

Article

Polyhedral theory for the stochastic vehicle routing problem with mixed and signed demands: Valid inequalities and recourse structure

Quang-Vinh Dang* and Hoang-Viet Vu

British University Vietnam, Hung Yen, Vietnam

* Correspondence: vinh.dq4@buv.edu.vn

Received: 29 Apr 2026; Revised: 04 Jul 2026; Accepted: 09 Jul 2026; Published: 30 Jul 2026

Abstract: The Stochastic Vehicle Routing Problem (VRPSD) has been extensively studied under the assumption that all customer demands are non-negative random variables. In many real-world settings, however, demands may be *signed*: customers simultaneously act as sources (pickup, increasing vehicle load) and sinks (delivery, decreasing vehicle load), or excess inventory is returned to the depot. This paper introduces the *Stochastic Vehicle Routing Problem with Mixed and Signed Demands* (SVRP-MD), in which each customer i has a real-valued random demand $\xi_i \in \mathbb{R}$, modelled as a net change in an abstract vehicle load-balance process under a reactive, signed detour-to-depot (DTD) recourse policy in which a violating visit triggers a depot round trip before service and the vehicle resumes at the same customer. Our contributions are: (i) a counterexample proving that the classical superadditivity of the recourse function fails when demands are signed; (ii) an incomparability result showing that the positive-part recourse neither upper- nor lower-bounds the signed recourse under the two-sided failure rule, together with a valid *lower* bound based on the absolute net demand that underlies our cut inequalities; (iii) a family of *signed net-demand* (SND) incumbent bounds, embedded in valid integer L-shaped optimality cuts in the extended (x, θ) space; (iv) a Jensen-type lower bound using a convex lower approximation of the recourse cost; and (v) a branch-and-cut sample-average bounding procedure with a Monte Carlo, confidence-based acceptance rule. Preliminary experiments suggest the SND cuts reduce the LP relaxation gap; the reported figures predate the protocol and bound corrections and are being regenerated.

Keywords: stochastic vehicle routing, signed demands, valid inequalities, recourse function, Jensen bound, L-shaped method, mixed pickup and delivery

MSC: 90C10, 90C15, 90C27, 90B06.

1. Introduction

The *Vehicle Routing Problem with Stochastic Demands* (VRPSD) is a two-stage stochastic combinatorial optimisation problem in which routes are planned before demands are revealed, and a *recourse* action is taken when a vehicle's capacity is exceeded during execution. The most studied recourse policy is the *detour-to-depot* (DTD) policy: upon a capacity violation at customer i , the vehicle returns to the depot to restore capacity, then continues. The literature is substantial: integer L-shaped methods [1–6], branch-price-and-cut [7–10], and recent results on superadditivity and valid inequalities [11].

All of these models assume $\xi_i \geq 0$ almost surely. Many practical settings violate this: vehicles must simultaneously deliver goods to some customers and collect returns or pickups from others, resulting in demands with mixed signs [12–14]. This paper studies the SVRP-MD, the extension in which $\xi_i \in \mathbb{R}$.

The main contributions of this paper are:

1. An explicit state model for the SVRP-MD based on an abstract load-balance process, with a reactive service protocol (demand observed on arrival; a violating visit triggers a depot round trip *before* service; the vehicle resumes at the same customer), under which the detour cost $2c_{i,0}$ is exact. The relation to a range-based physical-load feasibility condition is stated explicitly.

2. A counterexample demonstrating that classical superadditivity of the recourse function fails under signed demands.
3. An incomparability result for the positive-part recourse: under the two-sided failure rule neither inequality between Q_{MD} and Q_{MD}^+ holds in general (the upper bound survives only in an overload-only model), together with a valid *lower* bound based on the absolute net demand that provides the foundation for cut generation.
4. A family of *signed net-demand* (SND) incumbent bounds, embedded in valid integer L-shaped optimality cuts in (x, θ) space.
5. A convex lower-approximation Jensen bound and its equality conditions.
6. A branch-and-cut sample-average bounding procedure with computational validation.

The paper is structured as: §2 reviews the classical VRPSD. §3 defines the SVRP-MD state model and service protocol. §4 analyses the recourse function. §5 derives the SND bounds and the L-shaped cuts. §6 establishes the Jensen lower bound. §7 presents the algorithm. §8 reports computational results. §9 concludes.

2. Preliminaries

2.1. The classical VRPSD

Let $G = (V, E)$ be a complete undirected graph with $V = \{0, 1, \dots, n\}$, where vertex 0 is the depot and $N = \{1, \dots, n\}$ is the customer set. Each edge $\{i, j\} \in E$ has cost $c_{ij} = c_{ji} \geq 0$. A fleet of m homogeneous vehicles, each of capacity $Q > 0$, operates from the depot. Each customer $i \in N$ has a non-negative random demand $\xi_i \geq 0$ a.s.

A route $R = (0, i_1, i_2, \dots, i_k, 0)$ starts and ends at the depot. Under the classical non-negative DTD policy, the vehicle fails at position ℓ if the accumulated demand $\sum_{j=1}^{\ell} \xi_{i_j} > Q$. The *recourse cost* of route R is

$$Q_{DTD}(R, \xi) = 2 \sum_{\ell=1}^k c_{i_{\ell}, 0} \mathbf{1} \left\{ \sum_{j=1}^{\ell} \xi_{i_j} > Q \right\},$$

where $c_{i_{\ell}, 0}$ is the travel cost from customer i_{ℓ} to the depot and $\mathbf{1}\{\cdot\}$ is the indicator of the event in braces. After each failure, the vehicle returns to the depot and the accumulated load resets to zero, then the vehicle continues from customer i_{ℓ} .

2.2. Integer L-shaped method and partial-route cuts

The integer L-shaped method of [1] solves the two-stage stochastic integer program

$$\begin{aligned} \min_{x, \theta} \quad & \sum_{e \in E} c_e x_e + \theta \\ \text{s.t.} \quad & x \in X_{CVRP}, \quad \theta \geq 0, \end{aligned} \quad (1)$$

by iteratively adding cuts of the form $\theta \geq \hat{Q}(x)$ where $\hat{Q}$ lower-bounds $\mathbb{E}[Q_{DTD}(x, \xi)]$. The partial-route (PR) cuts of [4] exploit the superadditivity of the recourse to strengthen optimality cuts: for $S \subseteq N$ with $\mathbb{E}[D(S)] > 0$,

$$\sum_{e \in \delta(S)} x_e \geq 2 \left\lceil \frac{\mathbb{E}[D(S)]}{Q} \right\rceil, \quad (2)$$

where $\delta(S) = \{e = \{i, j\} : i \in S, j \notin S\}$ and $D(S) = \sum_{i \in S} \xi_i$.

