
Accelerating Benders decomposition for solving a sequence of sample average approximation replications

Harshit Kothari · James R. Luedtke

Abstract Sample average approximation (SAA) is a technique for obtaining approximate solutions to stochastic programs that uses the average from a random sample to approximate the expected value that is being optimized. Since the outcome from solving an SAA is random, statistical estimates on the optimal value of the true problem can be obtained by solving multiple SAA replications with independent samples. We study techniques to accelerate the solution of this set of SAA replications, when solving them sequentially via Benders decomposition. We investigate how to exploit similarities in the problem structure, as the replications just differ in the realizations of the random samples. Our extensive computational experiments provide empirical evidence that our techniques for using information from solving previous replications can significantly reduce the solution time of later replications.

Keywords Benders decomposition · Stochastic Programming · Integer Programming

1 Introduction

We study methods for solving two-stage mixed-integer stochastic programs with continuous recourse and randomness only in the right-hand side of the second-stage problem:

$$\min_{x \in X} c^\top x + \mathbb{E}_{\xi}[Q(x, \xi)], \quad (1)$$

This research has been supported by ONR grant N00014-21-1-2574.

Harshit Kothari
Department of Industrial and Systems Engineering, Wisconsin Institute of Discovery, University of Wisconsin-Madison, Madison, WI, USA E-mail: hkothari2@wisc.edu

James R. Luedtke
Department of Industrial and Systems Engineering, Wisconsin Institute of Discovery, University of Wisconsin-Madison, Madison, WI, USA E-mail: jim.luedtke@wisc.edu

where $Q(x, \xi)$ is the optimal value of the second-stage problem and is defined by

$$Q(x, \xi) := \min_y \{q^\top y : Wy = h(\xi) - T(\xi)x, y \geq 0\}. \quad (2)$$

Here $x \in X \subseteq \mathbb{Z}_+^q \times \mathbb{R}_+^{n-q}$ is the first-stage decision vector, $y \in \mathbb{R}_+^m$ is the vector of recourse decisions and ξ is a random vector.

Unless ξ has a finite and small number of possible realizations, it is usually impossible to solve this general form of a stochastic program to optimality because the expected value is hard to compute. Indeed, in [19] the authors prove that two-stage linear stochastic programming problems are $\#P$ -hard. Sample average approximation (SAA) [27] is an approach for obtaining approximate solutions to stochastic programs which approximates the expectation with an average over a finite set of scenarios sampled from the distribution of the random vector. With K scenarios (ξ_1, \dots, ξ_K) , the SAA problem is given by:

$$\hat{z}_K = \min_{x \in X} c^\top x + \sum_{k=1}^K p_k Q(x, \xi_k). \quad (3)$$

Let \hat{z}_K, \hat{x}_K be the optimal value and an optimal solution of problem (3). As these quantities are random, the multiple-replications procedure (MRP) [34, 3] has been proposed as a method to determine a confidence interval on the optimal value of the problem. MRP calculates this confidence interval for a candidate solution by solving multiple SAA replications with different independent samples of ξ . When only a fixed set of realizations of the random data is available, [29] propose a method to estimate solution quality using bootstrap aggregating to generate multiple samples and solving the corresponding SAA replications. Solving multiple SAA replications with different random samples can also be used to find higher-quality feasible solutions, e.g., [44]. Stochastic decomposition [24, 48] and stochastic approximation [40] are attractive alternatives to SAA for *continuous* stochastic programs, but are not applicable to mixed-integer stochastic programs. With these motivations in mind, we focus on the problem of solving a set of SAA replications of the same underlying stochastic programming instance derived from different samples.

Our objective in this work is to investigate how information obtained from solving one SAA replication can be used to speed up solution of other replications in order to minimize the total computation time to solve a set of SAA replications. If these SAA replications are solved on a sufficiently large cluster of computers, the wall clock time (as opposed to total computation time) could be minimized by simply solving all the replications in parallel. However, even in this setting, we argue that minimizing total computation time is an appropriate goal, as computing time is typically a limited resource, and power usage grows with total computation time. In addition, when multiple machines are available, algorithms such as Benders decomposition can be implemented in parallel on these machines when solving a single replication. Thus, solving the SAA replications sequentially and using information from early replications

to reduce the solution time of later replications can also lead to reduced wall clock time.

One approach to solving an SAA replication is to solve its deterministic equivalent form [7]. However, this formulation can become too large to solve directly when the the number of scenarios (K) is large. Decomposition methods like Benders decomposition [4, 47] and dual decomposition [11] address this by decomposing the problem and solving a sequence of smaller problems, coordinating the results, and repeating. In this work, we focus on problems with continuous recourse and randomness appearing only in the right-hand side of the subproblem constraints. This problem structure is seen in many applications such as fleet planning [34], telecommunications network design [43] and melt control [18]. Benders decomposition is a leading technique for solving problems having this structure as it is able to exploit the convexity of the recourse function, and hence we study techniques for reusing information when using Benders decomposition to solve a set of SAA replications.

Our assumptions on the problem structure imply that the dual feasible region of the Benders subproblem is fixed across all possible scenarios. Thus, our first proposal is to reuse dual solutions from previous replications to generate Benders cuts for future SAA replications by storing dual solutions in a dual solution pool (DSP). Then, whenever we would normally solve a subproblem to generate a Benders cut, we first check the DSP to see if any dual solutions there define a violated cut, and if so, we add the cut and avoid solving the subproblem. The idea of reusing stored dual solutions to generate Benders cuts has been used in different contexts when solving a single stochastic program, such as in stochastic decomposition [24, 48] and Benders decomposition [42, 2]. This is the first time this technique has been used in the context of solving multiple different SAA replications. We support this idea theoretically by estimating the number of solutions that need to be in the DSP to assure that a nearly most-violated cut can be found for a given first-stage solution.

While using the DSP can reduce time spent solving subproblems to generate Benders cuts, we make several additional contributions that reduce the computation time significantly beyond this. First, in preliminary computational studies we observed that the DSP tends to grow excessively large as the number of replications increases, as more distinct dual extreme points are discovered across replications. This makes the process of checking the DSP for violated cuts time-consuming. To address this, we propose a method for curating the DSP by retaining only some of the dual solutions in the pool. Second, we propose two techniques for choosing Benders cuts to include in the Benders main problem at the start of the algorithm. We tested these methods on two-stage stochastic linear and integer programs on three test problems. The combination of initialization techniques and DSP methods reduced the total time taken to solve these replications by half compared to using the DSP alone.

Our work contributes to a growing body of literature investigating techniques for improving methods for solving a sequence of closely related instances of an optimization problem. The surveys [6] and [16] provide an overview of

recent research investigating the use of machine learning (ML) to learn better methods for solving instances from a family of related instances. In stochastic programming, the authors in [25] train a support vector machines for the binary classification of the usefulness of a Benders cut and observe that their model allows for a reduction in the total solving time for a variety of two-stage stochastic programming instances. In [17], the authors propose to approximate the second-stage solution value with a feed-forward neural network. Recent work by [30] also leverages ML to estimate the scenario subproblem optimal values. In [35], the authors develop an ML approach to accelerate generalized Benders decomposition by estimating the optimal number of cuts that should be added to the main problem in the first iteration. There has also been work on using ML to quickly compute primal solutions for stochastic programs [5, 37]. ML-enhanced Benders decomposition has been used to accelerate solution times across various domains, including power systems [10], wireless resource allocation [31], network design problems [12], model predictive control [36], among others. These studies generally begin by solving optimization problems offline to collect data, followed by training an ML model to find algorithm parameters that speed up future solves. Our setting is fundamentally different. We solve a sequence of SAA replications of the same stochastic program to obtain statistical estimates (e.g., for confidence intervals). We do not assume we have solved any related instances prior to beginning the solution of these SAA replications, and hence there is no opportunity to train an ML model based on previous instance data to aid in solving these replications. Instead, information must be gathered and reused on the fly as each replication is solved sequentially.

This paper is organized as follows: In Section 2, we review the Benders decomposition method, for both two-stage stochastic linear and integer programs. In Section 3, we present our methods for accelerating Benders decomposition by reusing information from the solution of previous replications. In Section 4, we present our results from computational experiments.

2 Benders Decomposition

In this section, we describe the Benders decomposition method [4, 47] for solving the SAA (3). The dual of the subproblem (2) for scenario $k \in [K] := \{1, \dots, K\}$ is given by:

$$\max_{\pi} \{(h(\xi_k) - T(\xi_k)x)^\top \pi : W^\top \pi \leq q\}. \quad (4)$$

We denote the dual feasible region as $\Pi = \{\pi : W^\top \pi \leq q\}$, which is independent of the scenario $k \in [K]$.

We make the following assumptions about the stochastic program (1):

- The subproblem dual feasible region Π is non-empty.
- Relatively complete recourse: For every feasible first-stage solution $x \in X$ and every ξ in the support of the random variable ξ , there exists a feasible decision to subproblem (2).

We make the relatively complete recourse assumption mainly to simplify exposition. In Section 3.4.1, we introduce extensions of our methods to handle the case where relatively complete recourse does not hold.

Under these assumptions both the primal and the dual of the subproblem (2) have an optimal solution and from strong duality, we conclude that their optimal values are equal. Let \mathcal{V} denote the set of all the vertices of Π . Then,

$$Q(x, \xi_k) = \max_{\pi} \{(h(\xi_k) - T(\xi_k)x)^\top \pi : \pi \in \mathcal{V}\}. \quad (5)$$

Benders decomposition is based on a reformulation of (3), that introduces a new variable θ_k to represent the optimal value of subproblem k for each scenario $k \in [K]$. Using (5), the reformulation is as follows:

$$\begin{aligned} \min_{x \in X, \theta} \quad & c^\top x + \sum_{k=1}^K p_k \theta_k \\ \text{s.t.} \quad & Ax = b, \\ & \theta_k \geq (h(\xi_k) - T(\xi_k)x)^\top \pi, \quad \pi \in \mathcal{V}, \quad k \in [K]. \end{aligned} \quad (6)$$

In the following subsections we discuss how this reformulation is used within Benders decomposition to solve two-stage stochastic linear programs (LPs) (Section 2.1) and two-stage stochastic IPs with continuous recourse (Section 2.2).

2.1 Stochastic LPs

The size of the reformulation (6) depends on the number of vertices of the dual subproblem (\mathcal{V}), which is usually too large to explicitly enumerate. Thus, a delayed cut generation scheme is used to iteratively add these constraints. Specifically, Benders decomposition works with a “main problem” which has the form of (6) but at each iteration t only includes a subset $\mathcal{V}_{k,t}$ of the constraints for each scenario $k \in [K]$:

$$\text{MP}^t = \min_{x \in X} \quad c^\top x + \sum_{k=1}^K p_k \theta_k \quad (7a)$$

$$\text{s.t.} \quad Ax = b, \quad (7b)$$

$$\theta_k \geq (h(\xi_k) - T(\xi_k)x)^\top \pi_k, \quad \pi_k \in \mathcal{V}_{k,t}, \quad k \in [K]. \quad (7c)$$

The constraints (7c) in the reformulation are called Benders cuts.

Algorithm 1 outlines the Benders decomposition algorithm for LPs. We start by solving the main problem with some initial cuts of the form (7c) included. For simplicity of exposition, we assume these initial cuts are sufficient to assure problem (7) is bounded and hence has an optimal solution, $(x^t, \{\theta_k^t\}_{k \in [K]})$. Next, we solve subproblems (2), with $\xi = \xi_k$, to check for violated cuts and evaluate the objective for each subproblem $k \in [K]$. If this objective value is greater than θ_k^t , we have identified a violated cut. This leads

to the generation of a Benders cut, defined by the dual solution π_k^t . After iterating through all the subproblems, and adding violated cuts to the main problem, we solve the updated main problem. This iterative process continues until no further violated cuts are found, leading to an optimal solution of the original problem.

At each iteration, the main problem objective provides a lower bound, L^t , to the problem because it is a relaxation to the reformulation (6). We also obtain an upper bound U^t in each iteration by solving the subproblems at every iteration. The difference between U^t and L^t can be used as a convergence condition, terminating when this difference falls below a tolerance denoted by ϵ .

Algorithm 1 Benders decomposition algorithm.

```

1: Initialize  $t := 0$ 
2: Initialize  $\mathcal{V}_{k,0}$  for all  $k \in [K]$                                  $\triangleright$  Initialization step
3: repeat
4:   cutAdded  $\leftarrow$  False
5:   Solve  $MP^t$  (7) and obtain  $(x^t, \{\theta_k^t\}_{k \in [K]})$ 
6:    $L^t \leftarrow c^T x^t + \frac{1}{K} \sum_{k=1}^K \theta_k^t$                  $\triangleright$  Calculate lower bound
7:   for all  $k \in [K]$  do
8:     Solve (2), with  $\xi = \xi_k$ ,  $x = x^t$  and obtain  $Q(x^t, \xi_k)$  and dual solution  $\pi_k^t$ 
9:     if  $Q(x^t, \xi_k) > \theta_k^t$  then
10:       $\mathcal{V}_{k,t+1} \leftarrow \mathcal{V}_{k,t} \cup \{\pi_k^t\}$        $\triangleright$  Store dual solution to add cut in next iteration
11:      cutAdded  $\leftarrow$  True
12:    end if
13:   end for
14:    $U^t \leftarrow c^T x^t + \frac{1}{K} \sum_{k=1}^K Q(x^t, \xi_k)$                  $\triangleright$  Calculate upper bound
15:    $t \leftarrow t + 1$ 
16: until cutAdded = False or  $U^t - L^t \leq \epsilon$ 

```

This version of Benders decomposition in which we introduce auxiliary variables θ_k for every scenario $k \in [K]$ is called the multi-cut version. An alternate version is the single-cut version [47] in which a single variable is used to represent the epigraph of the expected value of the subproblem objective taken over the full set of scenarios. We focus on the multi-cut version in this paper, but in Section 3.4.2 we discuss how the methods proposed here can be extended to the single-cut version of the algorithm.

2.2 Stochastic IPs

When dealing with stochastic programs featuring integer variables in the first-stage, solving the main problem can be computationally expensive because it is an integer program. Hence, instead of iteratively solving the main problem and adding cuts, stochastic IPs can alternatively be solved using the branch-and-cut method [21, 38].

