

Highlights

Fair Vehicle Routing via Bilevel Optimization

Ivana Ljubić, Justo Puerto, Alberto Torreyon

- We introduce a family of fair capacitated routing problems with TSP-optimal routes
- An ordered framework handles a wide range of fairness measures in a single model
- A bilevel model guarantees that every route is an optimal TSP tour
- We quantify the price of ignoring TSP-optimality and who pays for fairness
- We provide extensive computational and managerial insights into fair routing

Fair Vehicle Routing via Bilevel Optimization

Ivana Ljubić^a, Justo Puerto^b, Alberto Torrejon^{a,*}

^a*Department of Finance, Accounting and Statistics, Institute for Statistics and Mathematics, WU Vienna, Welthandelsplatz 1, Vienna, 1020, Austria*

^b*Department of Statistics and Operational Research, Institute of Mathematics of the University of Seville (IMUS), University of Seville, Edificio Celestino Mutis, Avenida Reina Mercedes s/n, Seville, 41012, Spain*

Abstract

We propose a novel approach to fairness in the Capacitated Vehicle Routing Problem (CVRP) based on objective functions that order route lengths. We consider fairness criteria such as min-max, range, Gini mean difference, variance, and absolute deviations, including both monotonic and non-monotonic measures. To avoid inefficient routes, we impose per-vehicle Traveling Salesman Problem (TSP) optimality, requiring each vehicle to use a shortest route for its assigned customers. TSP optimality links customer assignment and routing, which we model through a bilevel optimization framework. We develop mixed-integer programming formulations using value-function reformulations and valid coupling inequalities to solve the problem certifying optimality. To the best of our knowledge, this is the first exact approach that addresses route inconsistency, a major issue in fair vehicle routing, for a broad family of fairness measures. Computational results provide formulations that solve state-of-the-art instances for several fairness measures. For the range, the only non-monotonic measure for which a tailored exact method exists, our general approach is up to two orders of magnitude faster than the state-of-the-art branch-price-and-cut algorithm. Finally, we study the trade-off between fairness and cost, showing how the budget affects fairness and how the resulting workload is distributed among drivers.

Keywords: Vehicle routing, Fairness, Bilevel optimization, Ordered optimization, TSP-optimality

*Corresponding author.

Email addresses: `ivana.ljubic@wu.ac.at` (Ivana Ljubić), `puerto@us.es` (Justo Puerto), `alberto.torrejon.valenzuela@wu.ac.at` (Alberto Torrejon)

1. Introduction

Traditional optimization models emphasize efficiency, typically understood as the minimization of a global cost function. However, as noted by Savas (1978), this focus can be insufficient in settings such as public administration decision-making, where all segments of the population must receive uniform and non-discriminatory treatment. A second criterion must therefore enter the decision process, namely *fairness*. Rooted in Equity Theory (Adams, 1963), fairness concerns the perceived equity of outcomes, since a solution that performs well at the system level but imposes disproportionate burdens on specific agents may be socially rejected or prove unsustainable. Fairness has no universally accepted definition, and the metrics proposed to capture its different aspects, such as proportionality, envy-freeness or balance, are often mutually incompatible, so the choice of metric shapes the solution. Even the oldest and most widely accepted approach, which focuses on the worst-off entity (Rawls, 1971), has been criticized for producing globally inefficient solutions and large aggregate losses for everyone else. This tension is intrinsic and known as the price of fairness (Mandell, 1991; Bertsimas et al., 2011). A more balanced distribution generally requires some agents to accept outcomes worse than a purely efficiency-oriented plan would give them. A significant body of literature therefore advocates integrating fairness as a primary objective, so that decision-makers can navigate the trade-off between system efficiency and equitable distribution (see Hooker and Williams, 2012; Bertsimas et al., 2012, among others). One of the most extensively studied families of decision-making problems serves as an ideal testbed for this dichotomy: *vehicle routing problems*.

The Capacitated Vehicle Routing Problem (CVRP), introduced by Dantzig and Ramser (1959), has long been a central topic in Operations Research, and its significance keeps growing in the increasingly logistics-oriented landscape of contemporary supply chains and courier companies (see, e.g., Archetti et al., 2025). It generalizes the Traveling Salesman Problem (TSP), the search for a shortest tour that visits each customer exactly once and returns to the depot, to a fleet of capacity-constrained vehicles, and it is therefore NP-hard. Current advances on the CVRP have been largely driven by branch-and-price algorithms grounded in the set-partitioning formulation of Balinski and Quandt (1964), among which *VRPSolver* stands out (Pessoa et al.,

2020; Errami et al., 2024). For comprehensive surveys of exact and heuristic algorithms, see Laporte (2009), Toth and Vigo (2014) or Archetti et al. (2025).

While effective at optimizing global objectives such as total travel distance or demand served, CVRP solutions can exhibit severe imbalances in driver workload, route duration, distance traveled or service requirements. Such disparities decrease driver satisfaction, increase employee turnover, degrade service quality in the long run, and may foster perceptions of unfairness, or even envy, among drivers. Embedding fairness within routing models is therefore a strategic need for sustainable logistics operations rather than a merely ethical consideration.

Fairness is usually promoted through objective functions that are less sensitive to extreme values, for instance by minimizing differences between route lengths or by explicitly targeting envy-freeness, which requires that no driver perceives the allocation of another driver as more favorable than their own. The latter is known to be partially captured by the Gini mean difference, the average absolute difference between two workloads, which is the numerator of the Gini coefficient, see Blanco et al. (2024) for a successful application in facility location. Other common measures include the min-max, the range, the variance and the absolute deviation from the mean. The interested reader is referred to Marsh and Schilling (1994), Karsu and Morton (2015), Lozano et al. (2016) and Matl et al. (2018) for extensive catalogues of fairness measures and detailed analyses of their properties.

