in [0,1]^{k_t} : \boldsymbol{F}\boldsymbol{\xi}_t \leq \boldsymbol{g}\}$ and

$$\boldsymbol{q}_t = \mathbf{e} \in \mathbb{R}^k, \quad \mathcal{Q}_{t+1}(\boldsymbol{x}_t) = 0 \ \forall \boldsymbol{x}_t \in \mathbb{R}^{n_t}, \quad \boldsymbol{T}_t(\boldsymbol{\xi}_t) = \mathbf{0} \in \mathbb{R}^{2k \times 2k},$$

$$\boldsymbol{x}_{t-1} = \mathbf{0} \in \mathbb{R}^{2k}, \quad \boldsymbol{W}_t = \begin{bmatrix} \boldsymbol{I} \\ \boldsymbol{I} \end{bmatrix} \in \mathbb{R}^{2k \times k} \quad \text{and} \quad \boldsymbol{h}_t(\boldsymbol{\xi}_t) = \begin{bmatrix} \boldsymbol{\xi}_t \\ \mathbf{e} - \boldsymbol{\xi}_t \end{bmatrix} \in \mathbb{R}^{2k}.$$

We now show that the answer to an instance $(\boldsymbol{F}, \boldsymbol{g})$ of the IP Feasibility problem is affirmative if and only if the optimal value of the corresponding instance of (EC.1) is k. Assume first that the IP Feasibility problem is solved by $\boldsymbol{z} \in \{0,1\}^k$. By construction, $\boldsymbol{z}$ is feasible in (EC.1) and attains the objective value k. At the same time, the optimal value of any instance of (EC.1) is at most k since $\boldsymbol{x} = \mathbf{e}$ is always feasible in the inner minimization problem of (EC.1). We thus conclude that the optimal value of (EC.1) is indeed k, as desired. Assume now that the optimal value of (EC.1) is k. In that case, there must be a feasible solution $\boldsymbol{\xi} \in [0,1]^k$ with $\boldsymbol{F}\boldsymbol{\xi} \leq \boldsymbol{g}$ for which the inner minimization problem is optimized by $\boldsymbol{x} = \mathbf{e}$. From the constraints of the minimization problem we then conclude that $\boldsymbol{\xi} \in \{0,1\}^k$, that is, $\boldsymbol{\xi}$ solves the IP Feasibility problem. $\square$

EC.2. Multi-Stage Robust Optimization without Relatively Complete Recourse

In this section we relax the assumption (A2) from the main paper, which stipulates that the multi-stage robust optimization problem (1) has a relatively complete recourse. We continue to assume that problem (1) satisfies the other three conditions (A1), (A3) and (A4).

Since problem (1) no longer satisfies (A2), the worst-case cost to-go functions $\mathcal{Q}_t$ become extended real-valued functions that attain the value $\mathcal{Q}_t(\boldsymbol{x}_{t-1}) = +\infty$ whenever the decision $\boldsymbol{x}_{t-1}$ cannot be completed to a feasible solution $\boldsymbol{x}_{t-1}, \ldots, \boldsymbol{x}_T$ for all subsequent realizations $\boldsymbol{\xi}_t, \ldots, \boldsymbol{\xi}_T$ of the uncertain parameters. While it turns out that the upper and lower bounds $\overline{\mathcal{Q}}_t$ and $\underline{\mathcal{Q}}_t$ from the main paper retain their bounding property in this more general setting, the lower bounds by construction never attain the value $+\infty$, and the algorithm may thus fail to identify an optimal solution. To address this issue, we will replace the lower bounds $\underline{\mathcal{Q}}_t$ with $\underline{\mathcal{Q}}_t + \delta_{\hat{\mathcal{X}}_t}$, where $\hat{\mathcal{X}}_t$ represents an outer approximation of the *stage-t feasible region* $\mathcal{X}_t = \{\boldsymbol{x}_t \in \mathbb{R}^{n_t} : \mathcal{Q}_{t+1}(\boldsymbol{x}_t) < +\infty\}$.

One readily verifies that the stage-t feasible region satisfies the recursion $\mathcal{X}_T = \mathbb{R}^{n_T}$ and

$$\mathcal{X}_{t-1} = \left\{ \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}} : \begin{bmatrix} \forall \boldsymbol{\xi}_t \in \Xi_t \;\; \exists \boldsymbol{x}_t \in \mathcal{X}_t : \\ \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_{t-1} + \boldsymbol{W}_t \boldsymbol{x}_t \geq \boldsymbol{h}_t(\boldsymbol{\xi}_t) \end{bmatrix} \right\}, \; t = T, \ldots, 2.$$

Thus, for a previously implemented stage-$(t-1)$ decision $\boldsymbol{x}_{t-1}$ and a parameter realization $\boldsymbol{\xi}_t$, a stage-t decision $\boldsymbol{x}_t \in \mathbb{R}^{n_t}$ is feasible in stage t *and* can be extended to a feasible solution $\boldsymbol{x}_t, \ldots, \boldsymbol{x}_T$ if and only if $\boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_{t-1} + \boldsymbol{W}_t \boldsymbol{x}_t \geq \boldsymbol{h}_t(\boldsymbol{\xi}_t)$ and $\boldsymbol{x}_t \in \mathcal{X}_t$. Likewise, a first-stage decision $\boldsymbol{x}_1 \in \mathbb{R}^{n_1}$ is feasible and extendable to a complete decision policy if and only if $\boldsymbol{W}_1 \boldsymbol{x}_1 \geq \boldsymbol{h}_1$ and $\boldsymbol{x}_1 \in \mathcal{X}_1$.

Similar to the stage-wise worst-case cost to-go functions $\mathcal{Q}_t$, the stage-wise feasible regions $\mathcal{X}_t$ are unknown for $t < T$. We therefore have to confine ourselves to outer approximations $\hat{\mathcal{X}}_t$ of $\mathcal{X}_t$ that are refined throughout the algorithm. To this end, we define the *stage-t feasibility problem* for a given outer approximation $\hat{\mathcal{X}}_t$ of $\mathcal{X}_t$ as

$$\begin{aligned} \overline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1}) \;=\; \text{maximize} \;\; & \begin{bmatrix} \text{minimize} & \mathbf{e}^\top \boldsymbol{v}_t \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\xi}_t)\, \boldsymbol{x}_{t-1} + \boldsymbol{W}_t \boldsymbol{x}_t + \boldsymbol{v}_t \geq \boldsymbol{h}_t(\boldsymbol{\xi}_t) \\ & \boldsymbol{x}_t \in \hat{\mathcal{X}}_t, \;\; \boldsymbol{v}_t \in \mathbb{R}^{m_t}_+ \end{bmatrix} \\ \text{subject to} \;\; & \boldsymbol{\xi}_t \in \Xi_t. \end{aligned}$$

The stage-t feasibility problem $\overline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1})$ closely resembles the upper bound problem $\overline{Q}_t(\boldsymbol{x}_{t-1})$, except that it is restricted to second-stage decisions $\boldsymbol{x}_t \in \hat{\mathcal{X}}_t$ and that its objective function minimizes the sum of the constraint violations incurred in $\overline{Q}_t(\boldsymbol{x}_{t-1})$. Since we only solve the problem when $\hat{\mathcal{X}}_t \neq \emptyset$, the inner minimization in $\overline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1})$ is guaranteed to be feasible for every realization of the uncertain parameters even if $\overline{Q}_t(\boldsymbol{x}_{t-1})$ is infeasible, which allows us to solve the problem via mixed-integer linear programming techniques or vertex enumeration, see Section 3.2.