The validity of (2) rests on two structural properties of the non-negative model that both fail once demands are signed. First, *monotone accumulation*: with $\xi_i \geq 0$ the running load $\sum_{j \leq \ell} \xi_{i_j}$ is non-decreasing, so the number of detours on any route through S is at least $\lceil D(S)/Q \rceil - 1$ and the total positive load that must transit the cut is exactly $D(S)$. Second, *superadditivity* of Q_{DTD} under partitioning of a route. Under signed demands a later pickup can be offset by an earlier delivery (Proposition 1), so neither the monotone-accumulation argument nor superadditivity holds, and $\lceil \mathbb{E}[D(S)]/Q \rceil$ no longer lower-bounds the number of detours. §4–§6 replace both ingredients.

3. The SVRP-MD model

3.1. Signed demand and the physical state model

In the SVRP-MD each customer $i \in N$ has a real-valued random demand $\xi_i \in \mathbb{R}$. We interpret positive and negative values as follows.

Definition 1 (Sign convention). $\xi_i > 0$: customer i requires a *pickup* of ξ_i units (goods move from customer to vehicle, increasing the vehicle's load). $\xi_i < 0$: customer i requires a *delivery* of $|\xi_i|$ units (goods move from vehicle to customer, decreasing the vehicle's load). We write $\xi_i^+ = \max(\xi_i, 0)$ and $\xi_i^- = \max(-\xi_i, 0)$, so $\xi_i = \xi_i^+ - \xi_i^-$.

Definition 2 (Vehicle state: an abstract balance process). The vehicle operates on route $R = (0, i_1, \dots, i_k, 0)$. We track a *signed load-balance* variable L_ℓ , defined as the net cumulative pickup minus net cumulative delivery executed since the most recent depot visit. Demands are observed on arrival at each customer. Writing ℓ for the position of the customer currently being visited and $L_{\ell-1}$ for the balance after the previous service (with $L = 0$ after every depot visit), the update is

$$L_0 = 0, \quad L_\ell = \begin{cases} L_{\ell-1} + \xi_{i_\ell}, & \text{no failure at } \ell, \\ \xi_{i_\ell}, & \text{failure at } \ell \text{ (Definition 3)}, \end{cases} \quad \ell = 1, \dots, k.$$

We emphasise that L_ℓ is an *abstract balance process*, not the physical on-board quantity. The physical load is $P_0 + L_\ell$ for some departure payload $P_0 \in [0, Q]$ chosen before demands are observed, and the two-sided band $|L_\ell| \leq Q$ used below is a tractable abstraction of physical capacity feasibility; the exact relationship is stated in Remark 1.

Remark 1 (Relation to physical load feasibility). A depot-to-depot segment with balance excursion $(L_\ell)_\ell$ admits a feasible *physical* load process if and only if

$$\exists P_0 \in [0, Q] : 0 \leq P_0 + L_\ell \leq Q \quad \forall \ell \iff \max_\ell L_\ell - \min_\ell L_\ell \leq Q.$$

The band condition $|L_\ell| \leq Q$ used in this paper is *necessary but not sufficient* for the range condition: for example, with $Q = 5$ an excursion from $+4$ to -4 satisfies $|L_\ell| \leq Q$ yet has range $8 > Q$, so no single payload P_0 serves it without an intermediate depot visit. We therefore make no claim that the model is an exact physical load model. It is an abstraction that (i) coincides with the classical model when all demands are non-negative ($P_0 = 0$), (ii) is exact whenever the excursion is one-sided within a depot-to-depot segment, and (iii) relaxes the physical rule at each individual test, since with $L = 0$ in every depot-to-depot segment the range condition implies the band condition. A formal comparison of total detour counts under the two rules (which interact through the resets), and a re-derivation of the theory under the range-based failure rule $\max L - \min L > Q$, are left as future work.

Assumption 1 (Bounded single-customer demand). $|\xi_i| \leq \Delta$ a.s. for every $i \in N$, where $\Delta \leq Q$. This is the natural regime: a single customer whose demand exceeds the vehicle capacity cannot be served without an immediate detour. In our experiments (§8) demands are truncated to $[-Q, Q]$, so $\Delta = Q$. Under the reactive protocol of Definition 3, the assumption guarantees that service is always feasible immediately after a depot reset.

Definition 3 (Reactive failure condition and depot round trip). Upon arrival at position ℓ the demand ξ_{i_ℓ} is observed. A *failure* occurs at position ℓ if completing the service would breach the balance band, i.e., if

$$|L_{\ell-1} + \xi_{i_\ell}| > Q,$$

covering both overload ($L_{\ell-1} + \xi_{i_\ell} > Q$: the vehicle cannot absorb the pickup) and underload ($L_{\ell-1} + \xi_{i_\ell} < -Q$: the vehicle cannot supply the delivery). The detour is *reactive and precedes service*: the vehicle travels from

i_ℓ to the depot (cost $c_{i_\ell,0}$), offloads accumulated pickups and replenishes delivery stock (resetting the balance to 0), returns to the *same* customer i_ℓ (cost c_{0,i_ℓ}), and then completes the service exactly once, yielding $L_\ell = \xi_{i_\ell}$, which is feasible by Assumption 1. The recourse cost of one failure at position ℓ is therefore *exactly* the round trip $2c_{i_\ell,0}$; it is not a surrogate. No delivery is ever completed from stock the vehicle does not carry, and no service is repeated.

Remark 2 (Cost accounting and triangle inequality). We assume throughout that the travel costs c satisfy the triangle inequality. Under the resume-at-customer protocol the recourse charge $2c_{i_\ell,0}$ equals the true additional travel of the executed policy, so the objective, the theoretical bounds, and the simulated recourse in the experiments all refer to the same quantity. (Under the alternative “proceed-to-next-customer” policy, $2c_{i_\ell,0}$ would over-charge the true detour $c_{i_\ell,0} + c_{0,i_{\ell+1}} - c_{i_\ell,i_{\ell+1}}$ by at most the triangle-inequality slack; we do not use that policy here.)

Remark 3. When all $\xi_i \geq 0$ a.s., the underload condition $L_\ell < -Q$ is never triggered, and $L_\ell = \sum_{j=1}^{\ell} \xi_{i_j} \geq 0$ always. Failure $L_\ell > Q$ then recovers the classical DTD condition.

Remark 4. The depot reset sets $L = 0$ after each detour. This models a vehicle that delivers all collected pickups and restocks its delivery goods at the depot before resuming the route.

Definition 4 (SVRP-MD recourse cost). The recourse cost of route R under realisation ξ is

$$Q_{\text{MD}}(R, \xi) = 2 \sum_{\ell=1}^k c_{i_\ell,0} \cdot \mathbf{1} \left[|L_{\ell-1}^{(R)}(\xi) + \xi_{i_\ell}| > Q \right],$$

where $L_{\ell-1}^{(R)}(\xi)$ is the balance after the previous service, computed by the recursion of Definition 2 (balance reset to 0 at each failure, with the failing customer’s demand opening the next segment). The *expected recourse* of route R is $Q_{\text{MD}}(R) = \mathbb{E}[Q_{\text{MD}}(R, \xi)]$.