A comprehensive and systematic study of fairness in routing, in particular concerning *workload equity*, was conducted by Matl et al. (2018). They deconstruct the notion into a fairness metric, which is the aspect of workload to be equalized (e.g., route length, service time or demand served), and a fairness measure, which is the function quantifying inequality across agents (e.g., range, variance or maximum). Their combination gives the fairness objective, the function actually optimized. A critical property of a fairness objective is its monotonicity. The objective is monotonic if a unilateral increase in the workload of a single driver cannot improve the overall equity value, which guarantees that worsening the condition of an individual never appears beneficial in terms of equity. Whereas the average (which is egalitarian rather than equitable) and the min-max or k -max measures are monotonic, many widely adopted and intuitive measures, such as the range, the variance and the Gini mean difference, are not.

Non-monotonicity induces a counterintuitive modeling dilemma, which we call route inconsistency. A model driven by a non-monotonic measure may

74 favor a solution in which the route of a driver is made needlessly inefficient,
 75 that is, it is not a TSP-optimal tour for the set of assigned customers. This
 76 may occur simply because the resulting increase in workload brings that route
 77 closer to the others and spuriously reduces the overall measure of dispersion.
 78 Such tours directly contradict the minimization of individual workloads.
 79 More generally, a non-monotonic objective can improve even in the extreme
 80 case in which all workloads get worse, a paradoxical manifestation of the
 81 price of fairness that Matl et al. (2018) study under the name of workload
 82 inconsistency. None of this happens under a monotonic objective, for which
 83 there always exists an optimal solution consisting only of TSP-optimal tours.
 84 Furthermore, in their theoretical analysis, they note that “it remains unclear
 85 to what extent workload inconsistency and non-TSP-optimality are a relevant
 86 concern in practice for nonmonotonic measures”, and their enumeration of
 87 Pareto sets on instances with up to 14 customers then reveals a large share
 88 of non-TSP-optimal and workload inconsistent solutions. We take up the
 89 question of non-TSP-optimality on larger instances, with an exact method
 90 and under a cost budget.

91 Consequently, much of the fairness-oriented literature in vehicle routing
 92 focuses on monotonic criteria, predominantly the min-max objective, which
 93 has been addressed with exact methods by Campbell et al. (2008), Huang
 94 et al. (2011) and Bertazzi et al. (2015), among others. The trade-off between
 95 efficiency and equity has also been modeled explicitly as a bi-objective problem,
 96 notably in the vehicle routing problem with route balancing, first studied
 97 by Jozefowiez et al. (2002) and later solved with exact methods by Reiter
 98 and Gutjahr (2012), Halvorsen-Weare and Savelsbergh (2016) and Artigues
 99 et al. (2018). TSP-optimal routing has also been studied in the multiple
 100 Traveling Salesman Problem (mTSP) of Svestka and Huckfeldt (1973), where
 101 fairness has again been approached mostly through monotonic criteria such as
 102 min-max (Applegate et al., 2002) or ℓ_p norms (Bektaş and Letchford, 2020).

103 Work on non-monotonic objectives is comparatively scarce. It includes the
 104 enumeration-based analysis of Matl et al. (2018) and post-hoc evaluations in
 105 heuristic (Lozano et al., 2016) and exact (López Sánchez et al., 2024) methods.
 106 Recently, van Rossum et al. (2026a) propose a branch-and-price method for
 107 the range and extend it to other order-based objectives such as the Gini mean
 108 difference, handling TSP-optimality either heuristically, by replacing each
 109 route of the final solution with its TSP-optimal counterpart, or exactly within
 110 the pricing problem, which becomes intractable for larger instances. In a
 111 follow-up work, van Rossum et al. (2026b) enforce TSP-optimality exactly for

the range, with lifted cutting planes that forbid arc sequences dominated by a shorter tour within a branch-price-and-cut algorithm.

Given a fairness measure X , we call X -CVRP-TSP the problem of assigning customers to capacitated vehicles, within a budget on the total distance, so that every route is a TSP-optimal tour for its customers and the measure X of the route lengths is minimized. The Range-CVRP-TSP has thus already been studied by van Rossum et al. (2026b), but models for more general measures, such as the Gini mean difference, absolute deviations or the variance, are presented here for the first time. To the best of our knowledge, ours is the first work that provides a unified framework for a broad range of fairness measures in the CVRP, monotonic, non-monotonic and even quadratic, while enforcing TSP-optimal routing with provable optimality guarantees.

Contributions. The main contributions of this paper are the following.

A unified ordered framework for fairness in routing. We propose a general modeling framework based on the principles of ordered optimization. By sorting route lengths, one directly identifies the most and the least burdened drivers, which enables fairness measures that explicitly account for inequality. Within Operations Research, ordered measures have been studied most extensively in location theory (Nickel and Puerto, 2006), where a large body of work addresses both their solution methods (Boland et al., 2006; Labbé et al., 2017; Puerto et al., 2017; Deleplanque et al., 2020; Marín et al., 2020; Pozo et al., 2024; Ljubić et al., 2024) and their fairness-oriented variants (Mesa et al., 2003; Blanco and Gázquez, 2023; Blanco et al., 2024; Tsang and Shehadeh, 2025). The routing literature is by contrast scarce. Ordered measures have recently been applied to the TSP by Ljubić et al. (2026), and to the CVRP by van Rossum et al. (2026a). Following Blanco et al. (2026), our framework accommodates both linear and non-linear measures, such as the variance, as well as criteria based on absolute or squared deviations from reference values such as the minimum or the median workload. It also allows negative entries in the weight vector, which most non-monotonic measures require. Thus, our framework provides decision makers with a diverse pool of fair CVRP solutions by simply varying the coefficients of the objective function, rather than requiring a tailored method for each fairness measure.

A bilevel model enforcing TSP-optimality. Our main methodological contribution addresses the route inconsistency inherent to non-monotonic objectives. We model the X -CVRP-TSP as a bilevel program (see Dempe, 2002; Dempe and Zemkoho, 2013; Kleinert et al., 2021, for further details on bilevel optimization). Its upper level assigns customers to drivers under

150 capacity and budget constraints while optimizing the fairness objective, and
 151 its lower level solves a collection of independent TSP problems that determine
 152 the exact length of each route. This yields exact approaches that enforce
 153 TSP-optimal routing for every driver.