A standard branch-and-cut algorithm for solving two-stage stochastic integer programs with continuous recourse is outlined in Algorithm 2. The al-

Algorithm 2 Branch-and-cut for IPs.

```

1: Initialize  $\mathcal{N} \leftarrow \{0\}$ ,  $\bar{z} \leftarrow +\infty$ ,  $(x^*, \{\theta_k^*\}_{k \in [K]}) \leftarrow \emptyset$ 
2: Initialize  $LP_0$  to be the relaxation of main problem with initial cuts
3: while  $\mathcal{N} \neq \emptyset$  do
4:   Choose a node  $i \in \mathcal{N}$ ,  $\mathcal{N} \leftarrow \mathcal{N} \setminus \{i\}$ 
5:   Solve  $LP_i$ . If feasible, obtain optimal solution  $(\hat{x}, \{\hat{\theta}_k\}_{k \in [K]})$  and optimal value  $\hat{z}$ 
6:   if  $LP_i$  feasible and  $\hat{z} < \bar{z}$  then
7:     if  $\hat{x} \in X$  then                                 $\triangleright$  Search for violated Benders cuts
8:       cutAdded  $\leftarrow$  False
9:       for all  $k \in [K]$  do
10:        Solve (2), with  $\xi = \xi_k$ ,  $x = \hat{x}$  and obtain  $Q(\hat{x}, \xi_k)$  and dual solution  $\pi_k$ 
11:        if  $Q(\hat{x}, \xi_k) > \hat{\theta}_k$  then
12:          Add Benders cut:  $\theta_k \geq (h(\xi_k) - T(\xi_k)x)^\top \pi_k$ 
13:          cutAdded  $\leftarrow$  True
14:        end if
15:      end for
16:      if cutAdded = True then
17:        go to step 5
18:      else
19:        Update incumbent solution  $x^* \leftarrow \hat{x}$ , and  $\bar{z}$            $\triangleright (\hat{x}, \hat{\theta})$  is feasible
20:      end if
21:      else
22:        Partition the problem and update  $\mathcal{N}$                    $\triangleright$  Branching step
23:      end if
24:    end if
25:  end while
26: return  $(x^*, \{\theta_k^*\}_{k \in [K]})$ 

```

gorithm begins by adding initial Benders cuts to the main problem. Starting the algorithm with some initial cuts defined can help speed up convergence of the algorithm [41,21]. The LP relaxation of this main problem is represented by the root node or LP_0 . This node is added to the list of candidate branch-and-bound nodes \mathcal{N} .

In each iteration, the algorithm selects a node from \mathcal{N} and solves the corresponding LP relaxation, a relaxation of the problem with added constraints from branching. This LP generates a candidate solution $(\hat{x}, \{\hat{\theta}_k\}_{k \in [K]})$. If \hat{x} violates the integrality constraints (i.e., $\hat{x} \notin X$), the algorithm creates two new subproblems by subdividing the feasible region (branching step). These two new subproblems are then appended to \mathcal{N} .

If $\hat{x} \in X$, we solve subproblems (2) with $x = \hat{x}$ and for $\xi = \xi_k$ for each $k \in [K]$ to check if the solution $(\hat{x}, \{\hat{\theta}_k\}_{k \in [K]})$ is feasible to the original problem. If any violated Benders cuts are found, they are added to the main problem, and the relaxation at that node is solved again. This process of dynamically adding cuts can be implemented using a “Lazy constraint callback” in MIP solvers. If no violated cuts are identified, then the current solution is feasible. Whenever we find a feasible solution with a better objective value than the current best solution, we update the upper bound \bar{z} and the incumbent solution x^* . The algorithm terminates when there are no more nodes to explore in \mathcal{N} .

In [21], the authors describe an implementation of Benders decomposition using the branch-and-cut technique. They highlight the importance of initializing the algorithm with a good set of cuts included in the main problem to accelerate convergence. To achieve this, they first solve the LP relaxation of the original problem using Benders decomposition and retain all identified cuts in the initial main problem of the IP. We follow their approach but only retain the cuts which are active at the optimal solution of the LP relaxation to manage the size of the main problem.

3 Techniques for Reusing Information to Accelerate Benders Decomposition

We now consider a setting where we wish to solve a sequence of SAA replications of the form (3), each with an independently drawn set of scenarios. Suppose we wish to solve M such replications, and each having K scenarios, $(\xi_1^r, \dots, \xi_K^r)$ for $r = 1, \dots, M$. The question we investigate is, when solving the SAA (3) associated with a replication $r > 1$ with Benders decomposition, how can we use information obtained from solving replications $1, \dots, r-1$ to reduce the solution time?

We first consider a simple warm-start strategy for stochastic IPs: use the optimal first-stage solution (say \tilde{x}) from a previous replication as an initial feasible solution in the branch-and-bound process. When solving a new replication, we evaluate \tilde{x} by solving the subproblems (2) at (\tilde{x}, ξ_k) for all scenarios $k \in [K]$, and pass the resulting solution $(\tilde{x}, Q(\tilde{x}, \xi_k)_{k \in [K]})$ to the solver. This can help the solution process by providing an upper bound for pruning nodes.

In the remainder of this section we propose other techniques for reusing information to accelerate Benders decomposition. To do so, we consider the key computational tasks within Benders decomposition: solving the main problem and solving subproblems to generate cuts. We propose two techniques, the use of a dual solution pool (Section 3.1) and a curated version of this pool (Section 3.2) to reduce time solving the subproblems. We also propose initialization techniques in Section 3.3 aimed at accelerating convergence of the algorithm.

3.1 Dual Solution Pool

As discussed in Section 2, the extreme points of the dual of the subproblem (2) are used to generate Benders cuts. Given our assumption that W and q are fixed in the second-stage problem, the dual feasible region of the subproblem is the same for every possible scenario, as shown in (4). This implies that the dual solutions from previous replications can be used to generate valid Benders cuts for the current replication without solving subproblems.

To exploit this observation, as we discover dual solutions when solving an SAA replication, we store them in a dual solution pool (DSP) denoted by \mathcal{V}_{DSP} . Then, when solving a new SAA replication, these dual solutions can

be used to generate Benders cuts if they cut off the current primal solution, thereby potentially avoiding the need to solve the subproblem (2).

We first describe how the DSP is used when solving two-stage stochastic LPs. Algorithm 1 is modified by running Algorithm 3 before starting the loop of solving subproblems (line 7). After obtaining a main problem solution, x^t , we check the DSP, \mathcal{V}_{DSP} , for each scenario $k \in [K]$ to find if it contains any dual solutions that define a Benders cut violated by x^t . Specifically, we solve the following problem for each scenario $k \in [K]$ (line 3 of Algorithm 3):

$$\bar{Q}_k(x^t, \mathcal{V}_{DSP}) := \max_{\pi \in \mathcal{V}_{DSP}} \pi^\top (h(\xi_k) - T(\xi_k)x^t). \quad (8)$$

That is, for each $k \in [K]$, $\bar{Q}_k(x^t, \mathcal{V}_{DSP})$ is defined as the objective value of subproblem (5) evaluated at x^t , with the feasible region replaced by \mathcal{V}_{DSP} . If problem (8) identifies a violated cut, i.e., $\bar{Q}_k(x^t, \mathcal{V}_{DSP}) > \theta_k^t$ for some scenario $k \in [K]$, we add the identified cut for each such scenario and proceed with solving the updated main problem (lines 4-5 in Algorithm 1). If no violated cut is found from the DSP for any scenario $k \in [K]$, then Algorithm 3 proceeds with solving the scenario subproblems (line 7). If any of these subproblems yields a violated Benders cut, the dual solution that defines the cut is not in the DSP (otherwise we would have found the violated cut when running Algorithm 3). Thus, every dual solution that defines a violated cut is saved, and at the end of the replication we add these the DSP for use in the solution of the following SAA replications.

We propose to use the DSP in two ways when solving two-stage stochastic IPs. First, as discussed in Section 2.2, the LP relaxation of the stochastic IP is solved to obtain a set of initial cuts to include in the main problem before starting the branch-and-cut process. The DSP can be used exactly as described in the last paragraph for solving two-stage stochastic LPs to accelerate this process. Second, we can apply Algorithm 3 when an integer feasible solution $\hat{x} \in X$ is found in Algorithm 2 before solving subproblems (lines 9-15 of Algorithm 2). If a cut is found from the DSP for any scenario $k \in [K]$ it is added to the main problem and the LP relaxation at the node is re-solved (line 5). If no cuts are found in the DSP, then the subproblems are solved as usual.

Algorithm 3 Using DSP to look for a violated Benders cut.

```

1: Input: Current main problem first-stage solution:  $x^t$ , DSP:  $\mathcal{V}_{DSP}$ 
2: cutAdded  $\leftarrow$  False
3: for all  $k \in [K]$  do
4:   Evaluate  $\bar{Q}_k(x^t, \mathcal{V}_{DSP})$  using (8) and let  $\bar{\pi}_k$  be a dual solution achieving the max
5:   if  $\bar{Q}_k(x^t, \mathcal{V}_{DSP}) > \theta_k^t$  then
6:     Add violated cut:  $\theta_k \geq \bar{\pi}_k^\top (h(\xi_k) - T(\xi_k)x^t)$ 
7:     cutAdded  $\leftarrow$  True
8:   end if
9: end for

```

We next turn to a theoretical investigation of the size of the DSP that is required for it to be expected to successfully find violated Benders cuts. Let $\Pi \subset \mathbb{R}^d$ be a polytope with vertex set \mathcal{V} and define $z^* : \mathbb{R}^d \rightarrow \mathbb{R}$ as:

$$z^*(c) = \max_{\pi \in \Pi} c^\top \pi. \quad (9)$$

In addition, define $\pi^* : \mathbb{R}^d \rightarrow \mathcal{V}$, where $\pi^*(c)$ is an optimal solution of (9) for each $c \in \mathbb{R}^d$ (ties can be broken arbitrarily). Let $A_v, v \in \mathcal{V}$ be a partition of \mathbb{R}^d such that $A_v = \{c \in \mathbb{R}^d : \pi^*(c) = v\}$. Assume we are solving a sequence of N problems of the form (9) with $c = \mathbf{c}^i$ for $i = 1, \dots, N$ where each \mathbf{c}^i is a random vector drawn independently from the same distribution. Given the set of optimal solutions obtained from these problems (think of it as the DSP), we then obtain a new random coefficient $\hat{\mathbf{c}}$ from the same distribution and wish to understand whether any solution in this set is optimal for this new coefficient. The following lemma relates this probability to the geometry of the set Π .

Lemma 1 *Let $\mathbf{c}^1, \dots, \mathbf{c}^N$ and $\hat{\mathbf{c}}$ be independent random vectors identically distributed according to a distribution \mathcal{D} . For each vertex $v \in \mathcal{V}$, let $q_v = \mathbb{P}_{\mathcal{D}}(\mathbf{c} \in A_v)$. Define ‘Failure’ as the event that none of the solutions $\pi^*(\mathbf{c}^i)$ for $i = 1, \dots, N$ is optimal for $z^*(c)$. Then,*

$$\mathbb{P}(\text{Failure}) = \sum_{v \in \mathcal{V}} q_v (1 - q_v)^N.$$

Proof. We analyze the probability by conditioning on the realization of the target vector $\hat{\mathbf{c}}$. Suppose $\pi^*(\hat{\mathbf{c}}) = v$, and hence $\hat{\mathbf{c}} \in A_v$. The sample set fails to recover v if and only if no sample \mathbf{c}^i falls into A_v . For a single sample \mathbf{c}^i , the probability of missing region A_v is $1 - q_v$. Since the N samples are independent, the probability that all of them miss A_v is:

$$\mathbb{P}(\text{Failure} \mid \hat{\mathbf{c}} \in A_v) = (1 - q_v)^N.$$

To find the unconditional failure probability, we integrate over the possible realizations of $\hat{\mathbf{c}}$. Since the regions A_v partition the space, we apply the law of total probability summing over all $v \in \mathcal{V}$:

$$\begin{aligned} \mathbb{P}(\text{Failure}) &= \sum_{v \in \mathcal{V}} \mathbb{P}(\text{Failure} \mid \hat{\mathbf{c}} \in A_v) \mathbb{P}(\hat{\mathbf{c}} \in A_v) \\ &= \sum_{v \in \mathcal{V}} (1 - q_v)^N \cdot q_v. \end{aligned} \quad \square$$

If all vertices are equally likely to be optimal ($q_v = 1/|\mathcal{V}|$), then the failure probability in Lemma 1 simplifies to

$$\sum_{v \in \mathcal{V}} \frac{1}{|\mathcal{V}|} \left(1 - \frac{1}{|\mathcal{V}|}\right)^N = \left(1 - \frac{1}{|\mathcal{V}|}\right)^N \leq e^{-N/|\mathcal{V}|}.$$

Thus, to assure a failure probability less than ρ would require $N \geq |\mathcal{V}| \ln(1/\rho)$. This illustrates dependence on the geometry of Π in terms of its number of

vertices $|\mathcal{V}|$. For a standard simplex in \mathbb{R}^d , this is a modest bound as $|\mathcal{V}| = d$. However, in general, $|\mathcal{V}|$ may be exponential in d . On the other hand, the following corollary demonstrates that if a subset of vertices has high probability of being optimal, then a sample size on the order of the size of that subset is sufficient to achieve low failure probability.

Corollary 1 *Let $\delta > 0$ and assume $S \subseteq \mathcal{V}$ satisfies $\sum_{v \in S} q_v \geq 1 - \delta$. Then,*

$$\mathbb{P}(\text{Failure}) \leq \frac{|S|}{eN} + \delta.$$

Proof. Applying Lemma 1 and the assumption yield

$$\begin{aligned} \mathbb{P}(\text{Failure}) &= \sum_{v \in \mathcal{V}} q_v (1 - q_v)^N \\ &= \sum_{v \in S} q_v (1 - q_v)^N + \sum_{v \in \mathcal{V} \setminus S} q_v (1 - q_v)^N \\ &\leq \sum_{v \in S} q_v (1 - q_v)^N + \sum_{v \in \mathcal{V} \setminus S} q_v \leq \sum_{v \in S} q_v (1 - q_v)^N + \delta. \end{aligned} \quad (10)$$

Next observe that for $q \in [0, 1]$, $q(1 - q)^N \leq q e^{-Nq}$ and the maximum of $f(q) = q e^{-Nq}$ over $q \geq 0$ is obtained at $q = 1/N$. Indeed, $f'(q) = e^{-Nq}(1 - Nq)$ and hence the unique stationary point occurs at $q = 1/N$. Since $f''(q) = -N e^{-Nq} + (1 - Nq)(-N e^{-Nq}) = (-2N + N^2 q)e^{-Nq}$, $f''(1/N) < 0$ and hence this is the maximum. Thus, $q(1 - q)^N \leq (1/N)e^{-1} = 1/(eN)$. Substituting into (10) yields

$$\mathbb{P}(\text{Failure}) \leq \sum_{v \in S} \frac{1}{eN} + \delta = \frac{|S|}{eN} + \delta. \quad \square$$

To relate this analysis to our use of DSP for searching for a violated Benders cut, recall that to search for a violated cut from the DSP, we solve the problem (8) in the hopes of finding a violated cut, rather than solving the true dual subproblem (4). We thus consider \mathcal{V}_{DSP} as the set of the solutions, $\{\pi^*(\mathbf{c}^1), \pi^*(\mathbf{c}^2), \dots, \pi^*(\mathbf{c}^N)\}$, and the cost vector as $\hat{c} = h(\xi_k) - T(\xi_k)x^t$. If the solution x^t has been used in a previous replication to generate Benders cuts, then \hat{c} has the same distribution as the distribution of the cost vectors in the subproblems (4) in that previous replication, and hence \mathcal{V}_{DSP} will contain dual optimal solutions from K samples from the same distribution as \hat{c} . Then Lemma 1 provides a bound on the probability that the best solution in the DSP (i.e., the dual solution that provides the largest right-hand side value of a Benders cut) is optimal (i.e., the maximum possible Benders cut violation). This implies that if a violated Benders cut exists, one would likely be found in the DSP. This analysis is not directly applicable if the primal solution x^t was not seen in a previous replication. However, the DSP contains dual solutions derived from many other primal solutions, and hence may still be useful for generating Benders cuts, which we verify empirically in Section 4.