Similar to the lower bound problem $\underline{Q}_t(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t)$, we will utilize the optimal solution to the second-stage problem in $\overline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1})$ for a fixed parameter realization $\boldsymbol{\xi}_t \in \Xi_t$, which we obtain from

$$\underline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t) \;=\; \begin{bmatrix} \text{minimize} & \mathbf{e}^\top \boldsymbol{v}_t \\ \text{subject to} & \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_{t-1} + \boldsymbol{W}_t \boldsymbol{x}_t + \boldsymbol{v}_t \geq \boldsymbol{h}_t(\boldsymbol{\xi}_t) \\ & \boldsymbol{x}_t \in \hat{\mathcal{X}}_t,\ \ \boldsymbol{v}_t \in \mathbb{R}_+^{m_t} \end{bmatrix}.$$

For a polyhedral outer approximation $\hat{\mathcal{X}}_t$, this problem is a linear program with the associated dual

$$\begin{aligned} \text{maximize} \quad & [\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_{t-1}]^\top \boldsymbol{\mu}_t - \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t) \\ \text{subject to} \quad & \boldsymbol{\mu}_t \in [0,1]^{m_t}, \end{aligned}$$

where $\sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{y}_t) = \sup\{\boldsymbol{x}_t^\top \boldsymbol{y}_t \,:\, \boldsymbol{x}_t \in \hat{\mathcal{X}}_t\}$ is the support function of $\hat{\mathcal{X}}_t$. Strong duality holds between $\underline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t)$ and its dual problem since $\underline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t)$ is feasible and $\boldsymbol{\mu}_t = \mathbf{0}$ is feasible in the dual.

We continue to use the upper and lower bound problems $\overline{Q}_t$ and $\underline{Q}_t$ from the main paper, but we replace the domains $\boldsymbol{x}_t \in \mathbb{R}^{n_t}$ in the lower bound problems with $\boldsymbol{x}_t \in \hat{\mathcal{X}}_t$. We will only solve the upper bound problems when $\overline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1}) = 0$, which implies that the inner minimization in the upper bound problems will be feasible for every realization of the uncertain parameters, and we can solve them via mixed-integer linear programming or vertex enumeration, see Section 3.2. The incorporation of $\hat{\mathcal{X}}_t$ in the lower bound problem $\underline{Q}_t(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t)$ leads to the new dual:

$$\begin{aligned} \text{maximize} \quad & [\boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_{t-1}]^\top \boldsymbol{\pi}_t - \underline{\mathcal{Q}}_{t+1}^\star(\boldsymbol{W}_t^\top \boldsymbol{\pi}_t - \boldsymbol{q}_t - \boldsymbol{y}_t) - \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{y}_t) \\ \text{subject to} \quad & \boldsymbol{\pi}_t \in \mathbb{R}_+^{m_t},\ \ \boldsymbol{y}_t \in \mathbb{R}^{n_t} \end{aligned}$$

In the following, strong duality will hold between $\underline{Q}_t(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t)$ and its dual since we only solve $\underline{Q}_t(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t)$ when $\underline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t) = 0$, in which case $\underline{Q}_t(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t)$ is feasible. Note that $\sigma_{\mathbb{R}^{n_t}}(\boldsymbol{y}_t) = 0$ if

$\boldsymbol{y}_t = \boldsymbol{0}$ and $\sigma_{\mathbb{R}^{n_t}}(\boldsymbol{y}_t) = +\infty$ otherwise. Thus, the new dual problem reduces to the dual problem from the main paper if we set $\hat{\mathcal{X}}_t = \mathbb{R}^{n_t}$ and $\boldsymbol{y}_t = \boldsymbol{0}$.

We now present our revised RDDP scheme for instances of (1) violating the assumption (A2):

1. **Initialization:** Set $\underline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}) = -\infty \ \forall \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}}$, $t = 2, \ldots, T$ (lower bound),
 $\overline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}) = +\infty \ \forall \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}}$, $t = 2, \ldots, T$ (upper bound),
 $\underline{\mathcal{Q}}_{T+1}(\boldsymbol{x}_T) = \overline{\mathcal{Q}}_{T+1}(\boldsymbol{x}_T) = 0 \ \forall \boldsymbol{x}_T \in \mathbb{R}^{n_T}$ (zero terminal costs),
 $\hat{\mathcal{X}}_t = \mathbb{R}^{n_t}$, $t = 1, \ldots, T$ (outer approximation of feasible region).
2. **Stage-1 Problem:** Let $\underline{\boldsymbol{x}}_1$ be an optimal solution to the lower bound stage-1 problem
$$\begin{array}{ll} \text{minimize} & \boldsymbol{q}_1^\top \boldsymbol{x}_1 + \underline{\mathcal{Q}}_2(\boldsymbol{x}_1) \\ \text{subject to} & \boldsymbol{W}_1 \boldsymbol{x}_1 \geq \boldsymbol{h}_1 \\ & \boldsymbol{x}_1 \in \hat{\mathcal{X}}_1. \end{array}$$
 If this problem is infeasible, terminate: problem (1) is infeasible. Otherwise, if $\underline{\mathcal{Q}}_2(\underline{\boldsymbol{x}}_1) = \overline{\mathcal{Q}}_2(\underline{\boldsymbol{x}}_1)$, terminate: $\underline{\boldsymbol{x}}_1$ is optimal with costs $\boldsymbol{q}_1^\top \underline{\boldsymbol{x}}_1 + \underline{\mathcal{Q}}_2(\underline{\boldsymbol{x}}_1)$. Else, set $t = 2$ and go to Step 3.
3. **Forward Pass:** Let $\boldsymbol{\xi}_t^{\mathrm{F}}$ be an optimal solution to the stage-t feasibility problem $\overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1})$.
 (a) If $\overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1}) > 0$, then let $\boldsymbol{\mu}_t$ be the shadow price of an optimal solution to the problem $\underline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t^{\mathrm{F}})$ and update the outer approximation of the stage-$(t-1)$ feasible region
$$\hat{\mathcal{X}}_{t-1} \leftarrow \hat{\mathcal{X}}_{t-1} \cap \left\{ \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}} : [\boldsymbol{h}_t(\boldsymbol{\xi}_t^{\mathrm{F}}) - \boldsymbol{T}_t(\boldsymbol{\xi}_t^{\mathrm{F}})\, \boldsymbol{x}_{t-1}]^\top \boldsymbol{\mu}_t - \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t) \leq 0 \right\}.$$
 If $\hat{\mathcal{X}}_{t-1} = \emptyset$, terminate: problem (1) is infeasible. Otherwise, repeat Step 3 for stage $t-1$ (if $t > 2$) or go back to Step 2 (if $t = 2$).
 (b) If $\overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1}) = 0$, then let $\overline{\boldsymbol{\xi}}_t^{\mathrm{fw}}$ be optimal in $\overline{Q}_t(\underline{\boldsymbol{x}}_{t-1})$, let $\underline{\boldsymbol{x}}_t$ be optimal in $\underline{Q}_t(\underline{\boldsymbol{x}}_{t-1}; \overline{\boldsymbol{\xi}}_t^{\mathrm{fw}})$, and repeat Step 3 for stage $t+1$ (if $t < T-1$) or go to Step 4 (if $t = T-1$).
4. **Backward Pass:** For $t = T, \ldots, 2$, let $\overline{\boldsymbol{\xi}}_t^{\mathrm{bw}}$ be an optimal solution to the upper bound problem $\overline{Q}_t(\underline{\boldsymbol{x}}_{t-1})$. If $\overline{Q}_t(\underline{\boldsymbol{x}}_{t-1}) < \overline{\mathcal{Q}}_t(\underline{\boldsymbol{x}}_{t-1})$, then update the upper bound
$$\overline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}) \leftarrow \operatorname{env}\left(\min\left\{\overline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}),\, \overline{Q}_t(\underline{\boldsymbol{x}}_{t-1}) + \delta_{\{\underline{\boldsymbol{x}}_{t-1}\}}(\boldsymbol{x}_{t-1})\right\}\right) \qquad \forall \boldsymbol{x}_{t-1} \in \mathbb{R}_{n_{t-1}}.$$
 Let $(\boldsymbol{\pi}_t, \boldsymbol{y}_t)$ be the shadow price of an optimal solution to the lower bound problem $\underline{Q}_t(\underline{\boldsymbol{x}}_{t-1}; \overline{\boldsymbol{\xi}}_t^{\mathrm{bw}})$. If $\underline{Q}_t(\underline{\boldsymbol{x}}_{t-1}; \overline{\boldsymbol{\xi}}_t^{\mathrm{bw}}) > \underline{\mathcal{Q}}_t(\underline{\boldsymbol{x}}_{t-1})$, then update the lower bound
$$\begin{aligned} \underline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}) \leftarrow \max\Big\{ & \underline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}), [\boldsymbol{h}_t(\overline{\boldsymbol{\xi}}_t^{\mathrm{bw}}) - \boldsymbol{T}_t(\overline{\boldsymbol{\xi}}_t^{\mathrm{bw}})\, \boldsymbol{x}_{t-1}]^\top \boldsymbol{\pi}_t - \\ & \underline{\mathcal{Q}}_{t+1}^{\star}(\boldsymbol{W}_t^\top \boldsymbol{\pi}_t - \boldsymbol{q}_t - \boldsymbol{y}_t) - \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{y}_t) \Big\} \ \forall \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}}. \end{aligned} \tag{EC.2}$$

After Step 4 has been completed for all stages $t = T, \ldots, 2$, go back to Step 2.