154 *Dense and sparse single-level reformulations.* We explore two comple-
 155 mentary ways to handle the bilevel enforcement of TSP-optimality. Dense
 156 formulations represent the lower-level routing problems explicitly, through
 157 routing variables and constraints that encode the TSP of each driver, which
 158 gives accurate models at the expense of size. Sparse formulations instead
 159 project the routing variables out and describe route lengths implicitly, through
 160 value-function representations and dynamically generated cuts. For the sparse
 161 model, we further derive coupling inequalities based on the submodularity of
 162 the TSP cost, which make sparse formulations competitive in our experiments.

163 *Route inconsistency and the price of fairness in practice.* We quantify, on
 164 the benchmark, the route inconsistency that Matl et al. (2018) identified for
 165 non-monotonic measures. It is the rule. On every instance and for every non-
 166 monotonic measure, a model that does not enforce TSP-optimality returns
 167 a plan with artificially lengthened routes, overstates the achievable fairness
 168 by 12% on average and by up to 99%, and repairing its routes afterwards
 169 leaves the plan suboptimal in most cases. We also show who pays for fairness.
 170 With a budget 10% above the minimum cost, two or three of five drivers are
 171 typically worse off than in the cost-efficient plan, and with a sufficiently large
 172 budget every driver can be. Some optimal fair plans are even dominated
 173 route by route by the cost-efficient plan, the workload inconsistency of Matl
 174 et al. (2018), which TSP-optimality does not prevent. We show that tight
 175 budgets limit this effect.

176 *An extensive computational study.* Our experiments show that our model-
 177 ing choices are decisive for tractability. The sparse formulation is more than
 178 an order of magnitude faster than the dense one, and the coupling inequalities
 179 reduce the search trees by two orders of magnitude and the solution times
 180 by up to three. On the Range-CVRP-TSP, the only non-monotonic measure
 181 for which a tailored exact method exists, our general method is up to two
 182 orders of magnitude faster than the state of the art, the branch-price-and-cut
 183 algorithm of van Rossum et al. (2026b) designed for this single measure, and
 184 solves instances with 35 customers. For seven different fairness measures it
 185 proves optimality on all instances with up to 30 customers when the budget
 186 is 10% above the minimum cost. These results help decision-makers assess
 187 how much fairness can be built into their planning models while keeping them

188 tractable.

189 *Structure of the paper.* Section 2 introduces the modeling framework based
 190 on the ordering of route lengths and shows how it accommodates a broad
 191 spectrum of linear and non-linear fairness measures. Section 3 discusses the
 192 role of non-monotonicity and TSP-optimality and establishes the need for
 193 a bilevel treatment. Section 4 presents the bilevel model, which enforces
 194 TSP-optimality for the route of each driver. Section 5 reformulates it as a
 195 single-level problem through value-function arguments, in a dense version
 196 that keeps the routing variables and a sparse one that projects them out.
 197 It also presents the coupling inequalities and replaces the remaining non-
 198 convex constraints by two families of valid inequalities. Section 6 reports
 199 the computational study, and Section 7 discusses the managerial insights,
 200 namely the price of ignoring TSP-optimality and the redistribution of workload
 201 among drivers. Section 8 presents concluding remarks and directions for future
 202 research. The proofs are given in Appendix A and the implementation details
 203 in Appendix B.

204 2. Ordered framework for fair vehicle routing

205 The CVRP generalizes the TSP. It looks for a set of minimum-cost routes for
 206 a fleet of identical vehicles that serve the demand of a set of customers without
 207 exceeding the vehicle capacity. It is defined on a complete directed graph
 208 $G = (V, A)$ with nodes $V = \{0, 1, \dots, n\}$ and arcs $A = \{(i, j) \in V \times V : i \neq j\}$.
 209 Node 0 is the depot of a fleet K of $m = |K|$ vehicles with capacity $Q > 0$, and
 210 $V' = V \setminus \{0\}$ is the set of customers. Customer $i \in V'$ has demand $q_i \geq 0$,
 211 and $q_0 = 0$. Each arc $(i, j) \in A$ has a positive cost c_{ij} . We assume that costs
 212 are symmetric, $c_{ij} = c_{ji}$, and satisfy the triangle inequality $c_{ij} + c_{jl} \geq c_{il}$ for
 213 all $i, j, l \in V$, and we set $c_{ii} = 0$. A solution consists of m routes such that
 214 each route starts and ends at the depot, each customer is visited exactly once,
 215 and the total demand of each route does not exceed Q .

216 We use the CVRP as a baseline and build on it a framework that covers a
 217 wide range of fairness measures. Let $\Phi_k \geq 0$ be the length of the k -th shortest
 218 route of a solution, $k \in K$. Sorting the route lengths in non-decreasing order
 219 and weighting them with a vector $\lambda = (\lambda_k)_{k=1}^m$, $\lambda_k \in \mathbb{R}$, gives a general
 220 objective, and different vectors λ give different measures. Table 1 lists the
 221 vectors of some well-known statistics. Besides the mean, the maximum and
 222 the quantiles, it includes measures such as the range, the Gini mean difference
 223 or deviations from a quantile, which are linear in the sorted variables Φ_k .

Table 1: Ordered measures and their λ -vectors for a fixed sample $a_1, \dots, a_n$ with order statistics $a_{(1)} \leq \dots \leq a_{(n)}$.