3.2 Curated DSP

Our preliminary experiments indicated that the number of distinct dual extreme points discovered and stored in the DSP tends to grow rapidly as we solve more SAA replications, even though the underlying dual polytope remains fixed. This growth makes it increasingly time-consuming to search the DSP for violated cuts. To address this, we propose to use a curated DSP in which we restrict the set of stored dual solutions to a more manageable size. This restricted set of stored dual solutions is denoted by \mathcal{V}_{cur} . We use the curated DSP exactly as the DSP is used as described in Section 3.1. The only difference is that when searching for cuts based on past dual solutions, we search the set \mathcal{V}_{cur} rather than the full DSP (\mathcal{V}_{DSP}).

Our approach for creating a curated DSP is detailed in Algorithm 4. After each SAA replication, the full DSP \mathcal{V}_{DSP} is partitioned into three sets: the permanent set (\mathcal{V}_{perm}), the trial set (\mathcal{V}_{trial}), and the remaining solutions which we refer to as “the bench”. The permanent set includes dual solutions that have generated violated cuts in multiple past replications. Once a solution is added to the permanent set, it remains there for all subsequent replications. The trial set consists of newly generated dual solutions discovered during the previous replication. These newly discovered solutions are included in the curated DSP in the following replication. If a solution in the trial set successfully identifies a violated cut, it is added to the permanent set in the next replication. The curated DSP, \mathcal{V}_{cur} , is the union of the permanent and trial sets. After finishing solving a replication, all dual solutions used to define violated cuts in the replication are added to a set called \mathcal{V}_{used} . After each replication, dual solutions are re-evaluated: those that were already in the DSP and defined a violated cut (i.e., they are in the set \mathcal{V}_{used}) are added to the permanent set \mathcal{V}_{perm} . Any dual solution that was newly generated during the current replication, and was not already part of the full DSP, is added to the trial set \mathcal{V}_{trial} . This allows new solutions to be tested and potentially included in future replications if they prove effective. For the next replication, the curated DSP is updated by combining the updated permanent and trial sets.

3.3 Initialization Techniques

The Benders decomposition method can be significantly accelerated if the main problem is initialized with a good set of Benders cuts. For linear programs, adding the “right” initial cuts can theoretically lead to convergence in just one iteration. Initializing the main problem with Benders cuts has also been observed to be important to solve stochastic IPs[41,21]. Thus, we investigate techniques for determining a set of Benders cuts to add to the initial Benders main problem. We focus our discussion on initialization techniques for two-stage stochastic IPs, and then discuss adaptations of these ideas that we propose for two-stage stochastic LPs.

Algorithm 4 Curated dual solution pool.

```

1: Initialize  $\mathcal{V}_{DSP}$  with dual solutions collected in the first replication
2: Initialize  $\mathcal{V}_{perm} \leftarrow \emptyset$ ,  $\mathcal{V}_{trial} \leftarrow \emptyset$ 
3: for all replications  $r = 2, \dots, M$  do
4:    $\mathcal{V}_{cur} \leftarrow \mathcal{V}_{perm} \cup \mathcal{V}_{trial}$ 
5:   Solve replication using Benders decomposition
6:   Add dual solutions used to generate cuts to  $\mathcal{V}_{used}$ 
7:   for  $\pi \in \mathcal{V}_{used}$  do                                 $\triangleright$  Re-evaluation of duals
8:     if  $\pi \in \mathcal{V}_{DSP}$  then
9:        $\mathcal{V}_{perm} \leftarrow \mathcal{V}_{perm} \cup \{\pi\}$ 
10:    else if  $\pi \notin \mathcal{V}_{DSP}$  then
11:       $\mathcal{V}_{trial} \leftarrow \mathcal{V}_{trial} \cup \{\pi\}$            $\triangleright \pi$  is newly discovered
12:       $\mathcal{V}_{DSP} \leftarrow \mathcal{V}_{DSP} \cup \{\pi\}$            $\triangleright$  Update full DSP
13:    end if
14:   end for
15:   Re-evaluate and update  $\mathcal{V}_{perm}$  and  $\mathcal{V}_{trial}$  for the next replication
16: end for

```

In our approach, we aim to leverage dual solutions collected from prior replications (\mathcal{V}_{DSP}) to generate initialization cuts for the current replication. Using all the collected dual solutions to generate a cut for every scenario would provide an initialization of the algorithm that provides the best possible bound given those dual solutions. However, this would also lead to a large number of initial cuts in the main problem, slowing down the solution of the main problem LP relaxations throughout the algorithm. Thus, we explore techniques for selecting, for each $k \in [K]$, a subset $\mathcal{V}_k^{sel} \subseteq \mathcal{V}_{DSP}$ of dual solutions from the DSP from which to add Benders cuts. To guide this selection, we use the first-stage solutions encountered during Benders decomposition in previous replications. We denote the set of all feasible solutions found during the solution process in past replications as X^{feas} (e.g., this includes integer first-stage solutions the solver finds via its internal heuristics and integer solutions discovered at nodes in the branch-and-bound search), and denote the set of optimal solutions of previous replications as X^{opt} . We propose two methods for choosing the initial cuts to add: static initialization and adaptive initialization, described in the following two subsections.

3.3.1 Static Initialization

The main idea behind this approach is that a previous optimal solution has a high likelihood of being near-optimal for this replication [26]. Therefore, we want to initialize the algorithm with cuts that maximize the objective value of the subproblem dual when evaluated on solutions in our set of previous optimal solutions X^{opt} . For each previous optimal solution, $\bar{x} \in X^{opt}$, and for each scenario $k \in [K]$, we identify a dual solution from the DSP that maximizes the subproblem objective value, i.e., a dual solution that achieves the maximum in (8). If there is a tie between multiple dual solutions, we randomly select one of them. We could add the Benders cut corresponding to this dual solution for each $\bar{x} \in X^{opt}$ and each $k \in [K]$. However, in our experiments, we found

that doing this for all previous optimal solutions yielded many cuts in the initialization that were not useful. Thus, we only add cuts for the first two optimal solutions in X^{opt} , leading to at most two cuts per scenario.

3.3.2 Adaptive Initialization

In this initialization technique, we use both the set of optimal solutions of previous replications, X^{opt} , and the full set of previously found feasible solutions, X^{feas} , to identify initial Benders cuts to include in the main model. This technique proceeds in two phases. In phase one, we find the solution x^{WS} which has the lowest objective value among the solutions in X^{opt} for the current SAA replication. In the second phase, we identify a set of cuts which ensure that the objective values of all other solutions in X^{feas} are suboptimal compared to x^{WS} when evaluated on the selected set of cuts. These cuts provide a strong initialization and ensure that these solutions are not encountered later in Benders decomposition, as, by design, their objective values in the model will be worse than the objective value of the initial solution we provide to the model. The hope is that the set of solutions, X^{feas} , serves as a representative approximation for the entire feasible region X , and constructing the cuts this way will lead to faster convergence. In the limit, if we had all the feasible primal solutions in X^{feas} , then this initialization would find the optimal solution.

Phase One. Let $z(x) = c^\top x + \sum_{k \in [K]} p_k Q(x, \xi_k)$ denote the objective value of a first-stage solution $x \in X$. In this phase, our goal is to find the solution $x^{WS} \in X^{opt}$ with the lowest true objective value, i.e., it satisfies

$$z(x^{WS}) \leq z(x) \quad \forall x \in X^{feas}.$$

While this could be accomplished by directly evaluating $z(x)$ for all $x \in X^{feas}$ this would be computationally expensive as it requires solving all scenario subproblems for each solution. Therefore, to find x^{WS} , we use an approximation of the true objective value,

$$\bar{z}(x, \mathcal{V}_{DSP}) = c^\top x + \sum_{k \in [K]} p_k \bar{Q}_k(x, \mathcal{V}_{DSP}),$$

where, for each $k \in [K]$, $\bar{Q}_k(x, \mathcal{V}_{DSP})$ is the approximate value of the subproblem for scenario k as defined in (8). $\bar{z}(x, \mathcal{V}_{DSP})$ represents the approximate objective value of x , with the subproblems being evaluated on the DSP (\mathcal{V}_{DSP}) instead of the full feasible region (Π). This approximation always underestimates the true objective value, i.e., $\bar{z}(x, \mathcal{V}_{DSP}) \leq z(x)$ as $\bar{Q}_k(x, \mathcal{V}_{DSP}) \leq Q(x, \xi_k)$ for all $k \in [K]$ because $\mathcal{V}_{DSP} \subseteq \Pi$.

Algorithm 5 outlines the process of finding x^{WS} . We start by calculating the approximate objective $\bar{z}(x, \mathcal{V}_{DSP})$ for each solution in X^{opt} and then arrange them in ascending order of this approximate objective value. We then check whether the solution with the smallest value of $\bar{z}(x, \mathcal{V}_{DSP})$, say \bar{x} , has the

lowest true objective value, by solving the scenario subproblems (2) with $x = \bar{x}$ and $\xi = \xi_k$ for each scenario $k \in [K]$ to find the true objective value of this solution. In this process, we may generate new dual solutions which are added to the DSP. If any new dual solutions are found, this will increase the value of $\bar{z}(\bar{x}, \mathcal{V}_{DSP})$ to $z(\bar{x})$, and \bar{x} might no longer be the solution with the lowest value of $\bar{z}(x, \mathcal{V}_{DSP})$ among X^{opt} . Therefore, we re-evaluate $\bar{z}(x, \mathcal{V}_{DSP})$ for all solutions on the updated \mathcal{V}_{DSP} to see if \bar{x} remains the best candidate. If so, the algorithm terminates; otherwise, we select the new minimizer and repeat the process. When the algorithm terminates it holds that

$$z(x^{WS}) = \bar{z}(x^{WS}, \mathcal{V}_{DSP}) \leq \bar{z}(x, \mathcal{V}_{DSP}) \leq z(x) \quad \forall x \in X^{opt}$$

and hence we have found the solution with the lowest true objective value in the set X^{opt} . This algorithm is guaranteed to converge in at most $|X^{opt}|$ iterations because in each iteration we expand \mathcal{V}_{DSP} such that $z(x) = \bar{z}(x, \mathcal{V}_{DSP})$ for a new $x \in X^{opt}$.

Algorithm 5 Adaptive initialization - phase one.

- 1: Evaluate $\bar{z}(x, \mathcal{V}_{DSP})$ for all $x \in X^{opt}$
- 2: **repeat**
- 3: $x^{WS} \leftarrow \arg \min \{\bar{z}(x, \mathcal{V}_{DSP}) : x \in X^{opt}\}$
- 4: Solve subproblem (2), with $x = x^{WS}$, $\xi = \xi_k$ for each $k \in [K]$
- 5: Update \mathcal{V}_{DSP} with newly found dual solutions
- 6: Re-evaluate $\bar{z}(x, \mathcal{V}_{DSP})$ for all $x \in X^{opt}$
- 7: **until** $\arg \min \{\bar{z}(x, \mathcal{V}_{DSP}) : x \in X^{opt}\} = x^{WS}$
- 8: **return** x^{WS}

Phase two. The goal of phase two is to identify a (hopefully small) set $\mathcal{V}_k^{sel} \subseteq \mathcal{V}_{DSP}$ for each scenario $k \in [K]$ such that, when the objective value of each solution in X^{feas} is evaluated using the Benders cuts defined by these solutions, the evaluation is higher than $z(x^{WS})$. As a result, when the Benders algorithm proceeds, none of these solutions will be identified as a candidate solution that might be better than x^{WS} . In this process, since we are considering more solutions (X^{feas}) than we considered in phase one (X^{opt}), we may find a solution $x \in X^{feas}$ that has a better objective value than x^{WS} identified in phase one, in which case we update x^{WS} .

Given a collection of sets of dual solutions $\mathcal{V}_k^{sel} \subseteq \mathcal{V}_{DSP}$ for $k \in [K]$, we define the lower bound approximation of the objective value of x ,

$$\bar{z}(x, \{\mathcal{V}_k^{sel}\}_{k \in [K]}) := c^\top x + \sum_k p_k \bar{Q}_k(x, \mathcal{V}_k^{sel}),$$

where, for each $k \in [K]$, $\bar{Q}_k(x, \mathcal{V}_k^{sel})$ is the objective of the subproblem k evaluated using the set of dual solutions \mathcal{V}_k^{sel} :

$$\bar{Q}_k(x, \mathcal{V}_k^{sel}) := \max_{\pi \in \mathcal{V}_k^{sel}} \pi^\top (h(\xi_k) - T(\xi_k)x).$$

Note that for all scenarios, $k \in [K]$, $\mathcal{V}_k^{sel} \subseteq \mathcal{V}_{DSP} \subseteq \Pi$ and so

$$\bar{Q}_k(x, \mathcal{V}_k^{sel}) \leq \bar{Q}_k(x, \mathcal{V}_{DSP}) \leq Q(x, \xi_k)$$

for all $x \in X$. Thus,

$$\bar{z}(x, \{\mathcal{V}_k^{sel}\}_{k \in [K]}) \leq \bar{z}(x, \mathcal{V}_{DSP}) \leq z(x) \quad \forall x \in X.$$

Using this notation, we restate the primary goal of phase two which is to find sets of dual solutions $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$ that satisfy

$$z(x^{WS}) \leq \bar{z}(x, \{\mathcal{V}_k^{sel}\}_{k \in [K]}) \quad \forall x \in X^{feas}. \quad (11)$$

The pseudocode for phase two is presented in Algorithm 6. To initialize $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$, we first add the dual solutions obtained by solving subproblems for x^{WS} . This ensures that $\bar{z}(x^{WS}, \{\mathcal{V}_k^{sel}\}_{k \in [K]}) = z(x^{WS})$. For IPs, we first solve the LP relaxation to add initialization cuts and thus we also include the dual solutions which defined active cuts at the optimal solution of the LP relaxation in the sets $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$. Next, we evaluate $\bar{z}(x, \{\mathcal{V}_k^{sel}\}_{k \in [K]})$ for all $x \in X^{feas}$. If the solution with least value of $\bar{z}(x, \{\mathcal{V}_k^{sel}\}_{k \in [K]})$, say \bar{x} , has the same objective value as $z(x^{WS})$, then we have converged having achieved our goal (11).

However, if $\bar{z}(\bar{x}, \{\mathcal{V}_k^{sel}\}_{k \in [K]}) < z(x^{WS})$ (line 7), this implies that we need to add more cuts to $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$ to increase $\bar{z}(\bar{x}, \{\mathcal{V}_k^{sel}\}_{k \in [K]})$ above $z(x^{WS})$. To do so, we first calculate $\bar{z}(\bar{x}, \mathcal{V}_{DSP})$. If $\bar{z}(\bar{x}, \mathcal{V}_{DSP}) \geq z(x^{WS})$, then we know there are cuts in the DSP which can be added to $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$ to achieve the goal of

$$\bar{z}(\bar{x}, \{\mathcal{V}_k^{sel}\}_{k \in [K]}) \geq z(x^{WS}). \quad (12)$$

Indeed, this would be achieved by adding the dual solution from \mathcal{V}_{DSP} that achieves the maximum in (8) to \mathcal{V}_k^{sel} for each $k \in [K]$. However, we heuristically try to minimize the number of dual solutions that are added to achieve (12). We arrange the scenarios in decreasing order of values of $\bar{Q}_k(\bar{x}, \mathcal{V}_{DSP}) - \bar{Q}_k(\bar{x}, \mathcal{V}_k^{sel})$. This quantity tells us how much $\bar{Q}_k(\bar{x}, \mathcal{V}_k^{sel})$ will increase if we add the dual solution which achieves $\bar{Q}_k(\bar{x}, \mathcal{V}_{DSP})$ to \mathcal{V}_k^{sel} . For each scenario k in this order, we add a dual solution from \mathcal{V}_{DSP} that achieves the maximum in (8) to \mathcal{V}_k^{sel} , and stop as soon as we achieve (12).