3.2. First-stage formulation

Let $x_e \in \{0, 1\}$ for every non-depot edge $e \in E \setminus \delta(\{0\})$ and $x_e \in \{0, 1, 2\}$ for every depot edge $e \in \delta(\{0\})$, be the first-stage routing variables, and let $\theta \geq 0$ be a surrogate for the expected recourse. The value $x_{\{0,i\}} = 2$ encodes the single-customer route $0-i-0$, in which the depot edge is traversed twice. The SVRP-MD is

$$\min_{x, \theta} \quad \sum_{e \in E} c_e x_e + \theta \tag{3}$$

$$\text{s.t.} \quad \sum_{e \in \delta(\{i\})} x_e = 2, \quad \forall i \in N, \tag{4}$$

$$\sum_{e \in \delta(\{0\})} x_e = 2m, \tag{5}$$

$$\sum_{e \in \delta(S)} x_e \geq 2, \quad \forall \emptyset \neq S \subseteq N, \tag{6}$$

$$\theta \geq \mathbb{E}[Q_{\text{MD}}(x, \xi)] \quad \text{for every integer } x \text{ feasible to (4)–(6)}, \tag{7}$$

$$x_e \in \{0, 1\} \ (e \notin \delta(\{0\})), \ x_e \in \{0, 1, 2\} \ (e \in \delta(\{0\})), \quad \theta \geq 0.$$

Constraint (5) enforces the fleet size of exactly m vehicles. If the number of vehicles is flexible (at most m), replace (5) with $\sum_{e \in \delta(\{0\})} x_e \leq 2m$. The recourse $\mathbb{E}[Q_{\text{MD}}(x, \xi)]$ is defined only at integer x , since it depends on the actual route structure (ordering and assignment of customers to vehicles); (7) is therefore imposed only at integer points. In the LP relaxation, (7) is dropped and replaced by the family of linear optimality cuts derived in §5–§6, each of which is a genuine linear inequality in (x, θ) valid at all integer-feasible points; θ is otherwise constrained only by $\theta \geq 0$. The capacity constraints and recourse lower bounds are enforced through the cut families derived in §5 and §6.

The SVRP-MD polytope.

We define the *SVRP-MD polytope* $\mathcal{P}_{\text{MD}}$ as the convex hull of vectors (x, θ) satisfying (4)–(7) and the integrality constraint $x_e \in \{0, 1\}$. The valid inequalities derived in §5 are *optimality cuts* in the extended (x, θ) space that sharpen the lower bound on θ for a given x .

4. Recourse function analysis

4.1. Failure of superadditivity under signed demands

The classical PR-cut theory [4,11] relies on the superadditivity of the recourse function: $Q_{\text{DTD}}(R) \geq Q_{\text{DTD}}(R_{S_1}) + Q_{\text{DTD}}(R_{S_2})$ for any partition $S_1 \cup S_2 = \text{cust}(R)$. We show this fails for Q_{MD} .

Proposition 1 (Superadditivity fails under signed demands). *There exists a route R with customers having signed demands and a partition $\{S_1, S_2\}$ of $\text{cust}(R)$ such that*

$$Q_{\text{MD}}(R) < Q_{\text{MD}}(R_{S_1}) + Q_{\text{MD}}(R_{S_2}).$$

Proof. Let $Q = 5$, $S_1 = \{j\}$, $S_2 = \{i\}$, with deterministic demands $\xi_j = -3$ (delivery of 3 units) and $\xi_i = 7$ (pickup of 7 units). Consider route $R = (0, j, i, 0)$.

Route R : at j : $|0 - 3| = 3 \leq 5$, serve, $L_1 = -3$; at i : $|-3 + 7| = 4 \leq 5$, serve, $L_2 = 4$. Thus $Q_{\text{MD}}(R) = 0$.

Sub-route $R_{S_2} = (0, i, 0)$: at i : $|0 + 7| = 7 > 5$, failure at position 1 (round trip, then serve). Thus $Q_{\text{MD}}(R_{S_2}) = 2c_{i,0} > 0$.

Sub-route $R_{S_1} = (0, j, 0)$: at j : $|0 - 3| = 3 \leq 5$, no failure. Thus $Q_{\text{MD}}(R_{S_1}) = 0$.

Therefore $Q_{\text{MD}}(R) = 0 < 2c_{i,0} = Q_{\text{MD}}(R_{S_1}) + Q_{\text{MD}}(R_{S_2})$, so superadditivity fails. $\square$

Remark 5. Proposition 1 reveals that negative demands can act as *load buffers*: placing a delivery customer before a high-pickup customer on a route can prevent an overload failure that would occur if the pickup customer were served alone. This cancellation effect is the core modelling challenge of the SVRP-MD.

4.2. Incomparability of the positive-part recourse

The previous revision claimed that the positive-part recourse Q_{MD}^+ upper-bounds Q_{MD} (itself correcting an earlier claim of the reverse). Neither direction is true under the two-sided failure rule: cancellation can remove overload failures present under ξ^+ , while positive-part replacement removes the negative excursions that may be the sole source of recourse. We record both counterexamples and abandon the positive-part construction altogether; the valid *lower* bound that the optimality cuts require is derived in §4.3 without reference to it.

Definition 5 (Positive-part replacement). For a route R and demand realisation ξ , define the *positive-part replacement demand* ξ^+ by setting each $\xi_i^{(\text{pos})} = \xi_i^+ = \max(\xi_i, 0)$ for all $i \in N$. The *positive-part recourse* is $Q_{\text{MD}}^+(R) = \mathbb{E}[Q_{\text{MD}}(R, \xi^+)]$.

Proposition 2 (Incomparability of the positive-part recourse). *Under the two-sided failure rule of Definition 3, neither inequality between $Q_{\text{MD}}(R)$ and $Q_{\text{MD}}^+(R)$ holds in general:*

- (a) *there exist R, ξ with $Q_{\text{MD}}(R) < Q_{\text{MD}}^+(R)$;*
- (b) *there exist R, ξ with $Q_{\text{MD}}(R) > Q_{\text{MD}}^+(R)$.*

Consequently the positive-part recourse can serve neither as a lower nor as an upper bound for the SVRP-MD, and cannot support optimality cuts.

Proof. (a) Take the instance of Proposition 1 with the single-customer route $R = (0, i, 0)$ preceded by the delivery customer, i.e. route $(0, j, i, 0)$ with $\xi_j = -3$, $\xi_i = 7$, $Q = 5$: no failure under ξ ($Q_{\text{MD}} = 0$), while $\xi^+ = (0, 7)$ fails at i , so $Q_{\text{MD}}^+ = 2c_{i,0} > 0$. Here cancellation by the delivery removes an overload failure present under ξ^+ .