Statistic	Expression	λ
Mean	$\frac{1}{n} \sum_{i=1}^n a_i$	$\frac{1}{n}(1, \dots, 1)$
Maximum	$\max a_i$	$(0, \dots, 0, 1)$
Minimum	$\min a_i$	$(1, 0, \dots, 0, 0)$
Median	$\text{median } a_i = \begin{cases} a_{(h+1)}, & n = 2h+1 \\ \frac{a_{(h)} + a_{(h+1)}}{2}, & n = 2h \end{cases}$	$(0, \dots, 0, \underbrace{1}_{(h+1)-th}, 0, \dots, 0)$ $(0, \dots, 0, \underbrace{\frac{1}{2}}_{h-th}, \underbrace{\frac{1}{2}}_{(h+1)-th}, 0, \dots, 0)$
q -th quantile (or τ -quantile)	$a_{(q)}, q = \lfloor \tau n \rfloor$	$(0, \dots, 0, \underbrace{1}_{q-th}, 0, \dots, 0)$
Range	$\max a_i - \min a_i$	$(-1, 0, \dots, 0, 1)$
General interquantile range (e.g., interquartile or interdecile ranges)	$a_{(s)} - a_{(r)}, (r < s)$	$(0, \dots, 0, \underbrace{-1}_{r-th}, 0, \dots, 0, \underbrace{1}_{s-th}, \dots, 0)$
Gini mean difference (mean absolute difference)	$\frac{1}{n^2} \sum_{i,j=1}^n a_i - a_j $	$\frac{2}{n^2} (2i - n - 1)_{i=1}^n$
Mean absolute deviation from the q -th quantile ($\text{MAD}_{(q)}$)	$\frac{1}{n} \sum_{i=1}^n a_i - a_{(q)} $	$\frac{1}{n} (-1, \dots, -1, \underbrace{2q - n - 1}_{q-th}, 1, \dots, 1)$

224

At a high level, the problem reads

$$l_{\text{ord}}(\Phi, \lambda) = \min \sum_{k \in K} \lambda_k \Phi_k, \quad (1a)$$

$$\text{s.t. } \Phi_k \leq \Phi_{k+1}, \quad \forall k \in K : k < m, \quad (1b)$$

$$\Phi_k \text{ is the length of route } k, \quad \forall k \in K. \quad (1c)$$

225

The objective (1a) is a weighted sum of the route lengths, which constraints (1b) keep sorted, and is usually called an ordered measure. Constraints (1c) state that Φ_k is the length of the corresponding route. We omit the routing constraints, which depend on the problem at hand, so the framework can be used in many other vehicle routing problems.

230

The framework also covers non-linear measures that penalize the deviation from a reference value, such as the mean or a quantile of the route lengths. It suffices to replace (1a) by a loss $\mathcal{L}(\Phi, l_{\text{ord}}(\Phi, \lambda))$ built on that reference, with $l_{\text{ord}}(\Phi, \lambda) = \sum_{k \in K} \lambda_k \Phi_k$. We use the absolute and the squared deviation,

233

$$\mathcal{L}(\Phi, l_{\text{ord}}(\Phi, \lambda)) = \sum_{k \in K} |\Phi_k - l_{\text{ord}}(\Phi, \lambda)|, \quad (2)$$

$$\mathcal{L}(\Phi, l_{\text{ord}}(\Phi, \boldsymbol{\lambda})) = \sum_{k \in K} (\Phi_k - l_{\text{ord}}(\Phi, \boldsymbol{\lambda}))^2. \quad (3)$$

234 If we set

$$\mathcal{L}(\Phi, l_{\text{ord}}(\Phi, \boldsymbol{\lambda})) = l_{\text{ord}}(\Phi, \boldsymbol{\lambda}), \quad (4)$$

235 we are simply minimizing the initial ordered measure defined by vector $\boldsymbol{\lambda}$.

236 3. Dealing with non-monotonicity and route inconsistency

237 For completeness, we start by reviewing the concept of monotonicity.

Definition 1 (Monotonicity). Let $\Phi, \Phi' \in \mathbb{R}_{\geq 0}^m$ be two vectors of route lengths, each associated with a feasible solution of the CVRP, where $\Phi = (\Phi_k)_{k \in K}$ and $\Phi' = (\Phi'_k)_{k \in K}$. Let the value of the fairness measure be defined by the loss function $\mathcal{L}(\Phi, l_{\text{ord}}(\Phi, \boldsymbol{\lambda}))$, where $l_{\text{ord}}(\Phi, \boldsymbol{\lambda})$ is an ordered measure. We say that $\mathcal{L}(\cdot)$ is monotone if the following condition holds. For any Φ' such that $\Phi'_k = \Phi_k + \delta_k$ for all $k \in K$, with $\delta_k \geq 0$ for all $k \in K$ and $\delta_k > 0$ for some $k \in K$, it follows that

$$\mathcal{L}(\Phi', l_{\text{ord}}(\Phi', \boldsymbol{\lambda})) \geq \mathcal{L}(\Phi, l_{\text{ord}}(\Phi, \boldsymbol{\lambda})).$$

238 In other words, worsening at least one route while keeping all others
239 no shorter cannot decrease the value of the fairness measure. Following
240 Definition 1, the mean, the median (or any other quantile) and the maximum
241 are monotonic measures.

242 Next, we adopt, without modification, the following companion notion of
243 Matl et al. (2018).

244 **Definition 2** (Workload inconsistency, Matl et al., 2018). Let $\Phi \in \mathbb{R}_{\geq 0}^m$ be
245 the route-length vector of a feasible solution of the CVRP, whose total cost is
246 $\sum_{k \in K} \Phi_k$ and whose value of the fairness measure is $\mathcal{L}(\Phi, l_{\text{ord}}(\Phi, \boldsymbol{\lambda}))$. We say
247 that Φ is workload inconsistent if there exists another feasible solution, with
248 route-length vector $\Phi' \in \mathbb{R}_{\geq 0}^m$, such that

$$\sum_{k \in K} \Phi_k > \sum_{k \in K} \Phi'_k \quad \text{and} \quad \Phi_k \geq \Phi'_k, \quad \forall k \in K,$$

249 and yet $\mathcal{L}(\Phi, l_{\text{ord}}(\Phi, \boldsymbol{\lambda})) < \mathcal{L}(\Phi', l_{\text{ord}}(\Phi', \boldsymbol{\lambda}))$. Otherwise, Φ is said to be
250 workload consistent. In words, Φ is workload inconsistent whenever it is
251 strictly costlier in total and no individual route is shorter, and yet the fairness
252 measure ranks it strictly better.