If $\bar{z}(\bar{x}, \mathcal{V}_{DSP}) < z(x^{WS})$ (line 13), this implies the dual solutions in \mathcal{V}_{DSP} are not sufficient for adding to $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$ to achieve (12). Indeed, \bar{x} may even have a lower true objective than the current x^{WS} . In this case, we solve the subproblems (2) for \bar{x} and each scenario $k \in [K]$ to calculate $z(\bar{x})$. For each $k \in [K]$, we add to \mathcal{V}_k^{sel} and \mathcal{V}_{DSP} , the optimal dual solution from the subproblem. If $z(\bar{x}) < z(x^{WS})$, we update the current best solution to be \bar{x} . Once the algorithm converges, the final x^{WS} is provided as an initial feasible solution when solving the SAA replication.

Adaptive initialization as described is designed primarily for initializing Benders decomposition when solving stochastic IPs. For stochastic LPs, in preliminary experiments we found that too much time is spent doing this

Algorithm 6 Adaptive initialization - phase two.

```

1: Input:  $x^{WS}$  from phase one
2: Initialize  $\mathcal{V}_k^{sel} \leftarrow \emptyset$  for all  $k \in [K]$ 
3: Add  $\pi$  which achieves the maximum in (8) at  $x^{WS}$  to  $\mathcal{V}_k^{sel}$  for each  $k \in [K]$ 
4: Add active LP cuts to  $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$ 
5:  $z^{WS} \leftarrow \bar{z}(x^{WS}, \mathcal{V}_{DSP})$ 
6: while True do
7:   Evaluate  $\bar{z}(x, \{\mathcal{V}_k^{sel}\}_{k \in [K]})$  for all  $x \in X^{feas}$ 
8:    $\bar{x} \leftarrow \arg \min \{\bar{z}(x, \{\mathcal{V}_k^{sel}\}_{k \in [K]} : x \in X^{feas}\}$ 
9:   if  $\bar{z}(\bar{x}, \{\mathcal{V}_k^{sel}\}_{k \in [K]}) < z^{WS}$  then
10:    Compute  $\bar{z}(\bar{x}, \mathcal{V}_{DSP})$ 
11:    if  $\bar{z}(\bar{x}, \mathcal{V}_{DSP}) \geq z^{WS}$  then
12:      Add enough duals from DSP to  $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$  so that  $\bar{z}(\bar{x}, \{\mathcal{V}_k^{sel}\}_{k \in [K]}) \geq z^{WS}$ 
13:    else
14:      Solve (2) for  $x = \bar{x}$  and  $\xi = \xi_k$  for each  $k \in [K]$  to compute  $z(\bar{x})$ 
15:      Add the optimal dual solution from (2) to  $\mathcal{V}_k^{sel}$  and  $\mathcal{V}_{DSP}$  for each  $k \in [K]$ 
16:      if  $z(\bar{x}) < z(x^{WS})$  then
17:        Update  $x^{WS} \leftarrow \bar{x}, z^{WS} \leftarrow z(\bar{x})$ 
18:      end if
19:    end if
20:  else
21:    Break
22:  end if
23: end while
24: return  $x^{WS}, \{\mathcal{V}_k^{sel}\}_{k \in [K]}$ 

```

initialization process relative to the savings it yields in the eventual algorithm. Thus, for LPs, we make some changes to the adaptive initialization. In phase one, we evaluate the previously collected optimal solutions (X^{opt}) on the DSP. Let \bar{x} denote the solution with the lowest approximate objective value $\bar{z}(x, \mathcal{V}_{DSP})$. Rather than solving scenario subproblems to verify whether \bar{x} has the lowest true objective value (line 4 of Algorithm 5), we directly declare \bar{x} to be x^{WS} . Phase two begins with this solution and proceeds as in Algorithm 6, with line 4 skipped as it is not relevant for LPs. The next change is after the else condition on line 13, which is run when $\bar{z}(\bar{x}, \mathcal{V}_{DSP}) < z^{WS}$. In the LP case, we do not solve subproblems at this point. Instead, we update x^{WS} to \bar{x} , add the dual solution that achieves the maximum in (8) at x^{WS} to \mathcal{V}_k^{sel} for each $k \in [K]$, and then terminate the initialization.

3.4 Extensions

In this section, we describe how our ideas for accelerating Benders decomposition using information from previous replications can be adapted for solving problems without relatively complete recourse (Section 3.4.1) and for the single-cut version of Benders decomposition (Section 3.4.2).

3.4.1 Relatively Complete Recourse

In the absence of relatively complete recourse, the reformulation (6) that is the basis of Benders decomposition needs to be augmented with Benders feasibility cuts [4]: $(h(\xi_k) - T(\xi_k)x)^\top r \leq 0 \quad \forall r \in \mathcal{R}, k \in [K]$, where \mathcal{R} is the set of extreme rays of the dual feasible region Π .

The Benders decomposition algorithm is modified such that if the subproblem (2) is infeasible, then an extreme ray of the dual feasible region is identified and used to add a Benders feasibility cut to the main problem.

Since the dual feasible region remains fixed according to our assumptions that W and c are fixed, the set of rays also remains constant across replications. Thus, as we seek to solve a sequence of SAA replications, we can store the dual extreme rays that are identified in a DSP just as we do for Benders optimality cuts. When we obtain a main problem solution, we search the DSP for violated cuts (both optimality and feasibility cuts). If no violated cut is found for any scenario, we proceed by solving the subproblems to generate a feasibility or optimality cut. Curating the DSP follows the same principles as we have discussed in Section 3.2.

For the static and adaptive initialization method, we can initialize the Benders optimality cuts exactly as described previously. In adaptive initialization, if a scenario subproblem (2) is infeasible, then we identify an extreme ray of the dual feasible region and add the associated Benders feasibility cut as an initial cut. For adding Benders feasibility cuts, we would check the DSP to determine if any primal solutions from previous replications X^{feas} are violated by any of the associated feasibility cuts, and add at least one such cut for each solution in X^{feas} that violates one of these cuts.

3.4.2 Single-Cut

In single-cut implementation of Benders decomposition, we introduce a variable Θ which represents the expected value of the subproblem objective. The single-cut Benders decomposition algorithm is based on a reformulation that includes Benders cuts of the form

$$\Theta \geq \sum_{k=1}^K p_k (h(\xi_k) - T(\xi_k)x)^\top \pi_k, \quad (13)$$

where $\pi_k \in \mathcal{V}$ for each $k \in [K]$.

In standard Benders decomposition, given a main problem solution \hat{x} , a Benders cut of the form (13) is found by solving the scenario subproblem (2) for each scenario $k \in [K]$, and then using the dual solution π_k from subproblem k for each $k \in [K]$ to define the cut (13). The single-cut version uses fewer variables in the main problem and only adds one cut per iteration. Hence, the main problem typically is more compact and hence solves faster than in the multi-cut approach, but it often requires more iterations to reach optimality.

Our proposal for using the DSP and the curated DSP directly adapts to the single-cut version. For every scenario $k \in [K]$, at a solution x^t , we find the dual solution with maximum $\bar{Q}_k(x^t, \mathcal{V}_{DSP})$ by solving (8) and then aggregate them to generate a cut (13) using these dual solutions. If it is violated by x^t , we add it to the main problem and continue with the algorithm. Otherwise, we solve all scenario subproblems and generate a cut using those dual solutions (and update the DSP with the newly identified dual solutions). We expect that this would speed up the Benders iterations because the full set of scenario subproblems do not need to be solved at every iteration.

To adapt static initialization in this context, for each previously collected optimal solution, we identify dual solutions from the DSP that maximize the subproblem objective value by solving (8) for each $k \in [K]$. Aggregating these yields a cut for that primal solution, and repeating this procedure for all previously collected optimal solutions leads to a set of initial cuts that can be added to the main problem.

For adaptive initialization, phase one remains unchanged, where the primary goal is to identify the best solution (x^{WS}) to provide as an initial feasible solution to the branch-and-cut algorithm. We follow the steps in Algorithm 5 to do this. To adapt this method for single-cut, we also maintain $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$ in phase one. Now, whenever we evaluate any solution, say \bar{x} , on the DSP to estimate its value, the dual solutions for each scenario $k \in [K]$ that correspond to $\bar{Q}_k(\bar{x}, \mathcal{V}_{DSP})$ are stored in \mathcal{V}_k^{sel} . Furthermore, any dual solutions that are found when solving the scenario subproblems are also incorporated into $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$. In phase two, we follow the same steps as outlined in Algorithm 6. This algorithm outputs both an updated $\{\mathcal{V}_k^{sel}\}_{k \in [K]}$ and the initial candidate solution, x^{WS} . In the multi-cut version, we would add a Benders cut to the main problem for every dual solution in \mathcal{V}_k^{sel} for each scenario $k \in [K]$. For single-cut, the main problem is first initialized with a cut of the form (13), with the π_k dual solutions defined according to the optimal dual solution of subproblem (2) with $x = x^{WS}$ and $\xi = \xi_k$. To determine which additional cuts to initialize the main problem with, we iterate through each primal solution $\bar{x} \in X^{feas}$. For each such \bar{x} , we find the maximum value of the right-hand side of the currently added cuts on \bar{x} . Let's call this value $\hat{\Theta}(\bar{x})$. If $c^\top \bar{x} + \hat{\Theta}(\bar{x})$ exceeds $z(x^{WS})$, we do nothing as the current cuts are sufficient to ensure that the objective value of \bar{x} in the model is higher than the objective value of x^{WS} in the model. Otherwise, we find the dual solution which achieves $\bar{Q}_k(\bar{x}, \{\mathcal{V}_k^{sel}\}_{k \in [K]})$ for each scenario $k \in [K]$ and use these to define a cut. This cut is added to the main problem, ensuring that, as a result of phase two, the updated $c^\top \bar{x} + \hat{\Theta}(\bar{x})$ will be at least $z(x^{WS})$.

4 Computational Study

This section presents a comprehensive computational study to evaluate the effectiveness of our proposed information reuse strategies.

4.1 Experimental Setup and Implementation Details

We compare the following approaches for reusing information from previous solves of a replication:

- **Baseline:** This approach represents the standard Benders decomposition algorithm. The only information reused from previous replications is that for IPs an initial feasible solution is provided based on the optimal solution of the most recent replication, as described at the beginning of Section 3. No information is reused for LPs.
- **DSP:** This approach stores the dual solutions collected in previous SAA replications and uses them to generate cuts as described in Section 3.1.
- **Curated DSP:** This approach refines the DSP by maintaining a smaller pool of dual solutions as described in Section 3.2.
- **Static init:** This approach extends the curated DSP approach by initializing the algorithm with cuts generated through static initialization as described in Section 3.3.1.
- **Adaptive init:** This approach extends the curated DSP approach by initializing the algorithm with cuts generated through adaptive initialization as described in Section 3.3.2.

For the computational experiments, we solve 26 replications of problem (3), each with an independently drawn set of scenarios. The first replication is used for data collection and is identical for all compared methods. Thus, to compare the impact of different strategies for reusing information, all results presented in the following sections are based on the 25 SAA replications excluding the first one. Every replication is given a time limit of one hour for each method. To reduce the time required to run the experiments, we run less than the full 25 replications for the baseline method because it is significantly slower than the other methods. This method is only tested on the 2nd, 14th and 26th replications, and results reported are averaged over these three runs instead of the full 25 as in the other methods. Since all replications are independent and the only information used from previous replications is an initial feasible solution for IPs, and nothing for LPs, an average of a metric taken over this subset is expected to be a close approximation of the average over the full set of replications.

We implemented Benders decomposition in Python using Gurobi 10.0.1 as the optimization solver for both LPs and IPs. We build a main problem model and a single subproblem model which is updated with the current primal solution (\hat{x}) and scenario data (ξ) whenever we need to look for a cut, saving model building time and enabling warm-starting of the subproblems.

The DSP is implemented as follows: each dual solution π is stored in an indexed array and assigned a unique integer index for efficient retrieval. We use Python’s hash function to determine if a dual solution obtained after solving a subproblem is already present in the DSP or not. When searching the DSP for dual solutions that potentially generate a violated cut, we must efficiently evaluate problem (8) for each scenario $k \in [K]$: determine which dual solution

from \mathcal{V}_{DSP} provides the tightest Benders cut for a given first-stage solution x^t . To avoid repeatedly computing $\pi^\top h(\xi_k)$ for all dual solutions $\pi \in \mathcal{V}_{DSP}$ and scenarios $k \in [K]$ for every primal solution x^t , we precompute and store these terms in a matrix of dimension $|\mathcal{V}_{DSP}| \times K$ at the beginning of each SAA replication using a single matrix multiplication operation. When evaluating a candidate first-stage solution x^t , we compute $\pi^\top (-T(\xi_k)x^t)$ for all dual solutions simultaneously using vectorized matrix operations via the `linalg` library in NumPy, yielding a vector of length $|\mathcal{V}_{DSP}|$. We then solve problem (8) by combining this vector with the precomputed $\pi^\top h(\xi_k)$ terms and finding the maximum for each scenario using `numba` to efficiently calculate the arg max. As discussed in the beginning of Section 3, whenever we solve a new SAA replication (after the first), we provide the solver with the optimal primal solution of the previous replication as a candidate solution. To do this, we solve the scenario subproblems given the new scenario data and this candidate solution \hat{x} , and then provide this solution $(\hat{x}, \{Q(\hat{x}, \xi_k)\}_{k \in [K]})$ to the solver. This initialization is done for all methods except for the adaptive initialization, in which the method generates its own candidate solution. We do this initialization even for the baseline method in order to better illustrate the impacts of the other techniques we propose for reusing information.

We terminate Benders decomposition when the optimality gap percentage is less than 10^{-4} : $(U^t - L^t)/L^t * 100\% \leq 10^{-4}$. A candidate Benders cut of the form $\theta_k \geq \alpha_k - \beta_k x$ is considered violated by the current main problem solution $(\hat{x}, \{\hat{\theta}_k\}_{k \in [K]})$ if it satisfies

$$\hat{\theta}_k - Q(\hat{x}, \xi_k) \geq 10^{-5} \|(1, \alpha_k, \beta_k)\|.$$

This relative cut violation threshold ensures that the violation exceeds a small tolerance scaled by the norm of the cut's coefficients.

For the branch-and-cut method (Section 2.2), we implement the algorithm using a callback provided by the solver Gurobi. Whenever the algorithm finds an integer feasible solution, the callback is called to verify if this solution is feasible to the true problem. In this callback, if we are using the DSP, we first check the DSP for violated cuts. If no violated cuts are found in the DSP, we solve subproblems to check if any cuts are violated. If DSP is not employed, then we directly solve subproblems to check for violated cuts. Due to the presence of callbacks, we set the `lazyconstraints` parameter to 1, and that avoids reductions and transformations which are incompatible with lazy constraints.