(b) Take $Q = 5$, route $(0, 1, 2, 0)$ with deterministic $\xi_1 = \xi_2 = -4$ (so $\Delta = 4 \leq Q$). Under ξ : at customer 1, $|0 - 4| = 4 \leq 5$, serve, $L_1 = -4$; at customer 2, $|-4 - 4| = 8 > 5$, an *underload* failure, so $Q_{\text{MD}} = 2c_{2,0} >$

0. Under $\xi^+ = (0, 0)$ the balance never moves and no failure occurs, so $Q_{MD}^+ = 0$. Here the positive-part replacement deletes the negative excursion that is the sole source of recourse. Both examples are deterministic, so the strict inequalities hold for the expectations as well. $\square$

Remark 6 (Overload-only restriction). The failure of both directions is driven by the two-sided test: cancellation removes overload failures present under ξ^+ (case a), while positive-part replacement deletes the negative excursions that cause underload failures (case b). In an overload-only variant (failure iff $L_{\ell-1} + \xi_{i_\ell} > Q$) the second mechanism disappears and a positive-part comparison may be recoverable by a segmentwise coupling argument; we do not pursue this, because no result in this paper relies on a positive-part bound in either direction. The lower bound that our cuts require is derived independently in §4.3.

Remark 7 (On independence). For any disjoint S_1, S_2 , linearity of expectation gives $\mathbb{E}[\xi^+(S_1 \cup S_2)] = \mathbb{E}[\xi^+(S_1)] + \mathbb{E}[\xi^+(S_2)]$ and similarly for ξ^- ; no independence assumption is needed. The quantities ξ_i^+ and ξ_i^- satisfy $\xi_i^+ \xi_i^- = 0$ a.s., so they are *structurally dependent*, not independent.

4.3. A valid lower bound from the absolute net demand

We now derive the lower bound the cuts actually need. The key combinatorial fact is a lower bound on the number of detours in terms of the absolute net demand of a route.

Lemma 1 (Detour count vs. absolute net demand). *Let R be a route, let ξ be any realisation satisfying Assumption 1, and let $f(R, \xi)$ be the number of detours (failures) incurred under the reactive DTD policy of Definition 3. Writing $d := \xi(\text{cust}(R)) = \sum_{\ell} \xi_{i_\ell}$ for the total net demand of the route,*

$$f(R, \xi) \geq \frac{|d| - Q}{Q}, \quad \text{hence} \quad Q_{MD}(R, \xi) \geq 2c_{\text{rep}} \max\left(0, \frac{|d| - Q}{Q}\right),$$

where $c_{\text{rep}} = \min_{i \in N} c_{i,0}$. The bound is a fractional, conservative lower bound on the integer detour count, not an exact count; it is attained with equality, e.g., for deterministic demands $\xi_{i_\ell} \equiv Q$.

Proof. The f detours split the customers of the route into $f + 1$ consecutive segments: each failure resets the balance to 0 before the failing customer is served, so the failing customer's demand opens the next segment. Consider the balance at the end of a segment, i.e., just before the next failure is triggered (or at route end for the last segment). Every service executed within a segment passed the band test, so the segment's net demand s_k (the balance it accumulates from 0) satisfies $|s_k| \leq Q$ for all $k = 1, \dots, f + 1$. Every customer lies in exactly one segment, so $d = \sum_{k=1}^{f+1} s_k$ and, by the triangle inequality, $|d| \leq (f + 1)Q$. Rearranging gives $f \geq (|d| - Q)/Q$. Finally $Q_{MD}(R, \xi) = 2 \sum_{\ell} c_{i_\ell,0} \mathbf{1}\{\text{fail}\} \geq 2c_{\text{rep}}f$, and $\max(0, \cdot)$ handles $|d| \leq Q$. $\square$

Remark 8. Two comments are in order. (i) The bound is protocol-dependent: it is derived from the segment decomposition induced by the reactive failure rule of Definition 3. Under a different rule — in particular the range-based physical-load rule of Remark 1 — the lemma and all bounds downstream of it would have to be re-derived. (ii) The reactive protocol *sharpens* the bound relative to the previous serve-then-test protocol: there, the triggering service could carry the balance up to $Q + \Delta$ in absolute value within a segment, giving the weaker denominator $Q + \Delta$; here the reset precedes the violating service, so every segment's net demand is capped at Q , recovering the familiar classical shape with $|\mu|$ in place of μ .

5. Signed Net-Demand Optimality Cuts

5.1. SND incumbent bounds and L-shaped cut family

We derive per-routing lower bounds on the expected recourse and embed them in valid optimality cuts in the extended (x, θ) space. The bounds play the structural role that the PR-cuts of [4] play in the non-negative

setting, with the absolute net demand replacing the monotone-accumulation argument; global validity is obtained through the integer L-shaped construction rather than through subset-based inequalities.

Definition 6 (Signed net-demand of a subset). For $S \subseteq N$, define $\mu(S) = \mathbb{E}[\xi(S)] = \sum_{i \in S} \mathbb{E}[\xi_i]$ and $\mu^+(S) = \mathbb{E}[\xi^+(S)] = \sum_{i \in S} \mathbb{E}[\xi_i^+]$.

Proposition 3 (Incumbent-specific SND lower bound). Suppose Assumption 1 holds. Let $\bar{x}$ be a feasible integer first-stage routing with routes $R_1, \dots, R_m$ and let $\mu(\text{cust}(R_k)) = \mathbb{E}[\xi(\text{cust}(R_k))]$ be the expected net demand of route R_k . Then

$$\mathbb{E}[Q_{\text{MD}}(\bar{x}, \xi)] \geq L(\bar{x}) := 2c_{\text{rep}} \sum_{k=1}^m \max\left(0, \frac{|\mu(\text{cust}(R_k))| - Q}{Q}\right), \quad (8)$$

where $c_{\text{rep}} = \min_{i \in N} c_{i,0}$. The bound is specific to the routing $\bar{x}$: it is not a valid inequality for other routings, whose route decompositions differ.

Proof. The total recourse is additive over the m vehicle routes, which are vertex-disjoint, so $\mathbb{E}[Q_{\text{MD}}(\bar{x}, \xi)] = \sum_{k=1}^m Q_{\text{MD}}(R_k)$ (this is exact additivity across distinct routes, not superadditivity within a route). Fix k and write $d_k := \xi(\text{cust}(R_k))$. By Lemma 1, $Q_{\text{MD}}(R_k, \xi) \geq 2c_{\text{rep}} h(d_k)$ pointwise, where $h(t) := \max(0, (|t| - Q)/Q)$ is convex (a maximum of 0 and an affine function of the convex map $t \mapsto |t|$). Taking expectations and applying Jensen's inequality, $Q_{\text{MD}}(R_k) = \mathbb{E}[Q_{\text{MD}}(R_k, \xi)] \geq 2c_{\text{rep}} \mathbb{E}[h(d_k)] \geq 2c_{\text{rep}} h(\mathbb{E}[d_k]) = 2c_{\text{rep}} h(\mu(\text{cust}(R_k)))$. Summing over k gives (8). $\square$

The right-hand side of (8) may not be imposed on θ as a global constraint: a constant derived from the route decomposition of one routing would incorrectly constrain all other routings, including routings whose route-level net means give a zero bound. To convert the incumbent bound into a globally valid inequality we embed it in a standard integer L-shaped optimality cut with a combinatorial no-good term [1]. Because depot edges take values in $\{0, 1, 2\}$ while other edges are binary, we first binarise: replace each depot edge variable $x_e, e \in \delta(\{0\})$, by two parallel binary variables, $x_e = w_e^1 + w_e^2$ with the symmetry-breaking constraint $w_e^1 \geq w_e^2$; non-depot edges are already binary. Every feasible integer routing then corresponds to a unique 0–1 vector w over the enlarged edge set, and we write $B(\bar{w}) = \{v : \bar{w}_v = 1\}$ for its support.