253 The notion is more demanding than comparing total costs, since it looks at
 254 every route length and not just their sum. Under a monotone measure (Defini-
 255 tion 1) no feasible solution can be workload inconsistent, because component-
 256 wise dominance already forces the measure not to decrease. Monotonicity also
 257 guarantees that some optimal CVRP solution consists of TSP-optimal tours
 258 (Matl et al., 2018). Non-monotone measures such as the range or the variance
 259 offer neither guarantee, so route inconsistency may arise and minimum-length
 260 routes have to be imposed explicitly.

261 Consider the CVRP instance of Table 2, with six customers, a central
 262 depot and two identical vehicles of capacity five. Table 3 reports the solutions
 263 obtained with and without TSP-optimality enforced on each route for three
 264 monotonic measures (mean, median and maximum) and four non-monotonic
 265 ones (range, Gini mean difference, mean absolute deviation from the median
 266 and variance). Figure 1 draws them. Under the monotonic measures the
 267 optimal solution is the cost-efficient one, whether or not TSP-optimality is
 268 imposed (Figure 1a). For the non-monotonic measures the two cases differ.
 269 Without TSP-optimality the model can make the two tours almost equal, but
 270 only by driving the second tour longer than its shortest tour, which increases
 271 the total distance considerably (Figure 1b). With TSP-optimality this route
 272 inconsistency disappears, and the optimal tours are TSP-optimal but less
 273 balanced (Figure 1c). The example shows two things. First, TSP-optimality
 274 removes route inconsistency but not workload inconsistency (Definition 2).
 275 Neither tour of the TSP-optimal fair solution is shorter than the corresponding
 276 tour of the cost-efficient solution, and yet every non-monotonic measure prefers
 277 the fair solution. Second, dropping TSP-optimality can increase the total
 278 cost, since the fair solution without TSP-optimality is longer in total than the
 279 one with it, so imposing it can even improve efficiency. Section 7.1 shows that
 280 this is not a peculiarity of the example, since every non-monotonic measure
 281 rewards such detours on every benchmark instance.

282 Our framework also admits negative entries in the weight vector λ , which
 283 is what most non-monotonic fairness measures in the literature require. The
 284 following result shows that, for the linear loss (4), negative weights are
 285 necessary as well as sufficient for non-monotonicity. In other words, such an
 286 objective admits a representation with nonnegative coefficients if and only
 287 if it is monotone. This delimits what formulations restricted to nonnegative
 288 weights can express. The equivalence is specific to the linear case, since the
 289 non-linear losses (2) and (3) may fail to be monotone even when $\lambda \geq 0$, as
 290 discussed after the theorem.

Table 2: Parameters of the route inconsistency example.

Nodes	0	1	2	3	4	5	6
Coordinates	(0,10)	(10,0)	(-10,0)	(0,-10)	(7.7,0)	(-7.7,0)	(0,-7.7)
Demands	0	1	1	3	1	1	3

Table 3: Results of the route inconsistency example for each fairness measure, without and with TSP-optimality. Obj is the value of the measure.

Measure	Monotonicity	no TSP-optimal				TSP-optimal			
		Obj	Tour 1	Tour 2	Total	Obj	Tour 1	Tour 2	Total
Mean	yes	47.05	45.03	49.06	94.09	47.05	45.03	49.06	94.09
Median	yes	47.05	45.03	49.06	94.09	47.05	45.03	49.06	94.09
Max	yes	49.06	45.03	49.06	94.09	49.06	45.03	49.06	94.09
Range	no	0.04	56.57	56.61	113.18	3.04	50.48	53.53	104.01
Gini	no	0.02	56.57	56.61	113.18	1.52	50.48	53.53	104.01
MAD _(med)	no	0.02	56.57	56.61	113.18	1.52	50.48	53.53	104.01
Variance	no	0.0004	56.57	56.61	113.18	2.31	50.48	53.53	104.01

291 **Theorem 1.** Let $l_{\text{ord}}(\Phi, \lambda) = \sum_{\ell \in K} \lambda_{\ell} \Phi_{(\ell)}$ be the ordered weighted objective
292 defined over the sorted route lengths $\Phi_{(1)} \leq \dots \leq \Phi_{(m)}$, and let the loss
293 function be the linear one given by (4), that is, $\mathcal{L}(\Phi, l_{\text{ord}}(\Phi, \lambda)) = l_{\text{ord}}(\Phi, \lambda)$.
294 Then $l_{\text{ord}}(\Phi, \lambda)$ is monotone if and only if $\lambda_{\ell} \geq 0$ for every $\ell \in K$.

295 The proof is given in Appendix A.1.

296 Consequently, for the linear loss function (4), route efficiency comes for
297 free when $\lambda \geq 0$, since some optimal solution is then TSP-optimal, whereas
298 TSP-optimality must in general be imposed when λ has at least one negative
299 entry (Theorem 1). This motivates the use of bilevel optimization models,
300 where, in the presence of non-monotonic measures, we must ensure that each
301 Φ_k represents the shortest route (i.e., the optimal TSP solution) for the set of
302 visited customers. Regarding the absolute and squared loss functions, (2) and
303 (3) respectively, there also exist cases where we need to impose TSP-optimality
304 in order to have optimal routes. This is the content of Remarks 1 and 2,
305 justified in Appendix A.2 and Appendix A.3.

306 **Remark 1.** TSP-optimality must be imposed to guarantee route efficiency
307 when the absolute deviation loss function (2) is minimized.

308 **Remark 2.** TSP-optimality must be imposed to guarantee route efficiency
309 when the squared deviation loss function (3) is minimized.




Figure 1: The three solutions of the route inconsistency example of Tables 2 and 3.

310 This analysis of monotonicity, route inconsistency and workload inconsis-
 311 tency yields a key insight. Monotonic fairness measures do not suffer from
 312 workload inconsistency and always admit an optimal solution free of route
 313 inconsistency, whereas non-monotonic measures may produce workload inconsis-
 314 tencies even with TSP-optimal routes. Monotonic measures have their own
 315 limitations, however, since, as outlined in the introduction, the min-max ob-
 316 jective captures less information about overall inequality than non-monotonic
 317 measures. Enforcing TSP-optimality of individual routes eliminates route
 318 inconsistency, although workload inconsistency may persist, implying a “price
 319 of fairness”. Neglecting route optimality only adds unnecessary inefficiency, so
 320 ensuring TSP-optimality is essential for efficient fairness and a fundamental
 321 step toward improving workload consistency in fair vehicle routing. Section 7
 322 returns to workload inconsistency and measures its extent on the solutions of
 323 our benchmark.