The algorithm largely resembles the RDDP scheme from the main paper. The key difference lies in the forward pass, which now iteratively constructs outer approximations $\hat{\mathcal{X}}_t$ of the stage-t feasible regions $\mathcal{X}_t$. To this end, we ensure that in each stage $t = 2, \ldots, T$, the selected decision $\underline{\boldsymbol{x}}_{t-1}$ can be extended to a decision $\boldsymbol{x}_t \in \hat{\mathcal{X}}_t$ satisfying $\boldsymbol{T}_t(\boldsymbol{\xi}_t)\underline{\boldsymbol{x}}_{t-1} + \boldsymbol{W}_t \boldsymbol{x}_t \geq \boldsymbol{h}_t(\boldsymbol{\xi}_t)$ under each possible parameter realization $\boldsymbol{\xi}_t \in \Xi_t$. If this is not possible, which is the case if and only if the feasibility problem $\overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1})$ does not evaluate to zero, then the previously selected decision $\underline{\boldsymbol{x}}_{t-1}$ cannot be completed to a nonanticipative decision policy under the current approximation of the feasible region. In that case, either $\hat{\mathcal{X}}_t = \emptyset$ or the approximation $\hat{\mathcal{X}}_{t-1}$ is updated through a feasibility cut. If $\hat{\mathcal{X}}_t = \emptyset$, then the problem must be infeasible. Otherwise, the introduced feasibility cut ensures that $\underline{\boldsymbol{x}}_{t-1}$ is no longer contained in $\hat{\mathcal{X}}_{t-1}$, and the forward pass is repeated for stage $t-1$.

We now analyze the convergence of the revised RDDP scheme. The algorithm starts with the trivial outer approximations $\hat{\mathcal{X}}_t = \mathbb{R}^{n_t}$ of the stage-wise feasible regions $\mathcal{X}_t$. These outer approximations are refined in the forward passes in Step 3.

LEMMA EC.1. *Each time the algorithm reaches Step 4, we have $\overline{Q}_t^F(\underline{\boldsymbol{x}}_{t-1}) = 0$ for all $t = 2, \ldots, T$.*

PROOF. By construction of the forward pass, we have $\overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1}) = 0$ for every $t = 2, \ldots, T$ *at some stage* in Step 3. We show that for every $t = 2, \ldots, T$, we also have $\overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1}) = 0$ *at the end* of Step 3. Assume to the contrary that $\overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1}) > 0$ for some $t = 2, \ldots, T$ at the end of Step 3. In that case, the outer approximation $\hat{\mathcal{X}}_t$ of the stage-t feasible region $\mathcal{X}_t$ must have changed after $\overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1})$ has been solved the last time. This is not possible, however, since $\hat{\mathcal{X}}_t$ is changed only when $\overline{Q}_{t+1}^{\mathrm{F}}(\underline{\boldsymbol{x}}_t) > 0$, in which case the problem $\overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1})$ is solved again. □

Lemma EC.1 ensures that the upper bound problems $\overline{Q}_t(\underline{\boldsymbol{x}}_{t-1})$ in the forward and backward passes are only solved when their inner minimization problems are feasible for every realization of the uncertain parameters, and we can thus continue to solve these problems via mixed-integer linear programming or vertex enumeration, see Section 3.2. The lemma also ensures that strong duality holds between the lower bound problems $\underline{Q}_t(\underline{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t)$ and their duals since the problems $\underline{Q}_t(\underline{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t)$ are only solved when they are feasible.

We now show that the bounding property of $\underline{\mathcal{Q}}_t$ and $\overline{\mathcal{Q}}_t$ is preserved in our revised RDDP scheme.

PROPOSITION EC.1 **(Bounding Property)**. *Throughout the algorithm, the bounds $\underline{\mathcal{Q}}_t$ and $\overline{\mathcal{Q}}_t$ as well as the outer approximations $\hat{\mathcal{X}}_t$ of the stage-wise feasible regions $\mathcal{X}_t$ satisfy*

$$\underline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}) + \delta_{\hat{\mathcal{X}}_{t-1}}(\boldsymbol{x}_{t-1}) \leq \mathcal{Q}_t(\boldsymbol{x}_{t-1}) \leq \overline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}) \qquad \forall \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}},\ \forall t = 2, \ldots, T,$$

that is, $\underline{\mathcal{Q}}_t + \delta_{\hat{\mathcal{X}}_{t-1}}$ and $\overline{\mathcal{Q}}_t$ bound $\mathcal{Q}_t$ from below and above, respectively.

PROOF. To prove that $\underline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}) + \delta_{\hat{\mathcal{X}}_{t-1}}(\boldsymbol{x}_{t-1}) \leq \mathcal{Q}_t(\boldsymbol{x}_{t-1})$ for all $\boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}}$ and $t = 2, \ldots, T$, we show that *(i)* $\delta_{\hat{\mathcal{X}}_{t-1}}(\boldsymbol{x}_{t-1}) = +\infty$ only when $\mathcal{Q}_t(\boldsymbol{x}_{t-1}) = +\infty$ and *(ii)* $\underline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}) \leq \mathcal{Q}_t(\boldsymbol{x}_{t-1})$. To prove the first statement, we show that $\mathcal{X}_t \subseteq \hat{\mathcal{X}}_t$ for all $t = 1, \ldots, T$ by backward induction on the time stage t. Indeed, we have $\mathcal{X}_T = \hat{\mathcal{X}}_T$ by construction. To show that $\mathcal{X}_t \subseteq \hat{\mathcal{X}}_t$ implies $\mathcal{X}_{t-1} \subseteq \hat{\mathcal{X}}_{t-1}$, we show that $\boldsymbol{x}_{t-1} \in \mathcal{X}_{t-1}$ satisfies every feasibility cut $\{\boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}} : [\boldsymbol{h}_t(\boldsymbol{\xi}_t^{\mathrm{F}}) - \boldsymbol{T}_t(\boldsymbol{\xi}_t^{\mathrm{F}})\,\boldsymbol{x}_{t-1}]^\top \boldsymbol{\mu}_t - \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t) \leq 0\}$ that is added to $\hat{\mathcal{X}}_{t-1}$ in Step 3 in some iteration. Indeed, we have

$$\begin{aligned}
\mathcal{X}_{t-1} &= \left\{ \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}} : \max_{\boldsymbol{\xi}_t \in \Xi_t} \min_{\boldsymbol{x}_t, \boldsymbol{v}_t} \left\{ \mathbf{e}^\top \boldsymbol{v}_t : \begin{bmatrix} \boldsymbol{T}_t(\boldsymbol{\xi}_t)\,\boldsymbol{x}_{t-1} + \boldsymbol{W}_t \boldsymbol{x}_t + \boldsymbol{v}_t \geq \boldsymbol{h}_t(\boldsymbol{\xi}_t) \\ \boldsymbol{x}_t \in \mathcal{X}_t,\ \ \boldsymbol{v}_t \in \mathbb{R}^{m_t}_+ \end{bmatrix} \right\} \leq 0 \right\} \\
&\subseteq \left\{ \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}} : \underline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1}; \boldsymbol{\xi}_t^{\mathrm{F}}) \leq 0 \right\} \\
&\subseteq \left\{ \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}} : \max_{\boldsymbol{\mu}_t \in [0,1]^{m_t}} \left\{ [\boldsymbol{h}_t(\boldsymbol{\xi}_t^{\mathrm{F}}) - \boldsymbol{T}_t(\boldsymbol{\xi}_t^{\mathrm{F}})\,\boldsymbol{x}_{t-1}]^\top \boldsymbol{\mu}_t - \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t) \right\} \leq 0 \right\} \\
&\subseteq \left\{ \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}} : [\boldsymbol{h}_t(\boldsymbol{\xi}_t^{\mathrm{F}}) - \boldsymbol{T}_t(\boldsymbol{\xi}_t^{\mathrm{F}})\,\boldsymbol{x}_{t-1}]^\top \boldsymbol{\mu}_t - \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t) \leq 0 \right\}.
\end{aligned}$$