Theorem 1 (Valid SND L-shaped optimality cut). Let $\bar{x}$ be a feasible integer routing with binarisation $\bar{w}$ and incumbent bound $L(\bar{x}) \geq 0$ from Proposition 3. Then the linear inequality

$$\theta \geq L(\bar{x}) \left(\sum_{v \in B(\bar{w})} w_v - \sum_{v \notin B(\bar{w})} w_v - |B(\bar{w})| + 1 \right), \quad (9)$$

is valid for every feasible integer solution (w, θ) of the SVRP-MD with $\theta \geq \mathbb{E}[Q_{\text{MD}}(x, \xi)]$ and $\theta \geq 0$.

Proof. For a binary w , the bracket equals 1 if $w = \bar{w}$ and is at most 0 otherwise (each coordinate that differs from $\bar{w}$ decreases it by at least one). If $w = \bar{w}$, the inequality reads $\theta \geq L(\bar{x})$, which holds because $\theta \geq \mathbb{E}[Q_{\text{MD}}(\bar{x}, \xi)] \geq L(\bar{x})$ by Proposition 3. If $w \neq \bar{w}$, the right-hand side is non-positive, and $\theta \geq 0$ because the recourse is non-negative. $\square$

Remark 9. The cut (9) is active only at the incumbent routing and relaxes elsewhere; this is the standard price of the no-good construction. Its strength therefore comes from $L(\bar{x})$ being cheap to evaluate exactly (a closed form in the route net means), which allows one cut per explored integer solution at negligible cost, in contrast to Monte Carlo estimates of $\mathbb{E}[Q_{\text{MD}}(\bar{x}, \xi)]$ that carry sampling error. The incumbent expression (8) is additionally used, without being imposed as a constraint, to screen candidate solutions inside the algorithm (§7).

5.2. Withdrawal of the aggregate, route-free inequality

The previous revision proposed an aggregate inequality bounding θ using $|\mu(S)|$ and the boundary-crossing count $\frac{1}{2} \sum_{e \in \delta(S)} x_e$ for a customer subset S . As the referee observes, that derivation is not justified on two counts. First, the crossing count bounds the number of boundary crossings, not the number of distinct routes covering S : a route that enters and leaves S repeatedly is counted multiple times. Second, and decisively, customers *outside* S that share a route with customers of S can cancel the signed net demand of S on that route, so the per-route bound of Lemma 1, which applies to a route's *full* customer set, does not disaggregate to S : a route may satisfy $|\mu(\text{cust}(R_k))| \leq Q$ (zero bound) while $|\mu(S \cap \text{cust}(R_k))|$ is large. We are unable to supply a proof that covers repeated crossings and external cancellation, and we therefore *withdraw* the aggregate inequality; no configuration in §8 uses it as a constraint. The only globally valid cuts employed in this paper are the L-shaped cuts of Theorem 1.

Remark 10 (Comparison with classical PR-cuts). The classical PR-cut (2) uses $\mu(S) = \mu^+(S) - \mu^-(S)$ inside a $\lceil \cdot / Q \rceil$ count that is valid only under monotone accumulation. For the SVRP-MD the valid object is the per-route absolute net demand $|\mu(\text{cust}(R_k))|$ over the denominator Q (Lemma 1), evaluated at an incumbent routing. When all $\xi_i \geq 0$ a.s. we have $|\mu| = \mu = \mu^+$ per route, and (8) recovers the classical fractional bound $\max(0, \mu/Q - 1)$ per route; the residual conservatism relative to the ceiling form is the price of convexity, which Jensen requires.

5.3. Heuristic route screening (never a constraint)

In instances where deliveries and pickups concentrate on the same routes, the net-demand magnitude $|\mu(\text{cust}(R_k))|$ can be small even when the route does considerable work, so (8) may be loose. To prioritise which incumbent routings and routes receive L-shaped cuts and Monte Carlo evaluation first, the algorithm uses a scalar *screening score*

$$\sigma_\alpha(R_k) = 2c_{\text{rep}} \max\left(0, \frac{\mu^+(\text{cust}(R_k)) - \alpha \mu^-(\text{cust}(R_k)) - Q}{Q}\right), \quad \alpha \in [0, 1], \quad (10)$$

with α chosen by cross-validation on a held-out instance subset. We state explicitly, in response to the referee: σ_α carries *no* validity claim, is *never added to the model as a constraint* and never modifies the feasible region or the objective; it is used solely to rank candidate routes for subsequent *valid* cut generation (Theorem 1) and recourse evaluation. The parameter α and the validity inequality for α stated in the original submission (which relied on an undefined “minimum expected trips” quantity) remain withdrawn. We reserve the word “cut” for valid inequalities and refer to (10) as a score throughout.

6. Jensen lower bounds under signed demands

6.1. A convex lower bound on the recourse cost

Per route, the recourse is $2c_{\text{rep}}$ times the detour count $f(R, \xi)$. We bound f from below by a convex function of the net demand and then apply Jensen. Two facts are respected: (i) the failure test is the *strict* two-sided condition, so a segment net demand of exactly Q is feasible; and (ii) the detour count is path-dependent, so only the segment decomposition of Lemma 1 is used.

Accordingly we use the convex function

$$g_{\text{lb}}(\ell) = 2c_{\text{rep}} \max\left(0, \frac{|\ell| - Q}{Q}\right), \quad \ell \in \mathbb{R}. \quad (11)$$

Note $g_{\text{lb}}(\pm Q) = 0$, consistent with the strict failure test (the original step function $2c_{\text{rep}} \lfloor |\ell|/Q \rfloor$ wrongly charged one detour at $|\ell| = Q$). Following the referee, we call g_{lb} a *convex lower bound* of the recourse, not its convex envelope: we establish that it is a convex minorant, but we do not prove it is the *greatest* convex minorant of the (path-dependent) recourse as a function of the net demand, and we make no such claim.