For stochastic IPs, we first solve the LP relaxation of the problem via Benders decomposition. Benders cuts that are active after solving the LP relaxation are retained in the main problem and used as part of the formulation that is given to the solver when it starts the branch-and-cut algorithm. The initialization methods add cuts in addition to these cuts. If we deploy any initialization method for the IP, then the same method is also used to initialize the LP relaxation of the problem. Initialization methods are always used in conjunction with curated DSP to check for violated cuts. Only duals from curated DSP are considered to generate initialization cuts.

The experiments were run on two Intel Core i7 machines: an i7-9700 CPU at 3.00GHz and an i7-10700 CPU at 2.90GHz.

4.2 Test Problems

The study investigates the performance of all these methods on three problem classes: stochastic capacitated facility location, stochastic network design, and stochastic unit commitment. We describe these problems at a high level below. Appendix A provides the detailed formulation of each problem. All test instances have 400 scenarios, unless mentioned otherwise.

Capacitated Facility Location Problem (CFLP) The CFLP has a set of facilities and a set of customers with uncertain demands. The objective of the problem is to minimize the total expected cost of building and operating facilities while ensuring that customer demand is met. In the first-stage, we decide which facilities to open. Every facility has a setup cost and also a capacity. In the second-stage, we decide how to allocate goods from open facilities to satisfy customer demand as much as possible. Unmet demands are penalized, and thus this model has relatively complete recourse.

Our instances use the data from [14] and the extension of these to create stochastic programming instances from [17]. They create a stochastic variant by first generating the first-stage costs and capacities, followed by generating scenarios by sampling K demand vectors using the distributions defined in [14]. In Tables 1 and 2, we list the number of facilities and sets of customers we consider for the IP and LP instances, respectively. We use larger test instances for the LP instances to provide a more difficult test for that problem class.

Facilities	Customers
15	{105, 125, 215}
25	{95, 105, 185}
35	{105, 185}
55	{125}
75	{105}

Table 1: CFLP instance data for IPs.

Facilities	Customers
25	{305, 355, 405, 455, 495}
55	{305, 355, 405, 455, 495}
85	{305}

Table 2: CFLP instance data for LPs.

Multi Commodity Network Design Problem (CMND) The CMND problem is defined on a directed network comprising of nodes (N), arcs (A), and commodities (\mathcal{K}). Each commodity must be routed from an origin node to a destination node in the network. The arcs are characterized by installation costs and capacity. The objective is to determine a subset of arcs for installation with the goal of minimizing the expected total cost. In the first-stage, binary decisions

are made for each arc to decide if it will be installed or not. In the second-stage, after the demand for each commodity is revealed, routing decisions are made for how to route the realized commodity amounts in the network.

We use the test instances in [15]. The instances were originally proposed for the deterministic fixed charge capacitated multi-commodity network design problem [23]. To generate stochastic programming instances, we adopt the approach outlined in [25]. They use the techniques described in [45] to create random samples for the demands of various commodities. In each scenario, the demand of a commodity follows a normal distribution with the mean set to the demand in the deterministic instances and standard deviation of 0.1 times the mean. To generate instances, we start with base instances given in Tables 12 and 13. These tables give the number of nodes, arcs and commodities for each base instance. For each base instance, we run experiments on three versions of these instances that differ in the ratio of fixed costs to variable costs. These are given by r02.1, r02.2, and r02.3 in the actual dataset documentation for the II base instance.

Problem Set	$ N $	$ A $	$ \mathcal{K} $
II	10	35	25
III	10	35	50
IV	10	60	10

Table 3: CMND instance data for IPs.

Problem Set	$ N $	$ A $	$ \mathcal{K} $
VI	10	60	50
IX	10	83	50
X	20	120	40

Table 4: CMND instance data for LPs.

Stochastic Unit Commitment Problem (UC) The UC problem schedules a fleet of thermal generators over a time horizon to meet uncertain electricity demand while minimizing expected total cost. In the first-stage, binary decisions determine the on/off commitment status for each generator in each time period, subject to minimum up-time and down-time constraints that couple decisions across periods. In the second-stage, once demand is realized, continuous dispatch variables determine the power output levels for all committed units, subject to generator capacity limits and the system-wide demand balance constraint. To guarantee relatively complete recourse, unmet demand is permitted at each time period but penalized at a sufficiently high cost. For a comprehensive survey on mixed-integer programming formulations and solution methods for UC problems, we refer the reader to [28].

For computational experiments, we adopt the classical thermal UC instance generation procedure of Borghetti et al. [9], distributed through Beasley's OR-Library [1]. It produces single-bus UC instances with a 2-day horizon (48 hourly periods) and fleets whose size we vary between 10, 20, 30, and 40 thermal units. We use difficulty levels 2 and 3 in the instance generator, which differ primarily in the minimum up- and down-time requirements. Unlike CFLP and CMND, we use the same problem sizes for both LP and IP variants of UC. From preliminary experiments, we noticed that we can solve instances of the

same size for both LP and IP, as the LP relaxation is close to optimal for the IP and generates strong initial cuts.

To construct stochastic instances, we generate K demand scenarios by perturbing the base load \bar{B}_t at each time period t with independent, zero-mean Gaussian noise.

4.3 Results

Metric	Description
Total T	Time taken to optimize a SAA replication
LP T	Time taken to solve the LPs to optimality
Init T	Time taken to initialize the problem. For IPs, this includes selecting active cuts from the LP and finding other initial cuts
IP T	Time taken to solve the IPs to optimality, excluding LP T and Init T
Iterations	(LP only) # of iterations needed to solve the problem
SP count	# of times subproblems are solved to generate cuts
DSP T	Time taken to search for a violated cut in the DSP
SP T	Time taken to solve subproblems and find a violated cut
Cut T	Time taken to find violated cuts to add to the main problem
Nodes	(IP only) # of branch-and-bound nodes
Root gap	(IP only) Gap closed at the root node of the IP from initialization cuts
Callback calls	(IP only) # of calls to the callback to check an integer feasible solution

Table 5: Benders decomposition metrics.

Table 5 summarizes the metrics used to evaluate different methods. Each metric represents the arithmetic mean of the quantity calculated over the 25 SAA replications after the first replication. For results that are aggregated over multiple instances, like in Tables 6 - 13, the quantity presented is a shifted geometric mean over all instances of that test problem, with a shift of 1 being applied. Geometric mean is chosen as the relevant mean when summarizing results over different base instances because there might be a lot of variation in the values for different instances. Arithmetic mean is used for summarizing results across the 25 replications of an individual instance because we expect the values to remain more consistent over the 25 replications. All the time-related measurements are done in seconds.

For IPs, note that LP T represents the time spent solving the initial LP relaxation before proceeding with the IP. IP T notes the time taken to solve the IP after initializing, and Total T includes the time taken for this entire process. For each IP instance, $\text{Total T} = \text{Init T} + \text{LP T} + \text{IP T}$ ¹.

DSP T and SP T describe how much time is being spent to generate Benders cuts via the DSP and by solving subproblems. Cut T tells us the total time taken to find and add Benders cuts: for each instance, $\text{Cut T} = \text{DSP T} + \text{SP T}$.

¹ This equality holds on a per-instance basis, but does not hold for the summary statistics because we use geometric mean to summarize across instances.

4.3.1 LP Results

Method	Total T	Iterations	SP count	Cut T	DSP T	SP T	Init T
Baseline	213.9	81.2	81.2	187.4	-	187.4	-
DSP	62.1	66.4	8.0	50.8	36.0	14.3	-
Curated DSP	46.3	58.6	8.4	37.8	22.3	15.2	-
Static init	23.3	25.1	8.4	20.7	5.4	15.2	0.1
Adaptive init	22.2	20.7	8.5	18.7	3.2	15.4	1.3

Table 6: LP results: CFLP.

Method	Total T	Iterations	SP count	Cut T	DSP T	SP T	Init T
Baseline	616.9	131.4	131.4	531.2	-	531.2	-
DSP	89.7	75.2	10.0	70.2	26.3	40.3	-
Curated DSP	63.4	66.2	10.2	47.3	4.9	41.4	-
Static init	49.9	31.5	10.4	44.4	2.0	42.2	0.1
Adaptive init	51.2	21.9	10.5	44.0	1.2	42.7	3.6

Table 7: LP results: CMND.

Method	Total T	Iterations	SP count	Cut T	DSP T	SP T	Init T
Baseline	265.0	44.8	44.8	71.7	-	71.7	-
DSP	159.0	47.5	6.2	36.4	26.5	10.3	-
Curated DSP	135.9	46.9	6.9	22.6	11.1	11.4	-
Static init	56.7	33.6	7.2	16.4	4.6	12.0	0.8
Adaptive init	54.0	31.5	7.4	15.6	3.6	12.3	4.0

Table 8: LP results: UC.

Tables 6 - 8 display the summary results of the different methods for solving the LP test instances for CFLP, CMND, and UC problems, respectively. The tables demonstrate the significant benefits of information reuse techniques across all three problem classes. The introduction of DSP drastically reduces **SP count**, suggesting that we are usually able to find violated cuts in the DSP, and only occasionally need to solve subproblems to generate cuts. This validates the presence of valuable dual solutions within the DSP, and their effectiveness in generating violated cuts. The reduction in **SP count** directly contributes to savings in **SP T**, leading to the observed savings in **Total T**. A somewhat surprising result is the decrease in **iterations** of the algorithm from baseline to DSP for CFLP and CMND.

As hoped, we observe curated DSP reduces **DSP T** and hence leads to a reduction in **Cut T** and ultimately **Total T**. Intuitively, one might expect **SP**

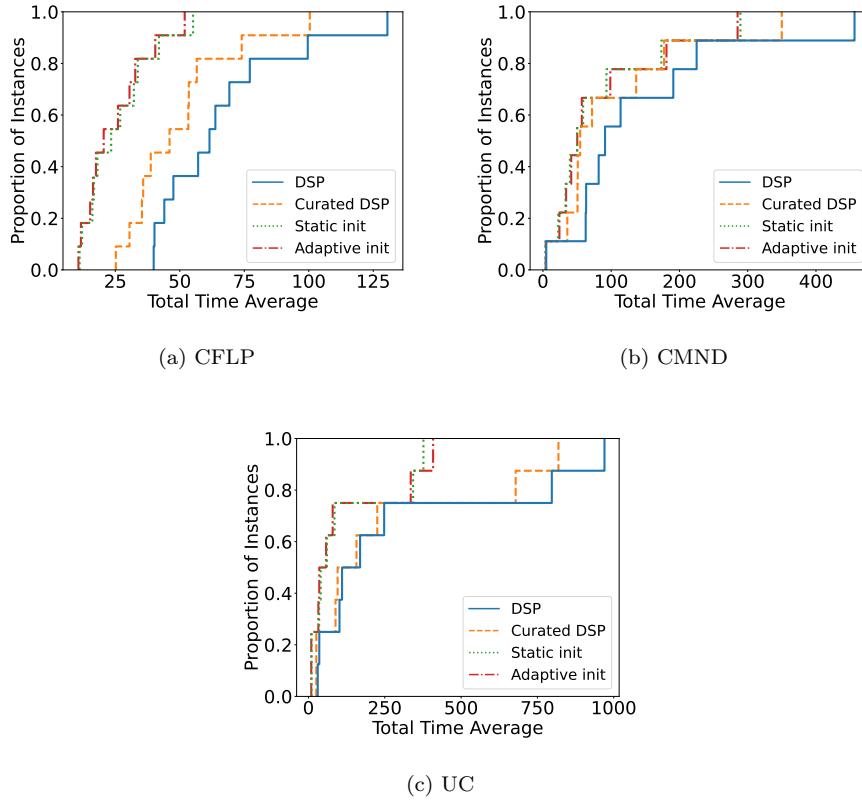


Fig. 1: Plots showing the fraction of solved LP instances over time for CFLP, CMND, and UC problems.

T to increase from DSP to curated DSP, as we have fewer dual solutions in the curated pool. However, the curation does not lead to a substantial increase in SP T suggesting that curated DSP obtains a good trade-off in the time saved from searching the DSP against the small extra time spent solving the subproblems

Both static and adaptive initialization methods consistently outperform baseline and DSP methods, needing fewer **iterations** to converge. This shows the value of initializing the algorithm with Benders cuts. This reduction in **iterations** directly contributes to these methods having the shortest overall **Total T**. Interestingly, adaptive initialization needs the fewest **iterations** to converge. This suggests that it is able to identify useful cuts. However, we do not see proportional decrease in **Total T** because adaptive initialization needs more time to find these initial cuts. Also, adaptive initialization usually adds more cuts in the main problem, leading to longer time to solve the main problem. Overall, using the combination of information reuse methods, we are able to solve CFLP and CMND problems approximately 10 times faster com-

pared to baseline on average, while for UC we observe approximately 5 times speedup. The more modest improvement for UC is because the subproblems are easier to solve. As a result, **Cut T** makes up a smaller share of **Total T**. Therefore, DSP and curated DSP, which primarily reduce **Cut T**, have limited impact on **Total T**. Hence, most of the improvement for UC comes from the initialization methods, rather than from DSP or curated DSP.

We complement the summary results presented in Tables 6 - 8 with the total time cumulative distribution function (CDF) plots shown in Figure 1. In these figures, the Y-axis shows the proportion of solved instances and the X-axis represents time. To focus on the relative improvements beyond just using DSP, these figures display only the four methods that reuse information and exclude the baseline method. Broadly, we observe that curated DSP improves significantly over DSP and that both initialization methods improve significantly over curated DSP, while they are comparable to each other.

4.3.2 IP Results

Method	Total T	IP T	LP T	Cut T	DSP T	SP T	Init T
Baseline	387.5	327.1	44.9	163.9	-	163.9	0.6
DSP	244.0	212.6	20.0	61.0	56.9	3.0	0.5
Curated DSP	216.9	191.4	14.7	46.2	41.4	3.8	0.5
Static init	194.1	182.4	6.1	43.9	39.2	3.8	0.6
Adaptive init	102.2	77.7	5.7	5.7	4.5	1.2	13.7

Table 9: IP results: CFLP - Part 1.

Method	Total T	IP T	LP T	Cut T	DSP T	SP T	Init T
Baseline	667.2	537.8	89.2	259.2	-	259.2	1.5
DSP	100.9	75.2	18.2	26.1	23.2	2.9	1.0
Curated DSP	61.8	45.8	9.8	10.5	7.0	3.4	1.0
Static init	58.5	47.2	6.5	9.4	6.3	2.9	1.1
Adaptive init	40.5	21.2	6.1	1.6	1.2	0.5	8.4

Table 10: IP Results: CMND - Part 1.

Tables 9 - 11 display the summary results of the different methods for solving the IP test instances for CFLP, CMND, and UC problems, respectively. For IPs, we solve the LP relaxation first to obtain a good initialization for branch-and-bound. Table 11 reveals that for UC, LP T dominates Total T. This computational profile differs from CFLP and CMND, where IP solving consumes the majority of time.

For CFLP and CMND, we find that DSP reduces SP T which translates into savings in Total T. Curated DSP reduces the DSP T because of a smaller

Method	Total T	IP T	LP T	Cut T	DSP T	SP T	Init T
Baseline	313.6	6.7	300.0	2.6	-	2.6	1.6
DSP	220.4	11.9	180.0	3.2	1.9	1.3	1.3
Curated DSP	194.5	11.1	145.2	1.9	0.6	1.2	1.3
Static init	101.6	10.8	65.5	1.8	0.5	1.3	1.9
Adaptive init	80.3	6.1	58.3	1.3	0.2	1.1	4.9

Table 11: IP results: UC - Part 1.

pool, in turn also helping to decrease **Total T**. Similar to the LP results, curated DSP does not lead to a significant increase in **SP T**, suggesting that curated DSP provides a good trade-off in the time checking the DSP against the time spent solving subproblems. For UC, **Cut T** remains similar across most methods because the tight LP relaxation requires few additional cuts during the IP solving phase.