Here, the first identity follows from the definition of $\mathcal{X}_{t-1}$ and $\mathcal{X}_t$. The first inclusion holds since $\underline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1}; \boldsymbol{\xi}_t^{\mathrm{F}})$ is equivalent to a variant of the max-min problem in the first row where we restrict Ξ_t to $\{\boldsymbol{\xi}_t^{\mathrm{F}}\}$ and replace $\mathcal{X}_t$ with $\hat{\mathcal{X}}_t$, which by the induction hypothesis contains $\mathcal{X}_t$. The second inclusion follows from weak linear programming duality, and the final inclusion holds since we restrict the feasible region $[0,1]^{m_t}$ of the maximization problem in the penultimate row to $\{\boldsymbol{\mu}_t\}$. Thus, $\boldsymbol{x}_{t-1} \in \mathcal{X}_{t-1}$ satisfies every feasibility cut introduced in Step 3, that is, $\mathcal{X}_{t-1} \subseteq \hat{\mathcal{X}}_{t-1}$.

The proof that $\underline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}) \leq \mathcal{Q}_t(\boldsymbol{x}_{t-1})$ and $\mathcal{Q}_t(\boldsymbol{x}_{t-1}) \leq \overline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1})$ directly follows from the proof of Proposition 2, which remains valid for our revised RDDP scheme due to Lemma EC.1. □

In analogy to Section 3.3, Proposition EC.1 implies that the stage-1 problem is a relaxation of problem (1'). Thus, infeasibility of the stage-1 problem implies infeasibility of problem (1'), and any solution $\underline{\boldsymbol{x}}_1$ satisfying $\underline{\mathcal{Q}}_2(\underline{\boldsymbol{x}}_1) = \overline{\mathcal{Q}}_2(\underline{\boldsymbol{x}}_1)$ must also satisfy $\underline{\mathcal{Q}}_2(\underline{\boldsymbol{x}}_1) = \mathcal{Q}_2(\underline{\boldsymbol{x}}_1)$ and therefore be optimal in problem (1'). We have thus arrived at the following result.

Corollary EC.1 **(Correctness)**. *If the RDDP scheme terminates, then it either returns an optimal solution to problem* (1) *or it correctly identifies infeasibility of the problem.*

To show that our revised RDDP scheme converges in finite time, we again assume that the conditions (C1) and (C2) from the main paper hold, and we make two additional assumptions.

(C3) The solution method for the stage-wise feasibility problems $\overline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1})$ returns an optimal extreme point of the uncertainty set Ξ_t.

(C4) The solution method for the stage-wise feasibility problems $\underline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t)$ returns an optimal basic feasible solution of the linear program corresponding to the epigraph reformulation of $\underline{Q}_t^{\mathrm{F}}(\boldsymbol{x}_{t-1};\boldsymbol{\xi}_t)$.

The assumptions closely resemble the conditions (C1) and (C2) from the main paper.

We first show that there are finitely many different outer approximations $\hat{\mathcal{X}}_t$ of each feasible region $\mathcal{X}_t$ that our revised RDDP scheme can generate during its execution.

Lemma EC.2 **(Outer Approximations)**. *For a fixed instance of problem* (1)*, there are finitely many different outer approximations $\hat{\mathcal{X}}_t$, $t=1,\ldots,T$, that the algorithm can generate.*

Proof. As in the proof of Lemma 1, we employ a backward induction on the time stage t. For $t=T$, the dual to the stage-T feasibility problem $\underline{Q}_T^{\mathrm{F}}(\underline{\boldsymbol{x}}_{T-1};\boldsymbol{\xi}_T^{\mathrm{F}})$ is

$$\begin{array}{ll} \text{maximize} & [\boldsymbol{h}_T(\boldsymbol{\xi}_T^{\mathrm{F}}) - \boldsymbol{T}_T(\boldsymbol{\xi}_T^{\mathrm{F}})\,\underline{\boldsymbol{x}}_{T-1}]^\top \boldsymbol{\mu}_T - \sigma_{\hat{\mathcal{X}}_T}(\boldsymbol{W}_T^\top \boldsymbol{\mu}_T) \\ \text{subject to} & \boldsymbol{\mu}_T \in [0,1]^{m_T}. \end{array}$$

Since $\hat{\mathcal{X}}_T = \mathcal{X}_T = \mathbb{R}^{n_T}$ throughout the algorithm, we have $\sigma_{\hat{\mathcal{X}}_T}(\boldsymbol{W}_T^\top \boldsymbol{\mu}_T) = 0$ if $\boldsymbol{W}_T^\top \boldsymbol{\mu}_T = \boldsymbol{0}$ and $\sigma_{\hat{\mathcal{X}}_T}(\boldsymbol{W}_T^\top \boldsymbol{\mu}_T) = +\infty$ otherwise. Thus, the dual problem can be rewritten as

$$\begin{array}{ll} \text{maximize} & [\boldsymbol{h}_T(\boldsymbol{\xi}_T^{\mathrm{F}}) - \boldsymbol{T}_T(\boldsymbol{\xi}_T^{\mathrm{F}})\,\underline{\boldsymbol{x}}_{T-1}]^\top \boldsymbol{\mu}_T \\ \text{subject to} & \boldsymbol{W}_T^\top \boldsymbol{\mu}_T = \boldsymbol{0} \\ & \boldsymbol{\mu}_T \in [0,1]^{m_T}. \end{array}$$

This is a linear program with finitely many basic feasible solutions, none of which depends on the candidate solution $\underline{\boldsymbol{x}}_{T-1}$ or the parameter realization $\boldsymbol{\xi}_T^{\mathrm{F}}$. Since the assumption (C4) ensures that the feasibility cut determined in Step 3 of the algorithm corresponds to a basic feasible solution of this dual problem, we thus conclude that for every $\boldsymbol{\xi}_T^{\mathrm{F}}$ there are finitely many different feasibility cuts that can be added to $\hat{\mathcal{X}}_{T-1}$ throughout the execution of the algorithm. The statement then follows for stage T since there are finitely many extreme points $\boldsymbol{\xi}_T^{\mathrm{F}} \in \Xi_T$ that can emerge as optimal solutions to the stage-T feasibility problem $\overline{Q}_T^{\mathrm{F}}(\underline{\boldsymbol{x}}_{T-1})$ according to the condition (C3).

For $t < T$, the dual to the stage-t feasibility problem $\underline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1}; \boldsymbol{\xi}_t^{\mathrm{F}})$ can be written as

$$\begin{aligned} \text{maximize} \quad & [\boldsymbol{h}_t(\boldsymbol{\xi}_t^{\mathrm{F}}) - \boldsymbol{T}_t(\boldsymbol{\xi}_t^{\mathrm{F}})\, \underline{\boldsymbol{x}}_{t-1}]^\top \boldsymbol{\mu}_t - \varphi \\ \text{subject to} \quad & \varphi \geq \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t) \\ & \boldsymbol{\mu}_t \in [0,1]^{m_t}, \end{aligned} \tag{EC.3}$$

where $\sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t)$ is a support function over the outer approximation $\hat{\mathcal{X}}_t$ of the stage-t feasible set $\mathcal{X}_t$. The set $\hat{\mathcal{X}}_t$ is polyhedral by construction, and the induction hypothesis implies that the set is described by finitely many halfspaces and that it takes on finitely many different shapes throughout the algorithm. Hence, the support function $\sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t)$ is a piecewise affine convex function with finitely many different pieces that takes on finitely many different shapes throughout the algorithm. Problem (EC.3) can then be cast as a linear program with finitely many basic feasible solutions, none of which depends on $\underline{\boldsymbol{x}}_{t-1}$ or $\boldsymbol{\xi}_t^{\mathrm{F}}$. A similar argument as in the previous paragraph then shows that there are finitely many different feasibility cuts that can be added to the outer approximation $\hat{\mathcal{X}}_{t-1}$ throughout the algorithm. □

We are now in the position to prove the finite convergence of our revised RDDP scheme.