324 4. Bilevel optimization model

325 We now present a bilevel model that imposes TSP-optimality on every route
 326 by a single-leader multiple-disjoint-followers structure. The leader assigns the
 327 customers to the m vehicles, that is, it chooses a partition $\{V_1, \dots, V_m\}$ of
 328 the customers that respects the vehicle capacity. The assignment is described
 329 by binary variables x_i^k , with x_i^k equal to 1 if customer i is assigned to vehicle
 330 k . Each follower (vehicle) then solves a TSP over its customers. The model
 331 thus separates the assignment of customers from the requirement that each
 332 route be as short as possible for its customers.

$$\min \mathcal{L}(\Phi, \nu), \quad (5a)$$

$$\text{s.t. } \nu = \sum_{k \in K} \lambda_k \Phi_k, \quad (5b)$$

$$x_0^k = 1, \quad k \in K, \quad (5c)$$

$$\sum_{i \in V'} x_i^k \geq 1, \quad k \in K, \quad (5d)$$

$$\sum_{k \in K} x_i^k = 1, \quad i \in V', \quad (5e)$$

$$\sum_{i \in V} q_i x_i^k \leq Q, \quad k \in K, \quad (5f)$$

$$\sum_{k \in K} \Phi_k \leq B, \quad (5g)$$

$$\Phi_k \leq \Phi_{k+1}, \quad k \in K : k < m, \quad (5h)$$

$$x_i^k \in \{0, 1\}, \quad i \in V, k \in K, \quad (5i)$$

333 where Φ_k is the optimal value of the lower-level problem of vehicle k ,

$$\phi_k(x^k) = \min \sum_{i \in V} \sum_{\substack{j \in V \\ j \neq i}} c_{ij} z_{ij}^k, \quad (5j)$$

$$\text{s.t. } \sum_{i \in V} z_{ij}^k = x_j^k, \quad j \in V, \quad (5k)$$

$$\sum_{j \in V} z_{ij}^k = x_i^k, \quad i \in V, \quad (5l)$$

$$\sum_{i \in S} \sum_{j \in V \setminus S} z_{ij}^k \geq x_h^k, \quad \emptyset \neq S \subset V, 0 \notin S, h \in S, \quad (5m)$$

$$z_{ij}^k + z_{ji}^k \leq x_i^k, \quad i, j \in V', \quad (5n)$$

$$z_{ij}^k \in \{0, 1\}, \quad i, j \in V : i \neq j. \quad (5o)$$

334 The objective (5a) minimizes a loss $\mathcal{L}(\Phi, \nu)$, where $\nu = \sum_{k \in K} \lambda_k \Phi_k$ is
 335 the ordered aggregation $l_{\text{ord}}(\Phi, \boldsymbol{\lambda})$ of Section 2 written as a variable. For the
 336 linear loss, the objective is ν itself and constraint (5b) can be dropped.

337 Constraints (5c)–(5i) define the upper level. Constraints (5c) make every
 338 vehicle start and end at the depot, constraints (5d) ensure that every vehicle

is used, and constraints (5e) assign each customer to exactly one vehicle. Constraints (5f) impose the capacity, and (5g) bounds the total length of the routes by a budget B . Hereinafter, we write $x \in X$ for constraints (5c), (5d), (5e), (5f) and (5i). Finally, constraints (5h) sort the route lengths, as the ordered objective requires. Without the budget constraint (5g), a non-monotonic measure could be improved by making the routes much longer, and the planner would have no control over the extra distance spent on fairness. The budget is therefore the lever that sets the price of fairness.

Constraints (5j)–(5o) define the lower level, which makes Φ_k the length of a shortest tour over the customers assigned to vehicle k . The binary variable z_{ij}^k equals 1 if vehicle k uses arc (i, j) . The objective (5j) is the length of the tour, and constraints (5k) and (5l) make vehicle k enter and leave exactly the nodes assigned to it. The cutset inequalities (5m) forbid cycles that do not contain the depot. For every set S of nodes without the depot and every $h \in S$ assigned to vehicle k , at least one arc must leave S , and the degree constraints then give it an entering arc as well. Constraints (5n) forbid using both (i, j) and (j, i) , which removes degenerate patterns early in the search. We write “ (x, z) is a tour” for constraints (5k)–(5o).

For a fairness measure X , given by a loss $\mathcal{L}$ and a weight vector λ , problem (5) is the X -CVRP-TSP. For example, the Range-CVRP-TSP uses the linear loss (4) with $\lambda = (-1, 0, \dots, 0, 1)$, and the Variance-CVRP-TSP uses the squared loss (3) with $\lambda = \frac{1}{m}(1, \dots, 1)$.

If the measure is monotone, the bilevel structure is not needed, because some optimal solution is TSP-optimal. Replacing every route of an optimal solution by a shortest tour over its customers does not increase the objective, and keeps the budget satisfied. The problem can then be solved as a single-level model that includes constraints (5k)–(5o) and the equations

$$\Phi_k = \sum_{i \in V} \sum_{j \in V, j \neq i} c_{ij} z_{ij}^k, \quad k \in K. \quad (6)$$

Most fairness measures (see Table 3 and Remarks 1 and 2) are not monotone, and without the lower level the leader could lengthen a route to improve the measure. The bilevel model forbids this, since every driver drives a shortest route.

The bilevel structure has further advantages. The fairness measure appears only in the leader’s objective, while the followers’ problems depend only on the assignment. The reformulations and valid inequalities of Section 5, which

are derived from the lower level, are therefore the same for every measure, and a single implementation solves the X-CVRP-TSP for every measure of the framework. The leader-follower view also matches practice, where the planner allocates customers and the drivers, who would not follow a needless detour, choose their own shortest tours. Finally, new requirements on the plan can be added to the upper level without changing the routing part, for instance constraints on the redistribution of workload among drivers.