For all the test problems, static and adaptive initialization methods consistently outperform baseline and DSP techniques. While the two initialization methods performed comparably for LPs, we observe a clear distinction for IPs. For CFLP, static initialization offers only a marginal improvement over curated DSP, whereas adaptive initialization demonstrates a substantial two-fold improvement in total time and a ten-fold reduction in cut generation time. This highlights the effectiveness of adaptive initialization in identifying strong initial cuts. Data for CMND shows a similar trend, with adaptive initialization again proving superior and static initialization providing only a marginal benefit over curated DSP. For UC, both initialization schemes significantly outperform the DSP methods, with adaptive initialization maintaining a slight edge over static initialization. For all the problem classes, although adaptive initialization has a higher initial computational cost (**Init T**), it delivers superior overall performance (**Total T**), making it the best choice for IPs.

Method	Nodes	Root gap (%)	Callback calls	SP count
Baseline	2814.5	2.9	115.7	115.7
DSP	2819.9	3.1	116.6	6.1
Curated DSP	2824.7	3.1	114.8	6.6
Static init	2645.7	3.1	111.0	6.5
Adaptive init	2129.8	2.8	13.4	1.9

Table 12: IP results: CFLP - Part 2.

Tables 12 - 14 present additional branch-and-bound metrics for all three problem classes. For CFLP and CMND, we observe significant reductions in **SP count** after employing DSP, similar to what we saw for LPs. Interestingly, deploying adaptive initialization leads to a significant reduction in both the **callback calls** and **SP count** compared to other methods. This suggests

Method	Nodes	Root gap (%)	Callback calls	SP count
Baseline	2404.9	12.7	140.1	140.1
DSP	1446.1	11.6	71.2	5.0
Curated DSP	1357.7	11.2	64.5	5.5
Static init	1420.7	9.3	62.5	5.2
Adaptive init	1379.0	3.9	10.8	1.2

Table 13: IP Results: CMND - Part 2.

Method	Nodes	Root gap (%)	Callback calls	SP count
Baseline	2.5	0.08	3.2	3.2
DSP	9.3	0.08	3.3	2.2
Curated DSP	8.3	0.08	3.6	2.1
Static init	8.6	0.08	3.2	2.2
Adaptive init	2.0	0.08	2.3	2.0

Table 14: IP Results: UC - Part 2.

that the initial cuts generated by adaptive initialization provide a strong relaxation of the original problem, minimizing the need for looking for additional cuts during the optimization process. Furthermore, adaptive initialization improves root gaps, particularly for CMND. Since `root gaps` help isolate the impact of initialization, this demonstrates that adaptive initialization closes the most gap among all methods. The resulting reduction in `callback calls` confirms that the initial cuts are highly effective. For UC, Table 14 shows extremely low root gaps (0.07%) and very few branch-and-bound node counts, confirming the near-integrality of the LP relaxation. This indicates that once the LP is solved, minimal additional work is required to obtain the IP solution. This behavior reinforces our observation that UC IP performance is dominated by LP solving.

Figure 2 presents the cumulative distribution function (CDF) plots of solution times for IP instances, again displaying the results only for the methods that reuse information. These plots further illustrate that the adaptive initialization method has the best performance. We also see that curated DSP generally leads to an improvement in total time compared to regular DSP.

4.3.3 Solution Quality with Equal Time Budgets

The results presented above demonstrate significant speedups in solution times when using our information reuse methods, particularly adaptive initialization. As these results focus on time to solve to optimality, they do not give a picture of how much improvement in solution quality and gap is achieved earlier in the solution process. To investigate this, we conduct additional experiments comparing solution quality obtained by the baseline method when limited to a time comparable to the time it takes the adaptive method to solve to optimality.

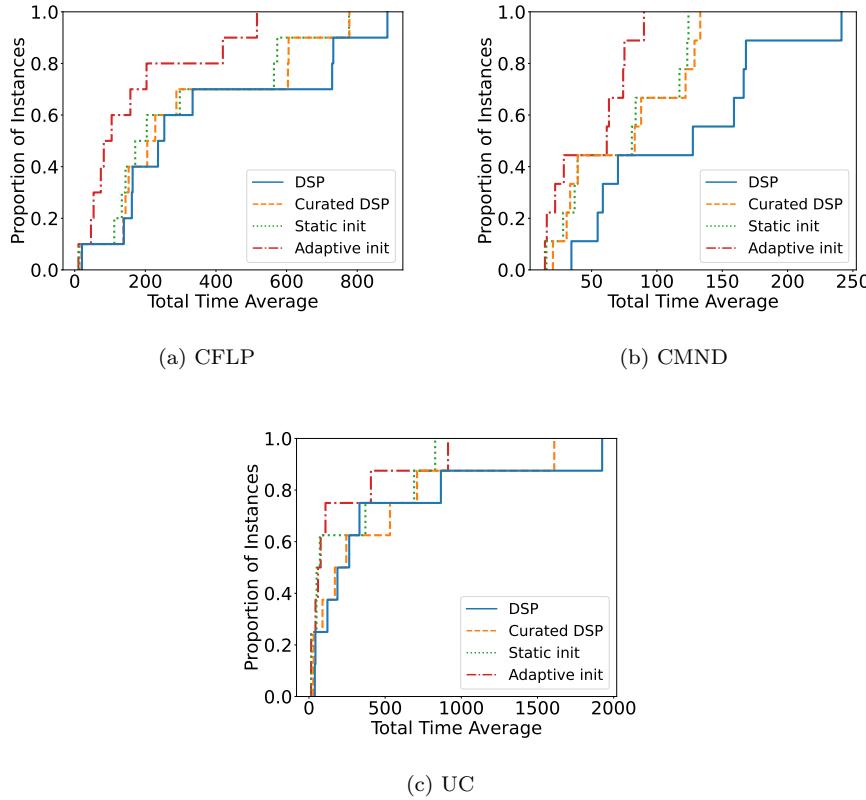


Fig. 2: Plots showing the fraction of solved LP instances over time for CFLP, CMND, and UC problems.

For each test instance, we first ran the adaptive initialization method on the 25 replications after the first one and recorded its average solution time T . We then ran the baseline method on the first replication of the same instance with two time limits: T and $2T$, recording whether it proved optimality and, if not, the final optimality gap at termination.

Tables 15 and 16 present the average optimality gaps across instances. The **Gap - T (%)** column shows the optimality gap when the baseline is limited to the time limit T , while **Gap - 2T (%)** shows the gap when the baseline is given twice that time limit. All runs hit the solver timeout without proving optimality.

These results confirm that the baseline's longer solution times are not simply due to closing a small amount of remaining optimality gap. Even when run for twice the time it takes the adaptive method to solve the instances to optimality, the baseline fails to achieve the solution quality that the adaptive initialization obtains. For IP instances, the gaps are smaller but still signif-

Problem	Gap - T (%)	Gap - 2T (%)
CMND	12474.8	4061.8
CFLP	15.0	6.5
UC	9.4	1.0

Table 15: LP relaxation: Average optimality gaps with time limits.

Problem	Gap - T (%)	Gap - 2T (%)
CMND	11.1	10.7
CFLP	3.5	2.4
UC	-	-

Table 16: IP: Average optimality gaps with time limits.

icant, demonstrating that our methods provide fundamental algorithmic improvements across both LP and IP formulations.

Note that for some IP instances, the baseline did not complete the initial LP relaxation within the time limit. For UC-IP, all instances timeout on the LP relaxation for both time limits T and $2T$, hence no gap values are reported in Table 16. For CMND-IP, four instances did not finish the LP relaxation within time T and 2 instances did not finish within $2T$. For CFLP-IP, one instance did not finish the LP relaxation within T , but all instances completed it within $2T$. The gap averages reported in Table 16 are calculated only for instances that completed the LP relaxation, as a gap cannot be computed otherwise.

4.3.4 Cut Distribution

Method	Initial cuts	SP cuts	DSP cuts	Total cuts
Baseline	-	39192	-	39192
DSP	-	448	33455	33903
Curated DSP	-	589	34139	34728
Static Init	784	608	33327	34719
Adaptive Init	13105	128	3610	16843

Table 17: Cut distribution for the CFLP instance with 35 facilities and 105 customers.

Our algorithm generates Benders cuts from three sources: cuts provided during initialization (**Initial cuts**), cuts derived by solving subproblems during the algorithm (**SP cuts**), and cuts obtained by searching a pool of dual solutions (**DSP cuts**). To provide more insight about our methods, we present in Tables 17 - 19 the distribution of cuts of each type for sample IP instances of the CFLP, CMND, and UC problems, respectively. To focus on differences between the methods, the **Initial cuts** and **Total cuts** exclude active LP cuts, which are nearly identical for all methods.

Method	Initial cuts	SP cuts	DSP cuts	Total cuts
Baseline	-	55721	-	55721
DSP	-	85	29851	29936
Curated DSP	-	128	27604	27732
Static Init	782	50	26455	27287
Adaptive Init	8723	24	3629	12376

Table 18: Cut distribution for the CMND instance r03.3.

Method	Initial cuts	SP cuts	DSP cuts	Total cuts
Baseline	-	667	-	667
DSP	-	116	1223	1339
Curated DSP	-	16	1011	1027
Static Init	726	48	161	935
Adaptive Init	1520	0	176	1696

Table 19: Cut distribution for the UC instance with 20 generators and difficulty 3.

We find that the inclusion of DSP in the Benders decomposition framework significantly reduces the number of subproblems solved to generate cuts, leading to a dramatic decrease in **SP cuts** across all problem types. This is expected and confirms the presence of useful dual solutions within the pool, capable of generating violated cuts. For CFLP and CMND, we observe that curated DSP requires slightly more **SP cuts** than DSP, which can be attributed to the reduced size of the curated DSP compared to the DSP. However, for the UC instance, curated DSP reduces **SP cuts** compared to DSP. Using either of the initialization methods leads to a further reduction in the number of subproblem solves required to generate cuts relative to the curated DSP. For CFLP and CMND, we observe a drastic reduction in the number of total cuts needed to reach the optimal solution when using adaptive initialization, which is explained by the significant decrease in the number of cuts added from the DSP. The adaptive initialization method requires very few additional cuts beyond the initial set introduced during problem initialization, which translates to significant time savings in solving the SAA replication. In contrast, for UC, the number of cuts added beyond the LP relaxation is relatively small across all methods, whether through initialization or by solving subproblems. This suggests that for UC instances, the LP relaxation cuts provide most of the constraints needed to get to the optimal IP solution, with initialization and DSP methods offering more modest incremental improvements.

4.3.5 Impact of Adaptive Initialization

Analysis of the cut distribution in Tables 17 - 19 reveals that adaptive initialization consistently introduces a higher number of initial cuts compared to static initialization. However, for the CFLP and CMND instances, the difference in the number of initial cuts between static and adaptive initialization

is more drastic compared to the UC instance. Therefore, to isolate the impact of cut quality from quantity on the observed performance improvements, we focus our analysis on CFLP and CMND and conduct an experiment in which we modify the static initialization method to add the same number of cuts as adaptive initialization for these two problem types. In this version of static initialization, which we refer to as boosted static initialization, we solve the LP relaxation first as usual, and then use static initialization to add $\lceil \frac{n}{K} \rceil$ cuts per scenario to the main problem, where n is the number of cuts that was added by the adaptive initialization method for the same instance. For each scenario within the new SAA replication, we identify the top $\lceil \frac{n}{K} \rceil$ dual solutions from the DSP that lead to the highest value of the subproblem objective value function, $\bar{Q}_k(\bar{x}, \mathcal{V}_{DSP})$, evaluated at the first two optimal solutions in X^{opt} . The cuts generated from these are then used to initialize the problem. This ensures both initialization approaches use the same number of cuts.

In Fig. 3, we track the root node gap closed and also the average time taken to solve the i^{th} SAA replication over 26 replications using boosted static initialization and adaptive initialization. The results demonstrate that adaptive initialization remains superior even when the static initialization adds the same number of cuts. Thus, we conclude that the adaptive nature of adaptive initialization is important, and in particular it seems to benefit from allowing the number of cuts added for each scenario to vary.

4.3.6 Impact of Curated DSP

In Figure 4, we plot the size of the DSP and curated DSP as we solve more SAA replications. For CMND and UC, we notice that, the DSP constantly increases in size, but the curated DSP has a sharp decrease in size in the second replication. Although it begins to grow again after that, the increase is not as significant compared to the DSP. This indicates that we have successfully achieved our goal for the curated DSP, as we are able to maintain a controlled size of the pool. In the case of CFLP, we notice similar trends, although the reduction in size from the first to the second replication is not as pronounced as in CMND. Additionally, by the end of the 25 replications, the size of the pool is nearly the same as it was after the first replication.

We next investigate the importance of our particular mechanism for choosing the dual solutions to keep in the curated DSP. To do this, we compared our method to a baseline that randomly selects the same number of dual solutions as were chosen in our curated DSP method. We refer to this method as the random curated DSP. In Table 20, we display the average total time taken

Method	CFLP	CMND	UC
Curated DSP	153.17	89.56	228.42
Random curated DSP	152.28	92.19	286.46

Table 20: Total time comparison for curated DSP and random DSP.

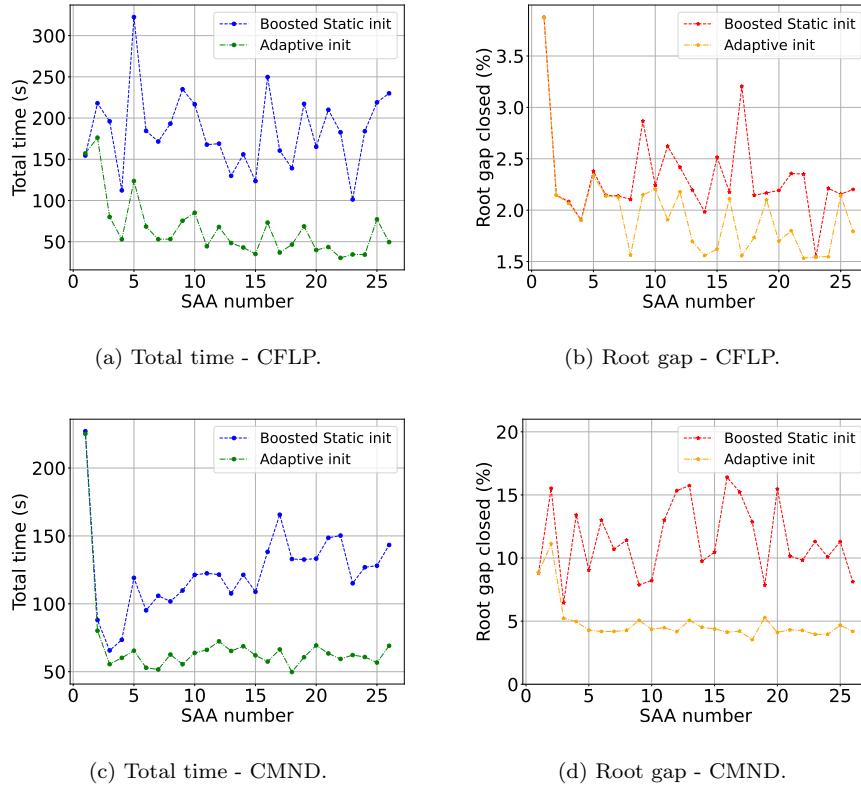


Fig. 3: Total time and root gap trends for boosted static initialization and adaptive initialization per replication. The CFLP results (top plots) are for the instance with 35 facilities and 105 customers. The CMND results (bottom plots) are for the r03.3 instance.

to solve 25 SAA replications after the first one with these two methods. We use the same instances to test these methods as the experiments on cut distribution in Sec. 4.3.4. For the UC instances, curated DSP achieves a notable improvement over random curated DSP. This suggests that for UC, the smart selection mechanism of curated DSP successfully identifies more effective dual solutions compared to random selection. In contrast, for CFLP and CMND instances, the performance is comparable between the two approaches, with no significant difference in solution times. These results indicate that the primary benefit of curated DSP comes from maintaining a smaller pool size. However, for some problem classes like UC, the intelligent selection mechanism provides additional computational gains by identifying more effective dual solutions. Importantly, a key feature of our method is that it automatically determines the number of dual solutions to retain in the curated DSP without requiring this number as input.