THEOREM EC.1 **(Finite Termination)**. *The revised RDDP scheme terminates in finite time.*

PROOF. We show that the forward pass in each iteration completes in finite time. The statement of the theorem then follows from the fact that the proofs of Lemmas 1 and 2 as well as Theorem 2 remain valid for the revised RDDP scheme due to Lemma EC.1.

Assume that the forward pass of an iteration would cycle indefinitely. In that case, there would be at least one stage $(t-1) \in \{1, \ldots, T-1\}$ such that the outer approximation $\hat{\mathcal{X}}_{t-1}$ of the stage-$(t-1)$ feasible region $\mathcal{X}_{t-1}$ is updated infinitely many times in Step 3. Fix any such update

$$\hat{\mathcal{X}}_{t-1} \leftarrow \hat{\mathcal{X}}_{t-1} \cap \left\{ \boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}} : [\boldsymbol{h}_t(\boldsymbol{\xi}_t^{\mathrm{F}}) - \boldsymbol{T}_t(\boldsymbol{\xi}_t^{\mathrm{F}})\, \boldsymbol{x}_{t-1}]^\top \boldsymbol{\mu}_t - \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t) \leq 0 \right\},$$

where $\boldsymbol{\mu}_t$ is the shadow price of an optimal solution to the problem $\underline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1};\boldsymbol{\xi}_t^{\mathrm{F}})$. Note that $\underline{\boldsymbol{x}}_{t-1} \in \hat{\mathcal{X}}_{t-1}$ before the update since $\underline{\boldsymbol{x}}_{t-1}$ minimizes $\underline{Q}_{t-1}(\underline{\boldsymbol{x}}_{t-2};\overline{\boldsymbol{\xi}}_{t-1}^{\mathrm{fw}})$. By construction, we have that

$$[\boldsymbol{h}_t(\boldsymbol{\xi}_t^{\mathrm{F}}) - \boldsymbol{T}_t(\boldsymbol{\xi}_t^{\mathrm{F}})\,\underline{\boldsymbol{x}}_{t-1}]^\top \boldsymbol{\mu}_t - \sigma_{\hat{\mathcal{X}}_t}(\boldsymbol{W}_t^\top \boldsymbol{\mu}_t) \;=\; \underline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1};\boldsymbol{\xi}_t^{\mathrm{F}}) \;=\; \overline{Q}_t^{\mathrm{F}}(\underline{\boldsymbol{x}}_{t-1}) \;>\; 0,$$

that is, the updated set $\hat{\mathcal{X}}_{t-1}$ no longer contains $\underline{\boldsymbol{x}}_{t-1}$. We thus conclude that each update would have to result in a different outer approximation $\hat{\mathcal{X}}_{t-1}$, which contradicts Lemma EC.2. $\square$

EC.3. Uncertain Objective Functions and Recourse Matrices

This section relaxes the assumption (A4) from the main paper, which stipulates that the objective coefficients and the recourse matrices of the multi-stage robust optimization problem (1) are deterministic. We also allow for generic (possibly non-convex) compact stage-wise uncertainty sets Ξ_t as long as the rectangularity condition $\Xi = \times_{t=1}^{T} \Xi_t$ remains satisfied. To simplify the proofs, we continue to assume that problem (1) satisfies the other two conditions (A1) and (A2).

For ease of exposition, we consider the following, slightly modified variant of problem (1):

$$\begin{array}{lll} \text{minimize} & \displaystyle\max_{\boldsymbol{\xi} \in \Xi} \sum_{t=1}^{T} \boldsymbol{q}_t^\top \boldsymbol{x}_t(\boldsymbol{\xi}^t) & \\ \text{subject to} & \boldsymbol{f}_1(\boldsymbol{x}_1) \leq \boldsymbol{0} & \forall \boldsymbol{\xi} \in \Xi \\ & \boldsymbol{f}_t(\boldsymbol{x}_{t-1}(\boldsymbol{\xi}^{t-1}), \boldsymbol{\xi}_t, \boldsymbol{x}_t(\boldsymbol{\xi}^t)) \leq \boldsymbol{0} & \forall \boldsymbol{\xi} \in \Xi,\ \forall t = 2, \ldots, T \\ & \boldsymbol{x}_t(\boldsymbol{\xi}^t) \in \mathbb{R}^{n_t},\ \boldsymbol{\xi} \in \Xi \text{ and } t = 1, \ldots, T, & \end{array} \tag{EC.4}$$

where $\boldsymbol{f}_1 : \mathbb{R}^{n_1} \mapsto \mathbb{R}^{m_1}$ and $\boldsymbol{f}_t : \mathbb{R}^{n_{t-1}} \times \mathbb{R}^{k_t} \times \mathbb{R}^{n_t} \mapsto \mathbb{R}^{m_t}$, $t = 2, \ldots, T$, are generic functions that describe the state transitions between two successive time stages.

We change the assumptions (A3) and (A4) from the main paper as follows:

(A3') Uncertainty Set. The uncertainty set Ξ is compact and stage-wise rectangular, that is, $\Xi = \times_{t=1}^{T} \Xi_t$ for compact (but possibly non-convex) sets $\Xi_t \subseteq \mathbb{R}^{k_t}$.

(A4') Uncertain Parameters. The transition functions $\boldsymbol{f}_t(\cdot, \boldsymbol{\xi}_t, \cdot)$, $t = 2, \ldots, T$, are jointly quasi-convex in their first and last arguments for every $t = 2, \ldots, T$ and $\boldsymbol{\xi}_t \in \Xi_t$. Moreover, for all $t = 2, \ldots, T$, $\boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}}$ and $\boldsymbol{\xi}_t \in \Xi_t$ there exists $\boldsymbol{x}_t \in \mathbb{R}^{n_t}$ such that $\boldsymbol{f}_t(\boldsymbol{x}_{t-1}, \boldsymbol{\xi}_t, \boldsymbol{x}_t) < \boldsymbol{0}$.

The assumption (A3') admits, among others, ellipsoidal and semidefinite uncertainty sets, as well as uncertainty sets involving discrete parameters. Likewise, the assumption (A4') allows for a broad range of linear and nonlinear state transitions. In particular, the choice $\boldsymbol{f}_t(\boldsymbol{x}_{t-1}, \boldsymbol{\xi}_t, \boldsymbol{x}_t) = \boldsymbol{h}_t(\boldsymbol{\xi}_t) - \boldsymbol{T}_t(\boldsymbol{\xi}_t)\boldsymbol{x}_{t-1} - \boldsymbol{W}_t(\boldsymbol{\xi}_t)\boldsymbol{x}_t$ recovers instances of problem (1) with uncertain recourse matrices. Using an epigraph reformulation, problem (EC.4) can also accommodate uncertain objective coefficients. The Slater conditions ensure that strong duality holds for the lower bound problems $\underline{\mathcal{Q}}_t(\boldsymbol{x}_{t-1}; \boldsymbol{\xi}_t)$, whose associated dual problems provide the cuts for the lower bounds $\underline{\mathcal{Q}}_t$ in the RDDP scheme.

We apply the following two modifications to our RDDP scheme from the main paper:

(i) We replace $\pm\infty$ in the initial bounds $\underline{\mathcal{Q}}_t$, $\overline{\mathcal{Q}}_t$ with a sufficiently large number $\pm M$.

(ii) We update the lower bounds $\underline{\mathcal{Q}}_t$ in both the forward and the backward passes.

The first modification is used as a technical device to ensure uniform convergence of the cost to-go approximations $\underline{\mathcal{Q}}_t$ and $\overline{\mathcal{Q}}_t$. Note that a finite bound M is guaranteed to exist since the feasible region and the stage-wise uncertainty sets Ξ_t are bounded. The second modification simplifies our convergence argument, but it can also be motivated by the following remark.