5. Handling the bilevel structure

This section turns the bilevel model (5) into a single-level mixed-integer program. Sections 5.1 and 5.2 present a dense reformulation, which keeps the routing variables, and a sparse one, which projects them out. Section 5.3 strengthens the sparse model with coupling inequalities, and Section 5.4 replaces the remaining non-convex constraint by valid inequalities. The proofs of all results of this section are given in Appendix A.

5.1. Dense value-function reformulation

The lower level (5j)–(5o) is not convex, so the bilevel problem cannot be reformulated with the Karush-Kuhn-Tucker conditions or with strong duality. We use instead the value-function reformulation

$$\min \mathcal{L}(\Phi, \nu), \tag{7a}$$

$$\text{s.t. } \nu = \sum_{k \in K} \lambda_k \Phi_k, \tag{7b}$$

$$x \in X, \tag{7c}$$

$$\sum_{k \in K} \Phi_k \leq B, \tag{7d}$$

$$\Phi_k \leq \Phi_{k+1}, \quad k \in K : k < m, \tag{7e}$$

$$(x, z) \text{ is a tour}, \quad k \in K, \tag{7f}$$

$$\sum_{i \in V} \sum_{j=0, j \neq i}^n c_{ij} z_{ij}^k = \Phi_k, \quad k \in K, \tag{7g}$$

$$\Phi_k \leq \phi_k(x^k), \quad k \in K, \tag{7h}$$

in which the only constraint inherited from the bilevel structure is the value-function constraint (7h). Constraints (7f) and (7g) already make Φ_k the

length of a tour over the customers of vehicle k , so $\Phi_k \geq \phi_k(x^k)$ always holds. Imposing TSP-optimality therefore only requires the upper bound $\Phi_k \leq \phi_k(x^k)$, which is non-convex.

One might hope to economize on separation effort by imposing the upper-bounding condition $\Phi_k \leq \phi_k(x^k)$ only for those $k \in K$ with $\lambda_k < 0$, since for $\lambda_k \geq 0$ a value of $\Phi_k > \phi_k(x^k)$ is never directly rewarded by the objective $l_{\text{ord}}(\Phi, \lambda) = \sum_{k \in K} \lambda_k \Phi_k$. The next proposition shows that this intuition is false.

Proposition 1. *The following statements hold.*

(a) Correctness of Model (7). *Model (7) is an exact reformulation of the bilevel problem (5).*

(b) Necessity of full enforcement. *Consider the linear loss (4), let $K^- = \{k \in K : \lambda_k < 0\}$ be non-empty and let $k_0 = \min K^-$. Dropping constraint (7h) for a single position $k \in K$ yields a strict relaxation of model (7), that is, a model whose optimal value is strictly smaller on some instances, if and only if $k \geq k_0$.*

The proof is given in Appendix A.4.

For the range, $k_0 = 1$, so (7h) is needed for every position, including those with $\lambda_k = 0$; the same holds for the Gini mean difference and the mean absolute deviation from the minimum, whose first weight is negative. For a monotone measure, K^- is empty and, as discussed in Section 4, the bilevel structure is not needed.

5.2. Sparse value-function reformulation

The dense model contains all the lower-level constraints and the $O(n^2m)$ binary routing variables z_{ij}^k , which can make it hard to solve. These variables can be projected out and replaced by constraints that describe the TSP cost as a function of the assignment x . Projecting z out of (7) gives the sparse model

$$\min \mathcal{L}(\Phi, \nu), \tag{8a}$$

$$\text{s.t. } \nu = \sum_{k \in K} \lambda_k \Phi_k, \tag{8b}$$

$$x \in X, \tag{8c}$$

$$\sum_{k \in K} \Phi_k \leq B, \tag{8d}$$

$$\Phi_k \leq \Phi_{k+1}, \quad k \in K : k < m, \quad (8e)$$

$$\phi_k(x^k) \leq \Phi_k, \quad k \in K, \quad (8f)$$

$$\Phi_k \leq \phi_k(x^k), \quad k \in K, \quad (8g)$$

where constraints (7f) and (7g) are replaced by the lower bound (8f). In the dense model this bound held automatically. Now that Φ_k is no longer tied to a tour, it has to be imposed explicitly. Constraint (8g) is the same as (7h) and is treated in Section 5.4.

Remark 3 (Correctness of the sparse model). If constraints (8f) and (8g) are both imposed, then $\Phi_k = \phi_k(x^k)$ for every $k \in K$, and model (8) is an exact reformulation of (5), as in the dense case. Proposition 1(b) also holds for the sparse model, with the same proof. It is even easier to violate here, since once (8g) is dropped for k , the continuous variable Φ_k is only bounded from below by (8f) and may take any value that the ordering and budget constraints allow.

Let $\mathcal{T}$ be the set of all routes in G . A route $T = (V_T, A_T) \in \mathcal{T}$ is given by its node set V_T , which contains the depot and at least one customer, and its arc set A_T . We write $C_T^* = \sum_{(i,j) \in A_T} c_{ij}$ for its length and $\text{TSP}(V_T)$ for the length of a shortest tour over V_T .

Constraint (8f) cannot be written directly, since $\phi_k(x^k)$ has no closed form. We replace it by the integer L-shaped cuts (Laporte and Louveaux, 1998)

$$\Phi_k \geq \underline{\Phi}_k + (\text{TSP}(V_T) - \underline{\Phi}_k) \left(\sum_{i \in V_T} x_i^k - |V_T| + 1 \right), \quad T \in \mathcal{T}, k \in K, \quad (9)$$

where $\underline{\Phi}_k$ is a lower bound on the length of any route (Appendix B.1), which is also used as a lower bound on the variable Φ_k . If vehicle k serves all customers of V_T , and possibly others, the cut gives $\Phi_k \geq \text{TSP}(V_T)$, because the TSP cost cannot decrease when customers are added. If a customer of V_T is missing, the right-hand side is at most $\underline{\Phi}_k$ and the cut says nothing new. Every route T that we separate, for this cut and for those of Section 5.4, is an optimal tour over the customers that a vehicle serves in an integer solution, which respects the capacity and serves at least one customer. All separated routes therefore satisfy $\underline{\Phi}_k \leq \text{TSP}(V_T) \leq \bar{\Phi}$, where $\bar{\Phi}$ is the upper bound on the length of any route used in Section 5.4.1 (Appendix B.1), so the gaps $\text{TSP}(V_T) - \underline{\Phi}_k$ and $\bar{\Phi} - \text{TSP}(V_T)$ that appear in the cuts are nonnegative.