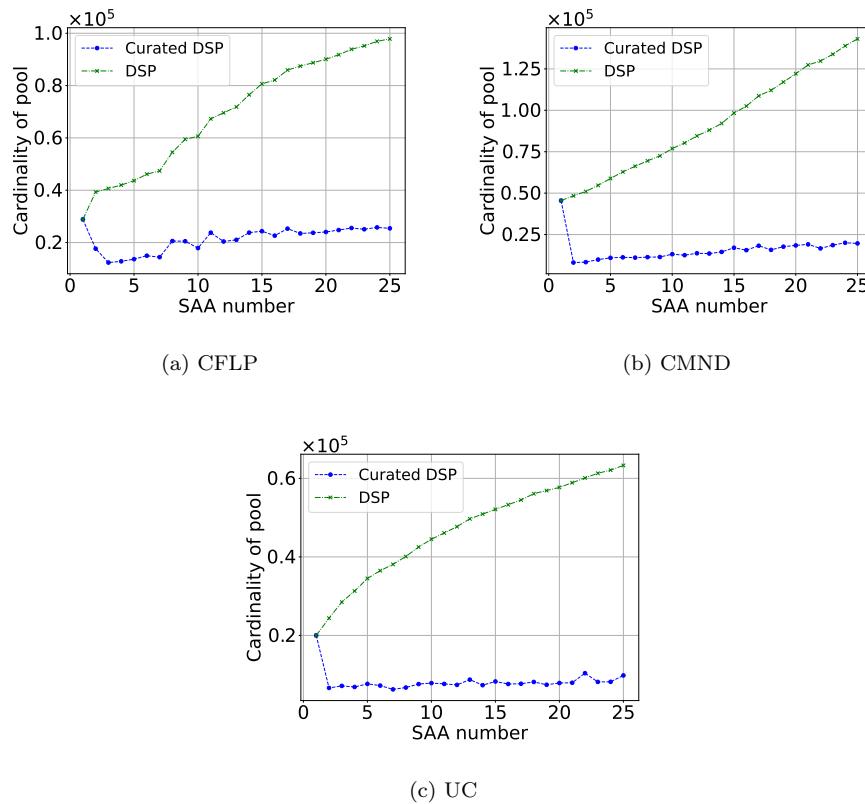


Fig. 4: Number of dual solutions in the pool for DSP and curated DSP.

4.3.7 Performance with Varying Number of Scenarios

All the previous experiments have been done on instances in which the number of scenarios in each replication is 400. We next investigate the impact of varying the number of scenarios in the replications on our conclusions.

For this experiment, we select one instance for each problem class and then run our methods for that instance with the number of scenarios as 200, 400, and 800. For network design, we use instance r09.2 for the LP case and r03.3 for the IP case. For facility location, we select the instance with 55 facilities and 495 customers for LP, while the IP instance has 35 facilities and 105 customers. For UC, we use an instance with 30 generators and difficulty 3.

Figure 5 displays the relative performance of our information reuse methods as a fraction of the execution time of the baseline method, across three scenario sets. We find that in general the relative performance of the different methods is similar across the different scenario sizes. We also find that the improvements from reusing information tend to be more significant for instances with more scenarios, which is not surprising considering that in such

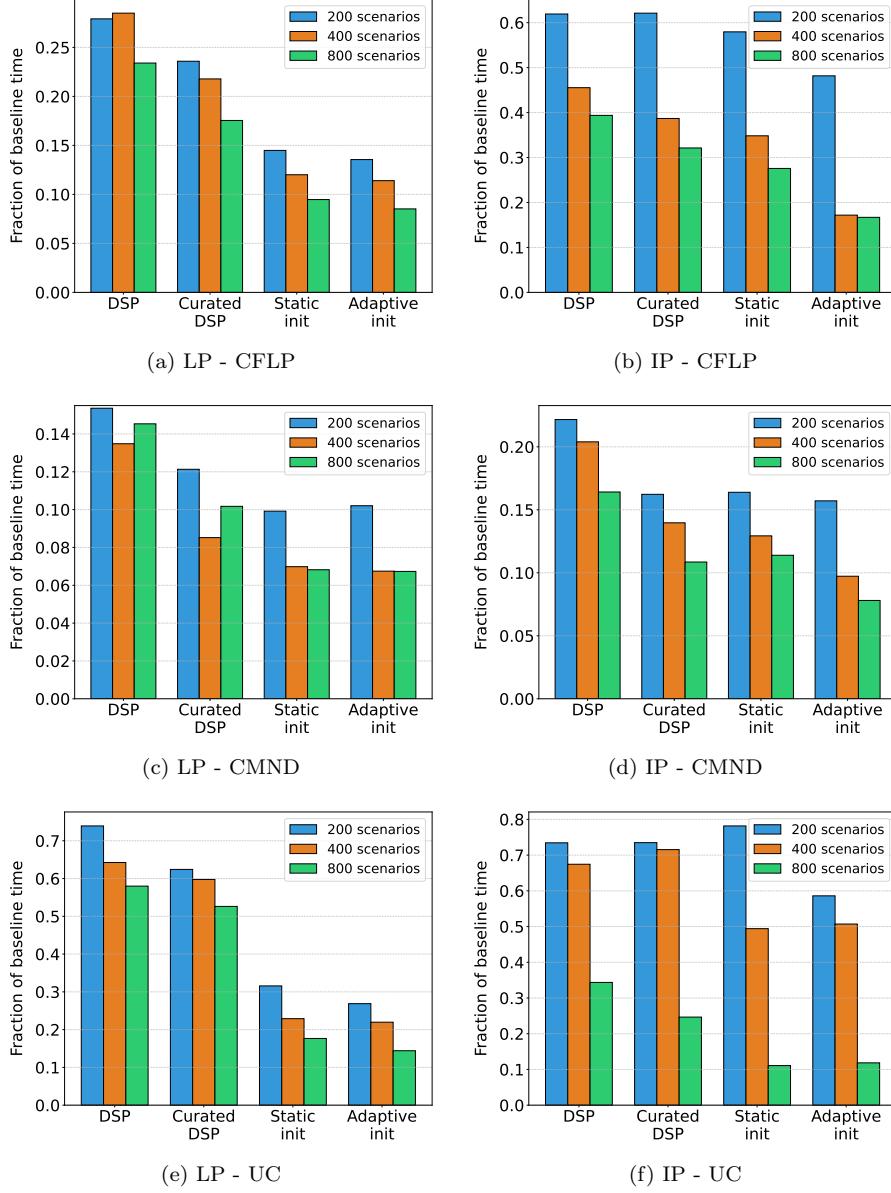


Fig. 5: Comparison of performance of methods as number of scenarios vary.

instances there is more work that has to be done in solving subproblems, and hence more opportunity to save time on that work with the proposed DSP and initialization techniques.

Method	Total T	Iterations	SP count	Cut T	DSP T	SP T	Init T
Baseline	621.7	374.1	374.1	621.1	-	621.1	-
DSP	406.4	207.5	10.6	401.0	380.4	20.3	-
Curated DSP	138.1	172.9	12.8	132.8	106.9	24.4	-
Static Init	130.5	151.2	13.8	123.8	95.4	27.2	1.0
Adaptive Init	68.7	33.9	9.9	35.4	16.0	19.3	27.8

Table 21: Single-cut LP results: CFLP.

Method	Total T	IP T	LP T	Cut T	DSP T	SP T	Init T
Baseline	263.9	176.3	86.8	176.1	-	176.1	0.4
DSP	230.2	167.1	60.5	166.9	164.3	2.2	0.6
Curated DSP	111.6	81.0	28.7	80.8	78.1	2.4	0.6
Static Init	99.2	74.4	22.0	74.1	71.4	2.4	1.0
Adaptive Init	76.5	9.5	5.4	9.1	7.2	1.6	51.6

Table 22: Single-cut IP results: CFLP - Part 1.

4.4 Single-Cut Benders Decomposition

As described in Section 3.4.2, our proposed information reuse strategies can be adapted to the single-cut variant of Benders decomposition. To evaluate whether these adaptations remain effective in the single-cut setting, we conducted experiments on CFLP instances for both LP and IP variants using the same experimental setup as described previously, with the key difference being that we use single-cut Benders decomposition to solve each SAA replication instead of multi-cut. In all experiments, we impose a time limit of 3600 seconds per replication. We do not report results for CMND or UC instances with single-cut Benders decomposition, as our experiments found the single-cut method to be impractical for these test instances. For our computational experiments on the CFLP, we selected a subset of instances from Tables 1 and 2. Specifically, for the LP variant, we tested instances with (warehouses, facilities) dimensions of (25, 305), (25, 495), and (55, 495), and for the IP variant, we tested instances with dimensions (15, 105), (15, 215), and (25, 95).

Table 21 presents results for CFLP-LP instances using single-cut Benders decomposition. The high iteration count in single-cut results in the discovery of a large number of dual solutions throughout the solution process, causing the DSP to grow considerably. Consequently, the DSP provides only marginal improvement over the baseline, as most of the time is spent searching through the large DSP. In contrast, the curated DSP approach maintains a smaller set of high-quality dual solutions, substantially reducing DSP search time. This demonstrates the importance of curation when a large number of dual solutions are collected. The initialization methods provide further improvements. Notably, adaptive initialization achieves a dramatic reduction in Benders iterations compared to the baseline, solving far fewer subproblems on average.

Method	Nodes	Root gap (%)	Callback calls	SP count
Baseline	4233.2	22.9	276.1	276.1
DSP	4020.5	16.8	272.6	3.2
Curated DSP	3842.9	14.3	265.0	3.5
Static Init	3518.5	9.1	252.9	3.5
Adaptive Init	1642.7	7.7	25.1	2.3

Table 23: Single-cut IP results: CFLP - Part 2.

Table 22 and Table 23 present complementary results for CFLP-IP instances using single-cut Benders decomposition. The pattern observed for LP instances is replicated here: the uncurated DSP provides modest improvement, while curated DSP achieves more substantial gains. Static initialization further improves performance, and adaptive initialization again delivers the best results. While adaptive initialization requires a significant initialization investment, it dramatically reduces both the LP relaxation time and the IP phase time.

Table 23 provides additional perspective through branch-and-bound metrics. The information reuse methods progressively improve the root gap, with adaptive initialization achieving the lowest root gap by adding high-quality cuts that effectively strengthen the problem formulation before branch-and-bound begins. This improvement in root gap quality translates directly to reduced node counts and callback invocations, with adaptive initialization exploring substantially fewer nodes and requiring far fewer callback invocations with minimal subproblem solves per callback.

These results demonstrate that our information reuse strategies remain highly effective for single-cut Benders decomposition across both LP and IP problem variants when applied to this problem class.

5 Conclusion and Future Directions

We presented methods to accelerate solving a sequence of SAA replications in two-stage stochastic programming, assuming randomness only in the right-hand sides of the subproblems. These methods are derived for Benders decomposition as the solution algorithm and we find that, for our test problems, it is possible to reduce the time to solve the replications after the first one by a factor of 10 for both stochastic LP and IP problems by using the information reuse techniques we have proposed.

One significant direction for future research is to consider problems in which the subproblems have uncertainty in either the objective coefficients or the recourse matrix. Such problems do not have the property that the dual feasible region is fixed across scenarios, and hence the techniques we proposed do not extend directly to such problems.

This paper focused on the Benders decomposition algorithm, as it is a leading algorithm for both two-stage stochastic LPs and for IPs with contin-

uous recourse. Future research could investigate techniques for accelerating different algorithms in this context of solving a sequence of SAA replications. For example, for two-stage stochastic LPs, the level method [32, 20] often performs better than Benders decomposition. We anticipate that the techniques presented here would be useful for accelerating this and other cut-based decomposition methods, but testing this hypothesis would be an interesting direction for future work. It would also be interesting to explore methods to accelerate alternative methods for solving two-stage stochastic IPs, such as dual decomposition [11] and methods that use different types of cuts [8, 13, 22, 33, 39, 46].