Remark EC.1 (Lower Bound Updates in Forward Pass). Our RDDP scheme updates the lower and upper worst-case cost to-go approximations $\underline{\mathcal{Q}}_t$ and $\overline{\mathcal{Q}}_t$ in the backward passes only. This is justified for the upper bounds $\overline{\mathcal{Q}}_t$ since the optimal values of the upper bound problems $\overline{Q}_t(\underline{\boldsymbol{x}}_{t-1})$ in the backward passes are guaranteed to be at most as large as those in the preceding forward passes. However, since $\overline{\boldsymbol{\xi}}_t^{\mathrm{fw}} \neq \overline{\boldsymbol{\xi}}_t^{\mathrm{bw}}$ in general, the supporting hyperplanes corresponding to the optimal solutions of $\underline{Q}_t(\underline{\boldsymbol{x}}_{t-1};\overline{\boldsymbol{\xi}}_t^{\mathrm{fw}})$ in the forward pass may still provide useful information after the supporting hyperplanes of the backward pass have been introduced. It may thus be beneficial to refine the lower cost to-go approximations $\underline{\mathcal{Q}}_t$ in both the forward and the backward passes.

Standard results from convex analysis, such as Propositions 2.9 and 2.22 from Rockafellar and Wets (2009), show that under the assumptions (A3') and (A4'), the stage-wise worst-case cost to-go functions $\mathcal{Q}_t$, $t=2,\ldots,T$, remain convex. However, the revised assumption (A4') implies that the functions $\mathcal{Q}_t$ may no longer be piecewise affine, even if the stage-wise uncertainty sets Ξ_t are polyhedral and all transition functions $\boldsymbol{f}_t$ are linear in $\boldsymbol{\xi}_t$ and jointly linear in $\boldsymbol{x}_{t-1}$ and $\boldsymbol{x}_t$. To illustrate this, consider the following instance of problem (EC.4):

$$
\begin{array}{lllll}
\text{minimize} & \max\limits_{\boldsymbol{\xi}\in\Xi} \; -2x_1(\xi^1) + x_{21}(\boldsymbol{\xi}^2) + x_{22}(\boldsymbol{\xi}^2) & & & \\
\text{subject to} & x_{21}(\boldsymbol{\xi}^2) = -\dfrac{2}{3}\xi_2 x'_{21}(\boldsymbol{\xi}^2), & x'_{21}(\boldsymbol{\xi}^2) = \xi_2 x''_{21}(\boldsymbol{\xi}^2), & x''_{21}(\boldsymbol{\xi}^2) = \xi_2 & \forall \boldsymbol{\xi}\in\Xi \\
& x_{22}(\boldsymbol{\xi}^2) = \xi_2 x'_{22}(\boldsymbol{\xi}^2), & x'_{22}(\boldsymbol{\xi}^2) = \xi_2 x_1(\xi^1) & & \forall \boldsymbol{\xi}\in\Xi \\
& x_1(\xi^1)\in[0,3],\; x_{21}(\boldsymbol{\xi}^2), x'_{21}(\boldsymbol{\xi}^2), x''_{21}(\boldsymbol{\xi}^2), x_{22}(\boldsymbol{\xi}^2), x'_{22}(\boldsymbol{\xi}^2)\in\mathbb{R},\; \boldsymbol{\xi}\in\Xi\,, & & &
\end{array}
\tag{EC.5}
$$

where the uncertainty set is $\Xi = \Xi_1 \times \Xi_2 = \{1\}\times[0,3]$. Problem (EC.5) is an instance of a linear multi-stage robust optimization problem with uncertain recourse matrices. In fact, if the recourse matrices in problem (EC.5) were deterministic, then the functions $\mathcal{Q}_t$ would be piecewise affine, and Theorems 2 and EC.1 would imply that the problem could be solved in finite time by our RDDP scheme. Straightforward variable substitutions reveal that the problem is equivalent to

$$\min_{x_1 \in [0,3]} \max_{\xi \in [0,3]} -\frac{2}{3}\xi^3 + \xi^2 x_1 - 2x_1,$$

and $-\frac{2}{3}\xi^3 + \xi^2 x_1 - 2x_1$ describes the tangent of the nonlinear function $\frac{1}{3}x_1^3 - 2x_1$ at $\xi \in [0,3]$. The cost to-go function of problem (EC.5) thus satisfies $\mathcal{Q}_2(x_1) = \frac{1}{3}x_1^3 - 2x_1$, and it is minimized at $x_1^\star$ satisfying the first-order condition $(x_1^\star)^2 - 2 = 0$, that is, at $x_1^\star = \sqrt{2}$.

Consider now a solution approach for the generic multi-stage robust optimization problem (EC.4) that operates on piecewise affine approximations of the cost to-go function $\mathcal{Q}_2$ in problem (EC.5). Assume that the slopes or the breakpoints of these approximations are derived from primal or dual linear programming formulations of the second-stage problem in (EC.5). These linear programs can be solved exactly both by the simplex algorithm and by interior point methods (by rounding to the nearest extreme point, which can be done in polynomial time for linear programs). All basic feasible solutions of these linear programs have a bit-length that is bounded by a polynomial of the size of the input data (Limongelli and Pirastu 1994). Consequently, all breakpoints of the cost to-go approximations have a polynomial bit-length as well. However, since the solution to problem (EC.5) is an irrational number, we cannot expect it to emerge as a breakpoint of a piecewise affine cost to-go approximation after finitely many iterations. Assuming that the solution approach selects breakpoints of the cost to-go approximations as candidate solutions, the solution approach thus cannot converge to the optimal solution of the problem (EC.5) in finitely many iterations.

We now show that our RDDP scheme asymptotically converges to an optimal solution of the generic multi-stage robust optimization problem (EC.4). Our convergence proof does not require the assumptions (C1) and (C2) from the main paper. For ease of exposition, however, we assume that all feasible stage-t solutions $\boldsymbol{x}_t$ are contained in a bounded set $\mathcal{X}_t$, and that there is a compact subset $\mathcal{X}_t^\infty \subseteq \operatorname{rel\,int} \mathcal{X}_t$ such that every accumulation point $\underline{\boldsymbol{x}}_t^\infty$ of our RDDP scheme satisfies $\underline{\boldsymbol{x}}_t^\infty \in \mathcal{X}_t^\infty$. The existence of the sets $\mathcal{X}_t$ is guaranteed by the assumption (A1). The stipulated existence of the sets $\mathcal{X}_t^\infty$ can be relaxed at the expense of further case distinctions in the proofs below.

In the following, we denote by $\underline{\mathcal{Q}}_t^{\mathrm{fw},\ell}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\ell}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\ell})$ the cost to-go approximations after the forward and the backward pass in iteration ℓ of the RDDP scheme, respectively. An inspection of the proof of Proposition 2 reveals that the bounding property of $\underline{\mathcal{Q}}_t^{\mathrm{fw},\ell}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\ell}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\ell})$ still holds for all $\boldsymbol{x}_{t-1} \in \mathbb{R}^{n_{t-1}}$. We now show that these approximations, as well as the cost to-go functions $\mathcal{Q}_t$, also satisfy certain regularity conditions.

Lemma EC.3. *The cost to-go approximations $\underline{\mathcal{Q}}_t^{fw,\ell}$ and $(\overline{\mathcal{Q}}_t^{bw,\ell}, \underline{\mathcal{Q}}_t^{bw,\ell})$, their limits $\underline{\mathcal{Q}}_t^{fw,\infty}$ and $(\overline{\mathcal{Q}}_t^{bw,\infty}, \underline{\mathcal{Q}}_t^{bw,\infty})$ as well as the true cost to-go functions $\mathcal{Q}_t$ are uniformly continuous over $\mathcal{X}_t^\infty$. Moreover, the functions $\underline{\mathcal{Q}}_t^{fw,\ell}$ and $(\overline{\mathcal{Q}}_t^{bw,\ell}, \underline{\mathcal{Q}}_t^{bw,\ell})$ converge uniformly over $\mathcal{X}_t^\infty$.*

Proof. The pointwise limits $\underline{\mathcal{Q}}_t^{\mathrm{fw},\infty}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\infty}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\infty})$ of the cost to-go approximations $\underline{\mathcal{Q}}_t^{\mathrm{fw},\ell}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\ell}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\ell})$ over $\mathcal{X}_t$ exist due to the monotone convergence theorem and the fact that the images of $\underline{\mathcal{Q}}_t^{\mathrm{fw},\ell}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\ell}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\ell})$ are restricted to the interval $[-M, +M]$.