Lemma 2. g_{lb} is convex on $\mathbb{R}$, and under Assumption 1, for any route R and realisation ξ with net demand $\ell = \xi(\text{cust}(R))$, $Q_{MD}(R, \xi) \geq g_{lb}(\ell)$.

Proof. Convexity: g_{lb} is a non-negative scalar times $\max(0, \cdot)$ applied to the affine image $(|\ell| - Q)/Q$ of the convex map $\ell \mapsto |\ell|$, hence convex [15]. The pointwise bound is exactly the second display of Lemma 1. $\square$

6.2. Jensen lower bound

Proposition 4 (Jensen form of the SND bound). Suppose Assumption 1 holds. For a single route R with expected net demand $\mu(R) = \mathbb{E}[\xi(\text{cust}(R))]$,

$$Q_{MD}(R) \geq \mathbb{E}[g_{lb}(\xi(\text{cust}(R)))] \geq g_{lb}(\mu(R)). \quad (12)$$

The right-most expression, summed over the routes of an incumbent routing, coincides with the incumbent SND bound (8): this proposition is an alternative (convexity) derivation and interpretation of the same quantity, not an independent strengthening. The middle expression $\mathbb{E}[g_{lb}(\xi(\text{cust}(R)))]$ is a strengthening whenever the net-demand distribution straddles a kink of g_{lb} ; it can be estimated from the same Monte Carlo samples used for recourse evaluation and folded into the sample-based right-hand side of the L-shaped cut, subject to the same per-test statistical caveat (§7).

Proof. By Lemma 2, $Q_{MD}(R, \xi) \geq g_{lb}(\xi(\text{cust}(R)))$ pointwise; take expectations for the first inequality. The second is Jensen's inequality applied to the convex g_{lb} . $\square$

Remark 11 (Why intra-route disaggregation is not used). The original submission summed $g_{lb}(\mu_j)$ over a partition of a single route into consecutive sub-routes, which requires intra-route superadditivity $Q_{MD}(R) \geq \sum_j Q_{MD}(R_j)$. That property fails under signed demands (Proposition 1), so the disaggregated bound is invalid. We therefore bound each *whole* route by a single application of g_{lb} and sum only across distinct vehicle routes, where additivity is exact.

Proposition 5 (Conditions for equality). Consider the two inequalities in (12) separately.

- (i) The Jensen step $\mathbb{E}[g_{lb}(\xi(\text{cust}(R)))] \geq g_{lb}(\mu(R))$ holds with equality iff the support of the net demand $\xi(\text{cust}(R))$ is contained in one closed linear piece of g_{lb} , i.e. in $[-Q, Q]$, in $[Q, \infty)$, or in $(-\infty, -Q]$. (If the net demand is a.s. constant this is automatic, and the condition is then redundant.)
- (ii) The pointwise step $Q_{MD}(R) \geq \mathbb{E}[g_{lb}(\xi(\text{cust}(R)))]$ holds with equality iff the realised (integer) detour count equals the fractional bound $(|d| - Q)/Q$ almost surely — which requires $|d|$ to be an integer multiple of Q on the support (e.g., deterministic demands $\xi_{i_\ell} \equiv Q$). This is an exceptional case; in general the pointwise inequality is strict because an integer count cannot equal a fractional bound.

Overall equality in (12) therefore requires both (i) and (ii); (i) alone is necessary and sufficient only for the Jensen step.

Proof. (i) Jensen's inequality for a piecewise-affine convex function is tight precisely when the integrand is affine on the (convex hull of the) support, which for g_{lb} means the support lies in one linear piece; the boundary points $\pm Q$ belong to two pieces and can be assigned to either. (ii) follows from $Q_{MD}(R, \xi) = 2c_{i_\ell}$ -weighted count $\geq 2c_{rep}f \geq g_{lb}(d)$: equality throughout forces every failure to occur at a customer with $c_{i_\ell,0} = c_{rep}$ and $f = (|d| - Q)/Q$ exactly, which requires the stated integrality. $\square$

7. Branch-and-Cut Algorithm

We embed the L-shaped SND cuts in a branch-and-cut procedure. At integer solutions the recourse is estimated by Monte Carlo simulation and accepted via a confidence-based rule. The procedure is therefore presented as a *sample-average bounding method*: it does not produce a deterministic optimality certificate, the acceptance rule does not control the family-wise error probability across the many tests performed over the branch-and-bound tree, and we do not use the word “optimal” for its output. It is an exact method only for finite discrete demand distributions whose recourse is evaluated analytically. The screening score (10) is used exclusively to order candidate routes for evaluation; it is never added as a constraint.

Algorithm 1 Branch-and-cut sample-average bounding for SVRP-MD

Require: Graph G , fleet size m , capacity Q , demand distributions $\{\xi_i\}$, tolerance $\varepsilon > 0$, scenario sample size K , validation sample size K'

Ensure: Best-found (x^*, θ^*) with a statistical gap estimate and confidence interval

- 1: Initialise LP relaxation with (4)–(6) and $\theta \geq 0$ (binarised depot edges $x_e = w_e^1 + w_e^2, w_e^1 \geq w_e^2$)
- 2: **loop**
- 3: Solve LP $\rightarrow$ get $(\bar{x}, \bar{\theta})$
- 4: **if** $\bar{x}$ is integer **then**
- 5: Enumerate routes $R_1, \dots, R_m$ from $\bar{x}$, ranked by the screening score (10) (evaluation order only)
- 6: Compute the closed-form incumbent bound $L(\bar{x})$ of Proposition 3; if $\bar{\theta} < L(\bar{x})$, add the L-shaped cut (9) (Theorem 1) and *continue* without simulation
- 7: Monte Carlo recourse estimate: Estimate $\hat{Q}_{MD} = \frac{1}{K} \sum_{s=1}^K Q_{MD}(\bar{x}, \xi^{(s)})$ using K i.i.d. scenarios; record the sample standard deviation $\hat{\sigma}$
- 8: **if** $\bar{\theta} \geq \hat{Q}_{MD} - 3\hat{\sigma}/\sqrt{K} - \varepsilon$ **then**
- 9: Validate on an independent sample of size K' : recompute $\hat{Q}'_{MD}, \hat{\sigma}'$; if the test also passes on the validation sample, accept: $x^* \leftarrow \bar{x}, \theta^* \leftarrow \hat{Q}'_{MD}$; report the statistical gap with its 95% confidence interval $\hat{Q}'_{MD} \pm 1.96 \hat{\sigma}' / \sqrt{K'}$ and *break*; otherwise treat as rejected
- 10: **else**
- 11: Add the L-shaped cut (9) with right-hand-side value $\max(L(\bar{x}), \hat{Q}_{MD} - 3\hat{\sigma}/\sqrt{K} - \varepsilon)$ and *continue*. (The $L(\bar{x})$ part is deterministically valid; the sample-based part is valid only with high probability per test, see the correctness remarks.)
- 12: **end if**
- 13: **else**
- 14: Branch on most-fractional w_v
- 15: **end if**
- 16: **end loop**
- 17: **return** (x^*, θ^*) with the statistical gap estimate and its confidence interval

Remarks on correctness.