References

1. OR-library: distributing test problems by electronic mail. *Journal of the operational research society* **41**(11), 1069–1072 (1990)
2. Adulyasak, Y., Cordeau, J.F., Jans, R.: Benders decomposition for production routing under demand uncertainty. *Operations Research* **63**(4), 851–867 (2015)
3. Bayraksan, G., Morton, D.P.: Assessing solution quality in stochastic programs. *Mathematical Programming* **108**, 495–514 (2006)
4. Benders, J.: Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik* **4**, 238–252 (1962). URL <https://doi.org/10.1007/BF01386316>
5. Bengio, Y., Frejinger, E., Lodi, A., Patel, R., Sankaranarayanan, S.: A learning-based algorithm to quickly compute good primal solutions for stochastic integer programs. In: *Integration of Constraint Programming, Artificial Intelligence, and Operations Research: 17th International Conference, CPAIOR 2020, Vienna, Austria, September 21–24, 2020, Proceedings* 17, pp. 99–111. Springer (2020)
6. Bengio, Y., Lodi, A., Prouvost, A.: Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research* **290**(2), 405–421 (2021)
7. Birge, J.R., Louveaux, F.: *Introduction to stochastic programming*. Springer Science & Business Media (2011)
8. Boland, N., Christiansen, J., Dandurand, B., Eberhard, A., Linderoth, J., Luedtke, J., Oliveira, F.: Combining progressive hedging with a Frank–Wolfe method to compute Lagrangian dual bounds in stochastic mixed-integer programming. *SIAM Journal on Optimization* **28**(2), 1312–1336 (2018)
9. Borghetti, A., Frangioni, A., Lacalandra, F., Nucci, C.A.: Lagrangian heuristics based on disaggregated bundle methods for hydrothermal unit commitment. *IEEE Transactions on Power Systems* **18**(1), 313–323 (2003)
10. Borozan, S., Giannelos, S., Falugi, P., Moreira, A., Strbac, G.: Machine learning-enhanced benders decomposition approach for the multi-stage stochastic transmission expansion planning problem. *Electric Power Systems Research* **237**, 110985 (2024). DOI <https://doi.org/10.1016/j.epsr.2024.110985>. URL <https://www.sciencedirect.com/science/article/pii/S0378779624008708>
11. Carøe, C.C., Schultz, R.: Dual decomposition in stochastic integer programming. *Operations Research Letters* **24**(1-2), 37–45 (1999)
12. Chan, T., Lin, B., Saxe, S.: A machine learning approach to solving large bilevel and stochastic programs: Application to cycling network design. Available at SSRN 4592562 (2023)
13. Chen, R., Luedtke, J.: On generating lagrangian cuts for two-stage stochastic integer programs. *INFORMS Journal on Computing* **34**(4), 2332–2349 (2022)
14. Cornuéjols, G., Sridharan, R., Thizy, J.M.: A comparison of heuristics and relaxations for the capacitated plant location problem. *European journal of operational research* **50**(3), 280–297 (1991)

15. Crainic, T.G., Fu, X., Gendreau, M., Rei, W., Wallace, S.W.: Progressive hedging-based metaheuristics for stochastic network design. *Networks* **58**(2), 114–124 (2011)
16. Deza, A., Khalil, E.B.: Machine learning for cutting planes in integer programming: A survey. arXiv preprint arXiv:2302.09166 (2023)
17. Dumouchelle, J., Patel, R., Khalil, E.B., Bodur, M.: Neur2sp: Neural two-stage stochastic programming. arXiv preprint arXiv:2205.12006 (2022)
18. Dupačová, J., Popela, P.: Melt control: Charge optimization via stochastic programming. In: Applications of stochastic programming, pp. 277–297. SIAM (2005)
19. Dyer, M., Stougie, L.: Computational complexity of stochastic programming problems. *mathematical programming* **106**, 423–432 (2006)
20. Fábián, C.I., Szőke, Z.: Solving two-stage stochastic programming problems with level decomposition. *Computational Management Science* **4**, 313–353 (2007)
21. Fortz, B., Poss, M.: An improved benders decomposition applied to a multi-layer network design problem. *Operations research letters* **37**(5), 359–364 (2009)
22. Gade, D., Küçükyavuz, S., Sen, S.: Decomposition algorithms with parametric Gomory cuts for two-stage stochastic integer programs. *Mathematical Programming* **144**(1-2), 39–64 (2014)
23. Ghamlouche, I., Crainic, T.G., Gendreau, M.: Cycle-based neighbourhoods for fixed-charge capacitated multicommodity network design. *Operations research* **51**(4), 655–667 (2003)
24. Higle, J.L., Sen, S.: Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of operations research* **16**(3), 650–669 (1991)
25. Jia, H., Shen, S.: Benders cut classification via support vector machines for solving two-stage stochastic programs. *INFORMS Journal on Optimization* **3**(3), 278–297 (2021)
26. Kim, S., Pasupathy, R., Henderson, S.G.: A guide to sample average approximation. *Handbook of simulation optimization* pp. 207–243 (2015)
27. Kleywegt, A.J., Shapiro, A., Homem-de Mello, T.: The sample average approximation method for stochastic discrete optimization. *SIAM Journal on optimization* **12**(2), 479–502 (2002)
28. Knueven, B., Ostrowski, J., Watson, J.P.: On mixed-integer programming formulations for the unit commitment problem. *INFORMS Journal on Computing* **32**(4), 857–876 (2020)
29. Lam, H., Qian, H.: Bounding optimality gap in stochastic optimization via bagging: Statistical efficiency and stability. arXiv preprint arXiv:1810.02905 (2018)
30. Larsen, E., Freijinger, E., Gendron, B., Lodi, A.: Fast continuous and integer l-shaped heuristics through supervised learning. *INFORMS Journal on Computing* (2023)
31. Lee, M., Ma, N., Yu, G., Dai, H.: Accelerating generalized benders decomposition for wireless resource allocation. *IEEE Transactions on Wireless Communications* **20**(2), 1233–1247 (2020)
32. Lemaréchal, C., Nemirovskii, A., Nesterov, Y.: New variants of bundle methods. *Mathematical programming* **69**, 111–147 (1995)
33. Lulli, G., Sen, S.: A branch-and-price algorithm for multistage stochastic integer programming with application to stochastic batch-sizing problems. *Management Science* **50**(6), 786–796 (2004)
34. Mak, W.K., Morton, D.P., Wood, R.K.: Monte carlo bounding techniques for determining solution quality in stochastic programs. *Operations research letters* **24**(1-2), 47–56 (1999)
35. Mitrai, I., Daoutidis, P.: Learning to initialize generalized benders decomposition via active learning. FOCAPO/CPC, San Antonio, Texas (2023)
36. Mitrai, I., Daoutidis, P.: Computationally efficient solution of mixed integer model predictive control problems via machine learning aided benders decomposition. *Journal of Process Control* **137**, 103207 (2024)
37. Nair, V., Dvijotham, D., Dunning, I., Vinyals, O.: Learning fast optimizers for contextual stochastic integer programs. In: UAI, pp. 591–600 (2018)
38. Naoum-Sawaya, J., Elhedhli, S.: An interior-point benders based branch-and-cut algorithm for mixed integer programs. *Annals of Operations Research* **210**, 33–55 (2013)
39. Rahmanni, R., Crainic, T.G., Gendreau, M., Rei, W.: The benders decomposition algorithm: A literature review. *European Journal of Operational Research* **259**(3), 801–817 (2017)

40. Robbins, H., Monro, S.: A stochastic approximation method. *The Annals of Mathematical Statistics* pp. 400–407 (1951)
41. Saharidis, G.K., Boile, M., Theofanis, S.: Initialization of the benders master problem using valid inequalities applied to fixed-charge network problems. *Expert Systems with Applications* **38**(6), 6627–6636 (2011)
42. Sakhavand, N., Gangammanavar, H.: Subproblem sampling vs. scenario reduction: efficacy comparison for stochastic programs in power systems applications. *Energy Systems* pp. 1–29 (2022)
43. Sen, S., Doverspike, R.D., Cosares, S.: Network planning with random demand. *Telecommunication systems* **3**(1), 11–30 (1994)
44. Sen, S., Liu, Y.: Mitigating uncertainty via compromise decisions in two-stage stochastic linear programming: Variance reduction. *Operations Research* **64**(6), 1422–1437 (2016)
45. Song, Y., Luedtke, J.R., Küçükyavuz, S.: Chance-constrained binary packing problems. *INFORMS Journal on Computing* **26**(4), 735–747 (2014)
46. van der Laan, N., Romeijn, W.: A converging Benders' decomposition algorithm for two-stage mixed-integer recourse models. *Operations Research* **72**, 2190–2214 (2023)
47. Van Slyke, R.M., Wets, R.: L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM journal on applied mathematics* **17**(4), 638–663 (1969)
48. Wang, S., Gangammanavar, H., Ekşioğlu, S., Mason, S.J.: Statistical estimation of operating reserve requirements using rolling horizon stochastic optimization. *Annals of Operations Research* **292**, 371–397 (2020)

Appendix A

We provide detailed descriptions of the problems considered in our computational study. We adopt the formulations and problem descriptions from [25] for both test problems.

A.1 Capacitated Facility Location Problem

Consider a set F of facilities, where a facility $i \in F$ has a setup cost f_i and a production capacity limit u_i . Additionally consider a set of customers C , where each customer j has an uncertain demand denoted by \tilde{d}_j . The vector of realizations of this uncertain demand in a scenario $k \in [K]$ is denoted by $d_k = [d_{kj} : j \in J]$. This demand can be met by shipping from any open facility i to customer j at a unit transportation cost c_{ij} . Any unmet demand for customer j incurs a lost-sale penalty, with a unit cost ρ_j . The objective is to select a subset of facilities to open in order to minimize the total expected cost.

This problem is modeled as a two-stage stochastic programming problem. The first-stage binary decision variables x_i indicate whether the facility i is opened or not. In the second-stage, after revelation of the demands, we introduce continuous decision variables $y_{ij} \geq 0, \forall i \in F, j \in C$, which represent goods transported from facility i to customer j . The model aims to find the best decisions to minimize the sum of facility setup cost, expected transportation cost, and expected lost-sale cost. The first-stage formulation is given by:

$$\begin{aligned} \min_x \sum_{i \in F} f_i x_i + \sum_{k \in [K]} p_k Q(x, d_k) \\ \text{s.t. } x_i \in \{0, 1\} \quad i \in F. \end{aligned}$$

The second-stage problem for each scenario k , $Q(x, d_k)$ is defined using transportation variables y_{ij} from a facility i to a customer j and auxiliary variables α_j that denote the amount of unmet demand of customer j . We have:

$$\begin{aligned} Q(x, d_k) = \min_{y, \alpha} \quad & \sum_{i \in F} \sum_{j \in C} c_{ij} y_{ij} + \sum_{j \in F} \rho_j \alpha_j \\ \text{subject to} \quad & \sum_{j \in C} y_{ij} \leq u_i x_i \quad \forall i \in F, \\ & d_{kj} - \sum_{i \in F} y_{ij} \leq \alpha_j \quad \forall j \in C, \\ & y_{ij} \geq 0 \quad \forall i \in F, j \in C, \\ & \alpha_j \geq 0 \quad \forall j \in C. \end{aligned}$$

By allowing unmet demand, the problem always has a feasible solution and Benders decomposition only requires optimality cuts.

A.2 Multi Commodity Network Design Problem

Consider a directed network with node set N , arc set A , and commodity set \mathcal{K} . Each commodity ℓ has an uncertain demand \tilde{v}_ℓ that must be routed from an origin node, $o_\ell \in N$, to its destination node, $d_\ell \in N$. The vector of demands in scenario $k \in [K]$ is denoted by $v_k = [v_k^\ell : \ell \in \mathcal{K}]$. For each arc $(i, j) \in A$, there is an installation cost f_{ij} and an arc capacity u_{ij} . The cost for transporting one unit of commodity ℓ on installed arc (i, j) is c_{ij}^ℓ . Any demand that is not met is penalized at a rate of $B > 0$ per unit.

The objective in the first-stage is to decide which subset of arcs to install to minimize the sum of arc installation cost and expected total transportation cost and penalty for unmet demand. In the second-stage, after demand is realized, the goal is to determine the optimal flow of commodities through the installed arcs to minimize the sum of transportation and unmet demand costs.

In the first-stage, we define binary decisions x_{ij} for all arcs $(i, j) \in A$ such that $x_{ij} = 1$ if we install arc (i, j) . The first-stage formulation is given by:

$$\begin{aligned} \min_x \quad & \sum_{(i,j) \in A} f_{ij} x_{ij} + \sum_{k \in [K]} p_k Q(x, v_k) \\ \text{s.t. } & x_{ij} \in \{0, 1\} \quad (i, j) \in A. \end{aligned}$$

In the second-stage, we define non-negative continuous decisions y_{ij}^ℓ to represent transportation units of commodity ℓ on arc (i, j) . Additionally, we introduce auxiliary variables α_i^ℓ to denote the unmet demand for commodity ℓ . For scenario k , the formulation is given by:

$$\begin{aligned} Q(x, v_k) = \min_{y, \alpha} \quad & \sum_{(i,j) \in A} \left[\sum_{\ell \in \mathcal{K}} c_{ij}^\ell y_{ij}^\ell + B \alpha_i^\ell \right] \\ \text{subject to} \quad & \sum_{j: (i,j) \in A} y_{ij}^\ell - \sum_{j: (j,i) \in A} y_{ji}^\ell = g_i^\ell (v_k^\ell - \alpha_i^\ell) \quad \forall i \in N, \ell \in \mathcal{K}, \\ & \sum_{\ell \in \mathcal{K}} y_{ij}^\ell \leq u_{ij} x_{ij} \quad \forall (i, j) \in A, \\ & y_{ij}^\ell \geq 0 \quad \forall (i, j) \in A, \ell \in \mathcal{K}, \\ & \alpha_i^\ell \geq 0 \quad \forall i \in N, \ell \in \mathcal{K}. \end{aligned}$$

The parameter g_i^ℓ is set to 1 if node i is the origin of the commodity ℓ , -1 if node i is the destination of the commodity ℓ , or 0 otherwise. By allowing unmet demand, the problem always has a feasible solution and Benders decomposition only requires optimality cuts.

A.3 Stochastic Unit Commitment Problem

Consider a set \mathcal{G} of thermal generators operating over a discrete time horizon $\mathcal{T} = \{1, \dots, T\}$ with time periods indexed by t . Each generator $g \in \mathcal{G}$

has power output bounds $\underline{P}_g, \bar{P}_g$ (minimum and maximum power output in MW when the generator is on), temporal constraints U_g (minimum up-time) and D_g (minimum down-time) measured in number of periods, and marginal production cost C_g (\$/MWh). The electricity demand at time t in scenario $k \in [K]$ is denoted by B_{tk} (MW). To ensure feasibility of the subproblems, we allow unmet demand with penalty cost M per unit. The objective is to determine the commitment schedule (which generators are on or off at each time period) to minimize expected total cost.

The first-stage binary decision variables $x_{gt} \in \{0, 1\}$ indicate the commitment status of generator g at time t (1 if on, 0 if off), for $g \in \mathcal{G}, t \in \{0, 1, \dots, T\}$, where x_{g0} represents the initial state at time 0. Additionally, binary variables w_{gt} and v_{gt} for $g \in \mathcal{G}, t \in \mathcal{T}$ track startup and shutdown events. In the second-stage for scenario k , after revelation of demand, continuous decision variables $p_{gtk} \geq 0, \forall g \in \mathcal{G}, t \in \mathcal{T}$, represent power output from generator g at time t , and auxiliary variables $\sigma_{tk} \geq 0, \forall t \in \mathcal{T}$ to denote the unmet demand at time t . The first-stage formulation is given by:

$$\begin{aligned} \min_x \quad & \sum_{k \in [K]} p_k Q(x, B_k) \\ \text{s.t.} \quad & x_{g0} = 0 \quad \forall g \in \mathcal{G}, \\ & x_{gt} = x_{g,t-1} + w_{gt} - v_{gt} \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, \\ & \sum_{s=\max\{t-U_g+1,1\}}^t w_{gs} \leq x_{gt} \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, \\ & \sum_{s=\max\{t-D_g+1,1\}}^t v_{gs} \leq 1 - x_{gt} \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, \\ & x_{gt}, w_{gt}, v_{gt} \in \{0, 1\} \quad \forall g \in \mathcal{G}, t \in \mathcal{T}. \end{aligned}$$

The second-stage problem for each scenario k , $Q(x, B_k)$ is defined as:

$$\begin{aligned} Q(x, B_k) = \min_{p, \sigma} \quad & \sum_{g \in \mathcal{G}} \sum_{t \in \mathcal{T}} C_g p_{gtk} + \sum_{t \in \mathcal{T}} M \sigma_{tk} \\ \text{subject to} \quad & \underline{P}_g x_{gt} \leq p_{gtk} \leq \bar{P}_g x_{gt} \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, \\ & \sum_{g \in \mathcal{G}} p_{gtk} + \sigma_{tk} \geq B_{tk} \quad \forall t \in \mathcal{T}, \\ & p_{gtk} \geq 0 \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, \\ & \sigma_{tk} \geq 0 \quad \forall t \in \mathcal{T}. \end{aligned}$$

By allowing unmet demand, the problem has relatively complete recourse.