The cost to-go approximations $\underline{\mathcal{Q}}_t^{\mathrm{fw},\ell}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\ell}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\ell})$, their limits $\underline{\mathcal{Q}}_t^{\mathrm{fw},\infty}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\infty}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\infty})$ and the true cost to-go functions $\mathcal{Q}_t$ are convex over $\mathcal{X}_t$ and hence continuous over $\mathcal{X}_t^\infty$. The Heine-Cantor theorem then implies that these functions are in fact uniformly continuous over $\mathcal{X}_t^\infty$.

The approximations $\underline{\mathcal{Q}}_t^{\mathrm{fw},\ell}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\ell}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\ell})$ constitute monotone sequences of functions. Since the functions in these sequences and their limits are continuous over $\mathcal{X}_t^\infty$, we conclude from Dini's theorem that $\underline{\mathcal{Q}}_t^{\mathrm{fw},\ell}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\ell}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\ell})$ converge to $\underline{\mathcal{Q}}_t^{\mathrm{fw},\infty}$ and $(\overline{\mathcal{Q}}_t^{\mathrm{bw},\infty}, \underline{\mathcal{Q}}_t^{\mathrm{bw},\infty})$ uniformly. □

Lemma EC.3 allows us to study the behavior of $\overline{\mathcal{Q}}_t^{\mathrm{bw},\ell}(\underline{\boldsymbol{x}}_{t-1}^\ell)$ and $\underline{\mathcal{Q}}_t^{\mathrm{bw},\ell}(\underline{\boldsymbol{x}}_{t-1}^\ell)$ as $\ell \to \infty$. It is instrumental in the following two results, which show that the upper and lower bounds $\overline{\mathcal{Q}}_t^{\mathrm{bw},\ell}(\underline{\boldsymbol{x}}_{t-1}^\ell)$ and $\underline{\mathcal{Q}}_t^{\mathrm{bw},\ell}(\underline{\boldsymbol{x}}_{t-1}^\ell)$ converge to the true costs in the vicinity of every accumulation point $\underline{\boldsymbol{x}}_t^\infty$.

Lemma EC.4. *Let $\underline{\boldsymbol{x}}_t^\infty$ be any accumulation point of the sequence $\underline{\boldsymbol{x}}_t^\ell$, and let $\overline{\mathcal{Q}}_t^{bw,\infty}$ be the limit functions of the sequences $\overline{\mathcal{Q}}_t^{bw,\ell}$, $t = 1, \ldots, T$ and $\ell = 1, 2, \ldots$ Then $\overline{\mathcal{Q}}_t^{bw,\infty}(\underline{\boldsymbol{x}}_{t-1}^\infty) \leq \boldsymbol{q}_t^\top \underline{\boldsymbol{x}}_t^\infty + \overline{\mathcal{Q}}_{t+1}^{bw,\infty}(\underline{\boldsymbol{x}}_t^\infty)$ for all $t = 2, \ldots, T$.*

Proof. By possibly going over to subsequences, we can assume that $\underline{\boldsymbol{x}}_t^\ell$ itself converges to $\underline{\boldsymbol{x}}_t^\infty$, $\ell = 1, 2, \ldots$ The upper bound then satisfies

$$
\begin{aligned}
\overline{\mathcal{Q}}_t^{\mathrm{bw},\infty}(\underline{\boldsymbol{x}}_{t-1}^\infty) &\overset{\text{(a)}}{=} \lim_{l\to\infty} \overline{\mathcal{Q}}_t^{\mathrm{bw},\ell+1}(\underline{\boldsymbol{x}}_{t-1}^{\ell+1}) \\
&\overset{\text{(b)}}{=} \lim_{l\to\infty} \max_{\boldsymbol{\xi}_t \in \Xi_t} \min_{\boldsymbol{x}_t \in \mathbb{R}^{n_t}} \left\{ \boldsymbol{q}_t^\top \boldsymbol{x}_t + \overline{\mathcal{Q}}_{t+1}^{\mathrm{bw},\ell+1}(\boldsymbol{x}_t) \,:\, \boldsymbol{f}_t(\underline{\boldsymbol{x}}_{t-1}^{\ell+1}, \boldsymbol{\xi}_t, \boldsymbol{x}_t) \leq \boldsymbol{0} \right\} \\
&\overset{\text{(c)}}{\leq} \lim_{l\to\infty} \max_{\boldsymbol{\xi}_t \in \Xi_t} \min_{\boldsymbol{x}_t \in \mathbb{R}^{n_t}} \left\{ \boldsymbol{q}_t^\top \boldsymbol{x}_t + \overline{\mathcal{Q}}_{t+1}^{\mathrm{bw},\ell}(\boldsymbol{x}_t) \,:\, \boldsymbol{f}_t(\underline{\boldsymbol{x}}_{t-1}^{\ell+1}, \boldsymbol{\xi}_t, \boldsymbol{x}_t) \leq \boldsymbol{0} \right\}
\end{aligned}
$$

$$\stackrel{\text{(d)}}{=} \lim_{l\to\infty} \min_{\boldsymbol{x}_t\in\mathbb{R}^{n_t}} \left\{\boldsymbol{q}_t^\top \boldsymbol{x}_t + \overline{\mathcal{Q}}_{t+1}^{\text{bw},\ell}(\boldsymbol{x}_t) \,:\, \boldsymbol{f}_t(\underline{\boldsymbol{x}}_{t-1}^{\ell+1}, \overline{\boldsymbol{\xi}}_t^{\text{fw},\ell+1}, \boldsymbol{x}_t) \le \boldsymbol{0}\right\}$$
$$\stackrel{\text{(e)}}{\le} \lim_{l\to\infty} \boldsymbol{q}_t^\top \underline{\boldsymbol{x}}_t^{\ell+1} + \overline{\mathcal{Q}}_{t+1}^{\text{bw},\ell}(\underline{\boldsymbol{x}}_t^{\ell+1})$$
$$\stackrel{\text{(f)}}{=} \boldsymbol{q}_t^\top \underline{\boldsymbol{x}}_t^{\infty} + \overline{\mathcal{Q}}_{t+1}^{\text{bw},\infty}(\underline{\boldsymbol{x}}_t^{\infty}).$$

By Lemma EC.3, the identity (a) holds since

$$\lim_{l\to\infty} \left\|\overline{\mathcal{Q}}_t^{\text{bw},\ell+1}(\underline{\boldsymbol{x}}_{t-1}^{\ell+1}) - \overline{\mathcal{Q}}_t^{\text{bw},\infty}(\underline{\boldsymbol{x}}_{t-1}^{\infty})\right\| \le \lim_{l\to\infty} \Bigg[\left\|\overline{\mathcal{Q}}_t^{\text{bw},\ell+1}(\underline{\boldsymbol{x}}_{t-1}^{\ell+1}) - \overline{\mathcal{Q}}_t^{\text{bw},\infty}(\underline{\boldsymbol{x}}_{t-1}^{\ell+1})\right\| + \left\|\overline{\mathcal{Q}}_t^{\text{bw},\infty}(\underline{\boldsymbol{x}}_{t-1}^{\ell+1}) - \overline{\mathcal{Q}}_t^{\text{bw},\infty}(\underline{\boldsymbol{x}}_{t-1}^{\infty})\right\| \Bigg],$$

where the first expression on the right-hand side converges to zero since $\overline{\mathcal{Q}}_t^{\text{bw},\ell}$ converges uniformly to $\overline{\mathcal{Q}}_t^{\text{bw},\infty}$, while the second expression converges to zero since $\underline{\boldsymbol{x}}_{t-1}^{\ell}$ converges to $\underline{\boldsymbol{x}}_{t-1}^{\infty}$ and $\overline{\mathcal{Q}}_t^{\text{bw},\infty}$ is continuous. The identity (b) follows from the update of the upper bound in the backward pass of iteration $\ell+1$. The inequality (c) holds since $\overline{\mathcal{Q}}_{t+1}^{\text{bw},\ell+1} \le \overline{\mathcal{Q}}_{t+1}^{\text{bw},\ell}$. The identity (d) follows directly from the definition of $\overline{\boldsymbol{\xi}}_t^{\text{fw},\ell+1}$. The inequality (e) holds since $\underline{\boldsymbol{x}}_t^{\ell+1}$ minimizes the problem

$$\min_{\boldsymbol{x}_t\in\mathbb{R}^{n_t}} \left\{\boldsymbol{q}_t^\top \boldsymbol{x}_t + \underline{\mathcal{Q}}_{t+1}^{\text{bw},\ell}(\boldsymbol{x}_t) \,:\, \boldsymbol{f}_t(\underline{\boldsymbol{x}}_{t-1}^{\ell+1}, \overline{\boldsymbol{\xi}}^{\text{fw}_t,\ell+1}, \boldsymbol{x}_t) \le \boldsymbol{0}\right\}$$

in the forward pass of iteration $\ell+1$, and it is therefore feasible in problem (d). The identity (f), finally, holds for the same reasons as the identity (a). □