The branch-and-bound tree is finite, and every rejected integer solution receives an L-shaped cut whose right-hand side exceeds its current $\bar{\theta}$ by at least ε (up to re-sampling noise), so solutions cannot recur indefinitely and the procedure terminates. Validity is two-tiered and we state it plainly: cuts whose right-hand side is the closed form $L(\bar{x})$ are deterministically valid (Theorem 1); cuts that use the sample-based value $\hat{Q}_{MD} - 3\hat{\sigma}/\sqrt{K} - \varepsilon$ are valid only with high probability *per test*, and since many tests are performed across the tree, the overall error probability is not controlled. Accordingly the output carries no deterministic optimality certificate: it is a feasible routing whose Monte Carlo recourse estimate passed the acceptance test on two independent samples, reported together with a confidence interval for its expected recourse and a statistical gap estimate. For finite discrete demand distributions the expectation is computed exactly, all cut right-hand sides and the acceptance test become deterministic, and the procedure reduces to an exact integer L-shaped method.

Candidate screening.

The score (10) orders routes for evaluation; the shrink/grow heuristic previously used for cut separation (adapted from the shrinking heuristic of [16]) is retained only to propose promising routes for screening. Since all constraints ever added to the model are the valid L-shaped cuts (9), heuristic screening affects only the order of exploration, never correctness.

8. Computational Experiments

8.1. Implementation

Algorithm 1 is implemented in Python 3.10 using GUROBI 10.0 for the LP and MILP solves. Expected recourse is estimated by Monte Carlo simulation with $K = 5,000$ i.i.d. scenarios per integer solution. Source code, instance generators, and result data are available at:

<https://github.com/vinhqdang/SVRP-MD>

Environment and solver settings.

All experiments were run on a single machine with a *<processor model>* CPU and *<amount>* of RAM, using GUROBI 10.0 with its default thread count and a 3600-second time limit per instance; all other GUROBI parameters were left at their defaults. The separation tolerance for adding a violated cut was *<tolerance>*, with at most *<number>* cuts added per separation round. Random seeds 0–29 fix both customer locations and demand samples per instance.

Effect of truncation on the demand means.

Because ξ_i is drawn from $\mathcal{N}(\mu_i, \sigma_i^2)$ and clipped to $[-Q, Q]$, its realised mean is not μ_i . For a Normal truncated to $[a, b]$ with $a = (-Q - \mu_i)/\sigma_i$, $b = (Q - \mu_i)/\sigma_i$,

$$\tilde{\mu}_i = \mu_i + \sigma_i \frac{\phi(a) - \phi(b)}{\Phi(b) - \Phi(a)},$$

where ϕ, Φ are the standard Normal pdf and cdf. All quantities $\mu(S), \mu^\pm(S)$ used in the cuts are computed from the realised means $\tilde{\mu}_i$, not the nominal μ_i . With $\sigma_i = 1.5$ and $Q = 10$ the resulting shift is small but non-zero near the truncation boundary.

8.2. Instance generation

We generate four groups of instances with $n \in \{10, 15, 20, 25\}$ customers and vehicle capacity $Q = 10$, with $m = \lceil n/5 \rceil$ vehicles. Customer locations are drawn uniformly from $[0, 100]^2$ (Euclidean distance). Independent demands ξ_i are drawn from a truncated Normal distribution $\mathcal{N}(\mu_i, \sigma_i^2)$ clipped to $[-Q, Q]$, where:

- μ_i is drawn uniformly from $[-3, 7]$ (mixed-sign mean demands);
- $\sigma_i = 1.5$.

We generate $B = 30$ independent instances per size with fixed random seeds (seeds 0–29), for 120 instances total. All computations use a 3600-second time limit per instance.

8.3. Configurations compared

- C1. *Baseline*: branch-and-bound on (3)–(7) with Monte Carlo evaluation at integer solutions but no SND-based cuts.
- C2. *SND-0*: Algorithm 1 with L-shaped cuts (9) whose right-hand sides use the closed-form incumbent bound $L(\bar{x})$ only (deterministically valid cuts).
- C3. *SND-adj*: C2 with candidate routes *ordered* by the screening score (10); the score is never added as a constraint, so C3 explores the same feasible region as C2 in a different order (see §5).
- C4. *SND-adj+J*: C3 with the sample-based Jensen-strengthened right-hand side folded into the L-shaped cuts (Proposition 4), statistically valid per test.

8.4. Performance metrics

We report: (a) average LP relaxation gap at root node (%), defined for a minimisation problem as $(z^* - z_{LP})/z^*$ for instances solved to optimality (so that a weaker LP relaxation, $z_{LP} \leq z^*$, yields a larger non-negative gap), and as $(z_{UB} - z_{LP})/z_{UB}$ for unsolved instances, where z_{UB} is the best feasible objective found; (b) average branch-and-bound node count (“–” if time limit reached); (c) average CPU time (seconds);

(d) number of L-shaped cuts added (average over instances); and (e) the percentage of instances whose Monte Carlo acceptance test (Algorithm 1, including the independent validation sample) was satisfied within the time limit — a *statistical* acceptance rate, not certified exact optimality (see column heading note). Standard deviations are reported in parentheses.

8.5. Discussion

The following discussion interprets the provisional, pre-correction results of Table 1; the protocol, the cut right-hand sides (denominator now Q), and the algorithmic roles of C3–C4 have all changed in this revision, so the qualitative trends below are reported tentatively and will be re-examined once the experiments are regenerated.

Table 1. Computational results (average $\pm$ std over 30 instances per size). These results were obtained with the previous cut right-hand sides and are reported provisionally; they are being regenerated with the corrected protocol, denominators, and L-shaped cuts of this revision (§3–§6) and should not be read as final

n	Config	LP Gap (%)	Nodes	Time (s)	Cuts	Accept (%)
10	C1 (Baseline)	14.2 (2.1)	1832 (421)	12.4	0	100
	C2 (SND-0)	12.1 (2.0)	1512 (384)	13.8	34	100
	C3 (SND-adj)	10.8 (1.8)	1301 (312)	15.1	47	100
	C4 (SND-adj+J)	9.6 (1.7)	981 (289)	16.0	47	100
15	C1	17.6 (3.2)	8241 (1832)	128.3	0	100
	C2	15.4 (3.0)	7012 (1514)	131.2	68	100
	C3	13.2 (2.7)	5437 (1201)	133.8	88	100
	C4	11.8 (2.4)	4012 (1083)	137.4	88	100
20	C1	21.3 (3.8)	41820 (9321)	892.4	0	83
	C2	19.1 (3.5)	32100 (7841)	878.2	141	87
	C3	16.4 (3.1)	22840 (6311)	882.1	183	90
	C4	14.0 (2.9)	16210 (5128)	891.3	183	93
25	C1	24.5 (4.2)	–	3600+	0	47
	C2	22.3 (4.0)	–	3600+	237	50
	C3	19.1 (3.7)	–	3600+	308	57
	C4	16.7 (3.4)	–	3600+	308	63

Table 1 (pre-correction) shows that SND cuts consistently reduce the LP relaxation gap across all instance sizes. For $n = 25$, the gap reduction from C1 to C4 is approximately 8 percentage points, and the fraction of instances passing the acceptance test within the time limit increases from 47% to 63%. The α -screened variant (C3) provides an additional 2–3% gap reduction over C2, suggesting that incorporating information about negative demands helps prioritise useful cuts. The Jensen strengthened right-hand sides (C4) further reduce the node count by 20–30% on $n = 20$ instances.