Lemma EC.5. *Let $\underline{\boldsymbol{x}}_t^{\infty}$ be any accumulation point of the sequence $\underline{\boldsymbol{x}}_t^{\ell}$, and let $\underline{\mathcal{Q}}_t^{bw,\infty}$ be the limit functions of the sequences $\underline{\mathcal{Q}}_t^{bw,\ell}$, $t=1,\ldots,T$ and $\ell=1,2,\ldots$ Then $\underline{\mathcal{Q}}_t^{bw,\infty}(\underline{\boldsymbol{x}}_{t-1}^{\infty}) \ge \boldsymbol{q}_t^\top \underline{\boldsymbol{x}}_t^{\infty} + \underline{\mathcal{Q}}_{t+1}^{bw,\infty}(\underline{\boldsymbol{x}}_t^{\infty})$ for all $t=2,\ldots,T$.*

Proof. By possibly going over to subsequences, we can assume that $\underline{\boldsymbol{x}}_t^{\ell}$ itself converges to $\underline{\boldsymbol{x}}_t^{\infty}$, $\ell=1,2,\ldots$ The lower bound then satisfies

$$\underline{\mathcal{Q}}_t^{\text{bw},\infty}(\underline{\boldsymbol{x}}_{t-1}^{\infty}) \stackrel{\text{(a)}}{=} \lim_{l\to\infty} \underline{\mathcal{Q}}_t^{\text{bw},\ell+1}(\underline{\boldsymbol{x}}_{t-1}^{\ell+1})$$
$$\stackrel{\text{(b)}}{\ge} \lim_{l\to\infty} \underline{\mathcal{Q}}_t^{\text{fw},\ell+1}(\underline{\boldsymbol{x}}_{t-1}^{\ell+1})$$
$$\stackrel{\text{(c)}}{=} \lim_{l\to\infty} \min_{\boldsymbol{x}_t\in\mathbb{R}^{n_t}} \left\{\boldsymbol{q}_t^\top \boldsymbol{x}_t + \underline{\mathcal{Q}}_{t+1}^{\text{bw},\ell}(\boldsymbol{x}_t) \,:\, \boldsymbol{f}_t(\underline{\boldsymbol{x}}_{t-1}^{\ell+1}, \overline{\boldsymbol{\xi}}_t^{\text{fw},\ell+1}, \boldsymbol{x}_t) \le \boldsymbol{0}\right\}$$
$$\stackrel{\text{(d)}}{=} \lim_{l\to\infty} \boldsymbol{q}_t^\top \underline{\boldsymbol{x}}_t^{\ell+1} + \underline{\mathcal{Q}}_{t+1}^{\text{bw},\ell}(\underline{\boldsymbol{x}}_t^{\ell+1})$$
$$\stackrel{\text{(e)}}{=} \boldsymbol{q}_t^\top \underline{\boldsymbol{x}}_t^{\infty} + \underline{\mathcal{Q}}_{t+1}^{\text{bw},\infty}(\underline{\boldsymbol{x}}_t^{\infty}).$$

Here, the identity (a) is explained by the same reasons as the identity (a) in the proof of Lemma EC.4. The inequality (b) holds since $\underline{\mathcal{Q}}_t^{\text{bw},\ell+1} \geq \underline{\mathcal{Q}}_t^{\text{fw},\ell+1}$. The identity (c) follows from the update in the forward pass of iteration $\ell+1$, the definition of $\overline{\boldsymbol{\xi}}_t^{\text{fw},\ell+1}$, as well as the fact that strong duality holds for the lower bound problems $\underline{Q}_t(\underline{\boldsymbol{x}}_{t-1}^{\ell+1};\overline{\boldsymbol{\xi}}_t^{\text{fw},\ell+1})$. The identity (d) holds by definition of $\underline{\boldsymbol{x}}_t^{\ell+1}$. The identity (e), finally, holds for the same reasons as (a). □

Equipped with the Lemmas EC.4 and EC.5, we can now prove the main result of this section.

THEOREM EC.2. *Let $\underline{\boldsymbol{x}}_t^{\infty}$ be any accumulation point of the sequence $\underline{\boldsymbol{x}}_t^{\ell}$, and let $\overline{\mathcal{Q}}_t^{bw,\infty}$ and $\underline{\mathcal{Q}}_t^{bw,\infty}$ be the limit functions of the sequences $\overline{\mathcal{Q}}_t^{bw,\ell}$ and $\underline{\mathcal{Q}}_t^{bw,\ell}$, respectively, $t=1,\ldots,T$ and $\ell=1,2,\ldots$ Then $\overline{\mathcal{Q}}_t^{bw,\infty}(\underline{\boldsymbol{x}}_{t-1}^{\infty}) = \underline{\mathcal{Q}}_t^{bw,\infty}(\underline{\boldsymbol{x}}_{t-1}^{\infty})$ for all $t=2,\ldots,T$.*

PROOF. Employing a backward induction on the time stage t, the statement follows from Proposition 2, Lemmas EC.4 and EC.5 as well as the fact that the terminal cost to-go bounds satisfy $\overline{\mathcal{Q}}_{T+1}^{\text{bw},\infty}(\boldsymbol{x}_T) = 0 = \underline{\mathcal{Q}}_{T+1}^{\text{bw},\infty}(\boldsymbol{x}_T)$ for all $\boldsymbol{x}_T \in \mathbb{R}^{n_T}$. □

From Proposition 2 we can thus conclude that any accumulation point of our RDDP scheme is an optimal solution to the generic multi-stage robust optimization problem (EC.4).

EC.4. Detailed Numerical Results

Table EC.1 provides further details of the numerical results from Section 5.

	RDDP				**ADR**
Instance	20%	10%	5%	1%	(%, time)
5-4-25	7.5s	8.7s	9.7s	12.3s	(36.4%, 23.0s)
5-4-50	26s	29s	33s	45s	(69%, 117s)
5-4-75	56s	63s	68s	80s	—
5-4-100	102s	113s	121s	133s	—
10-4-25	48s	57s	65s	81s	(68.3%, 83s)
12-4-25	105s	119s	135s	184s	(70.3%, 162s)
14-4-25	164s	197s	238s	316s	(79.4%, 177s)
16-4-25	228s	307s	381s	523s	(91.6%, 252s)
18-4-25	591s	658s	732s	896s	(66.0%, 434s)
20-4-25	978s	1,174s	1,422s	1,999s	(74.5%, 691s)
6-5-25	20s	22s	26s	33s	(41.6%, 38s)
7-6-25	52s	59s	65s	77s	(95.9%, 96s)
8-7-25	202s	213s	229s	276s	(103%, 106s)
9-8-25	568s	587s	627s	702s	(105%, 179s)
10-9-25	2,201s	2,324s	2,438s	2,714s	(99.8%, 2,162s)

Table EC.1 Optimization times required by the RDDP scheme to reduce the optimality gaps to 20%, 10%, 5% and 1% for different instances classes. The column 'ADR' reports the suboptimalilty of the affine decision rules and their runtimes. All numbers correspond to the median values over the randomly generated instances.