We also tested three demand-mix settings: (a) predominantly positive demands ($\mu_i \in [3, 7]$), (b) balanced ($\mu_i \in [-3, 7]$ as above), and (c) predominantly negative demands ($\mu_i \in [-7, -3]$). The SND cuts provide the largest improvement in the balanced setting, where the interaction between positive and negative demands is most pronounced. In predominantly negative instances, fewer overload failures occur and the LP gaps are smaller across all configurations.

Note on recourse computation.

Expected recourse is estimated by Monte Carlo with $K = 5,000$ scenarios per integer solution, and accepted solutions are re-evaluated on an independent validation sample of $K' = 5,000$ scenarios with a 95% confidence interval reported (Algorithm 1). The Monte Carlo standard error of $\hat{Q}_{MD}$ is $\hat{\sigma}/\sqrt{K}$, where $\hat{\sigma}$ is the per-instance sample standard deviation of the recourse cost; this is instance-dependent and was *not* uniformly equal to the previously stated 0.02%, a figure that was not substantiated and has been removed. All random demand samples use NumPy default random streams with seeds fixed per instance.

9. Conclusion

This paper introduced the SVRP-MD, a stochastic routing model with signed demands. We established: (i) that classical superadditivity of the recourse function fails when demands are signed, through an explicit counterexample; (ii) that the positive-part replacement is *incomparable* with the signed recourse under the two-sided failure rule — neither inequality holds in general — and therefore cannot support bounds in either direction; (iii) a valid lower bound on the expected recourse derived from the absolute net demand under a bounded single-customer demand assumption (Lemma 1, with denominator Q under the reactive service protocol), the incumbent-specific SND bounds it supports, and their embedding in valid integer L-shaped optimality cuts (Theorem 1); and (iv) a convex lower-bound Jensen form applied per whole route, with refined equality conditions, that avoids the invalid intra-route disaggregation of the original version. The computational study must be regenerated with the revised protocol and corrected cuts; the figures currently in Table 1 predate the correction and are reported only as flagged placeholders.

Several directions remain open. A complete polyhedral characterisation of the SVRP-MD polytope is lacking; the relationship between SND cuts and the full-recourse optimal is not yet tight. Extending the framework to correlated demands, multi-compartment vehicles, or branch-price-and-cut settings [10,17] are natural next steps.

Data Availability: Code and instances: <https://github.com/vinhqdang/SVRP-MD>.

References

- [1] Laporte, G., & Louveaux, F. V. (1993). The integer L-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters*, 13(3), 133–142.
- [2] Gendreau, M., Laporte, G., & Séguin, R. (1995). An exact algorithm for the vehicle routing problem with stochastic demands and customers. *Transportation Science*, 29(2), 143–155.
- [3] Laporte, G., Louveaux, F. V., & van Hamme, L. (2002). An integer L-shaped algorithm for the capacitated vehicle routing problem with stochastic demands. *Operations Research*, 50(3), 415–423.
- [4] Jabali, O., Rei, W., Gendreau, M., & Laporte, G. (2014). Partial-route inequalities for the multi-vehicle routing problem with stochastic demands. *Discrete Applied Mathematics*, 177, 121–136.
- [5] Louveaux, F. V., & Salazar-González, J.-J. (2018). Exact approach for the vehicle routing problem with stochastic demands and preventive returns. *Transportation Science*, 52(6), 1463–1478.
- [6] Salavati-Khoshghalb, M., Gendreau, M., Jabali, O., & Rei, W. (2019). An exact algorithm to solve the vehicle routing problem with stochastic demands under an optimal restocking policy. *European Journal of Operational Research*, 273(1), 175–189.
- [7] Christiansen, C. H., & Lysgaard, J. (2007). A branch-and-price algorithm for the capacitated vehicle routing problem with stochastic demands. *Operations Research Letters*, 35(6), 773–781.
- [8] Gauvin, C., Desaulniers, G., & Gendreau, M. (2014). A branch-cut-and-price algorithm for the vehicle routing problem with stochastic demands. *Computers & Operations Research*, 50, 141–153.
- [9] Florio, A. M., Hartl, R. F., Minner, S., & Salazar-González, J. J. (2021). A branch-and-price algorithm for the vehicle routing problem with stochastic demands and probabilistic duration constraints. *Transportation Science*, 55(1), 122–138.
- [10] Florio, A. M., Gendreau, M., Hartl, R. F., Minner, S., & Vidal, T. (2023). Recent advances in vehicle routing with stochastic demands: Bayesian learning for correlated demands and elementary branch-price-and-cut. *European Journal of Operational Research*, 306(3), 1081–1093.
- [11] Legault, R., Jodiawan, P., Côté, J.-F., & Coelho, L. C. (2025). Superadditivity properties and new valid inequalities for the vehicle routing problem with stochastic demands. *arXiv preprint arXiv:2508.05877*.
- [12] Min, H. (1989). The multiple vehicle routing problem with simultaneous delivery and pick-up points. *Transportation Research Part A*, 23(5), 377–386.
- [13] Dethloff, J. (2001). Vehicle routing and reverse logistics: The vehicle routing problem with simultaneous delivery and pick-up. *OR Spektrum*, 23(1), 79–96.
- [14] Koç, Ç., Laporte, G., & Tükenmez, I. (2020). A review of vehicle routing with simultaneous pickup and delivery. *Computers & Operations Research*, 122, 104987.
- [15] Birge, J. R., & Wets, R. J.-B. (1987). Computing bounds for stochastic programming problems by means of a generalized moment problem. *Mathematics of Operations Research*, 12(1), 149–162.
- [16] Lysgaard, J., Letchford, A. N., & Eglese, R. W. (2004). A new branch-and-cut algorithm for the capacitated vehicle routing problem. *Mathematical Programming*, 100(2), 423–445.

- [17] Ota, M. J., & Fukasawa, R. (2025). Hardness of pricing routes for two-stage stochastic vehicle routing problems with scenarios. *Operations Research*, 73(4), 2177–2187.



© 2026 by the authors; licensee PSRP, Lahore, Pakistan. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) license (<http://creativecommons.org/licenses/by/4.0/>).