449 Alternatively, a family of logic-based TSP-optimality cuts (Fachini and
 450 Armentano, 2020) remains informative when some customers of V_T are missing.
 451 The idea is that removing a customer i from a tour can shorten it by at most
 452 the cost of the two edges that i would need to be visited again. Since costs
 453 are symmetric, only the edges incident to i matter, and we write $c_i^{[1]} \geq c_i^{[2]}$
 454 for the two largest values among $\{c_{ij} : j \in V, j \neq i\}$. The cuts read

$$\Phi_k \geq \text{TSP}(V_T) - \sum_{i \in V_T} (1 - x_i^k) (c_i^{[1]} + c_i^{[2]}), \quad T \in \mathcal{T}, k \in K. \quad (10)$$

455 Validity follows from a sequential insertion argument. Starting from an
 456 optimal tour over the nodes assigned to vehicle k , each missing customer i
 457 of V_T can be inserted between two consecutive nodes $p \neq q$ of the current
 458 tour at a cost $c_{ip} + c_{iq} - c_{pq} \leq c_i^{[1]} + c_i^{[2]}$, and such a tour always has two
 459 distinct nodes because every vehicle serves at least one customer. This gives
 460 a tour over a superset of V_T , whose length is at least $\text{TSP}(V_T)$, and hence
 461 (10). Unlike (9), which only counts the missing customers and uses the gap
 462 $\underline{g}_k = \text{TSP}(V_T) - \underline{\Phi}_k$, this bound depends on which customers are missing.

463 **Proposition 2.** *For a route $T \in \mathcal{T}$ and $k \in K$, neither (9) nor (10)*
 464 *dominates the other in general.*

465 The proof is given in Appendix A.5.

466 5.3. Coupling inequalities

467 The cuts above are separated dynamically and enter the model only
 468 when an integer solution violates them. We now add two families of coupling
 469 inequalities, which are lower bounds on Φ_k that depend only on the assignment
 470 and on the data. They can be included in the initial model, so the linear
 471 relaxation uses them from the root node on.

472 The first family looks at pairs of customers. For $i \neq j \in V'$, let $T_{ij} =$
 473 $c_{0i} + c_{ij} + c_{j0}$ be the length of the tour through $\{0, i, j\}$. For every ordered
 474 pair $i \neq j$ and every $k \in K$ we impose the pair-envelope inequality

$$\Phi_k \geq \ell_{ij}(x^k) := 2c_{0i} x_i^k + (T_{ij} - 2c_{0i}) x_j^k, \quad i \neq j \in V', k \in K. \text{ (pair-env)}$$

475 In this subsection, an inequality $\Phi_k \geq h(x^k)$ is called valid if $h(\mathbf{1}_S) \leq \text{TSP}(S)$
 476 for every $S \subseteq V'$, where $\mathbf{1}_S$ is the incidence vector of S and $\text{TSP}(S)$ is the
 477 length of a shortest tour through the depot and S , that is, validity refers to
 478 the routing cost alone.

479 **Proposition 3.** *Inequalities (pair-env) are valid. Moreover, for every pair*
 480 *$i \neq j \in V'$ they are the strongest valid inequalities that involve only x_i^k and x_j^k ,*
 481 *in the sense that any valid inequality $\Phi_k \geq h(x^k)$ with h affine is dominated*
 482 *by (pair-env).*

483 The proof is given in Appendix A.6. In short, on the ground set $\{i, j\}$ the
 484 tour length $S \mapsto \text{TSP}(S)$ is submodular, because $T_{ij} \leq 2c_{0i} + 2c_{0j}$ by the
 485 triangle inequality; for a submodular function on two elements the Lovász
 486 extension coincides with the lower convex envelope (Lovász, 1983), and
 487 (pair-env) and its counterpart with i and j exchanged are precisely the two
 488 extended polymatroid inequalities (Edmonds, 1970) that describe it. Any
 489 valid inequality in x_i^k and x_j^k alone lies below this envelope.

490 The family has $m|V'|(|V'| - 1)$ inequalities with three nonzeros each
 491 and is added once to the initial model. The second coefficient, $T_{ij} - 2c_{0i} =$
 492 $c_{ij} + c_{j0} - c_{0i}$, is the cost of inserting j between i and the depot. The pair-
 493 envelope inequality is thus the lower-bounding counterpart of the insertion-
 494 based cuts of Section 5.4.

495 The second family gives constant lower bounds on the sorted route lengths.
 496 Let $D_{(1)} \geq D_{(2)} \geq \dots \geq D_{(|V'|)}$ be the depot costs c_{0i} , $i \in V'$, in non-
 497 increasing order, with $D_{(|V'|+1)} = 0$. For $t = 0, \dots, m - 1$, let $P^*(t)$ be the
 498 largest number of customers that t vehicles can serve, so that $P^*(0) = 0$.

499 **Proposition 4.** *Let $p_0, \dots, p_{m-1}$ be integers with $p_t \geq P^*(t)$ for every t .*
 500 *Every feasible solution of (8) satisfies*

$$\Phi_{m-t} \geq 2D_{(p_t+1)}, \quad t = 0, \dots, m - 1. \quad (\text{rank-lb})$$

501 The proof is given in Appendix A.7. In words, the t longest routes can serve
 502 at most p_t customers, so among the $p_t + 1$ customers farthest from the depot
 503 at least one is served by one of the other $m - t$ routes, and that route is at
 504 least as long as the round trip to this customer. The bound therefore applies
 505 to the $(m - t)$ -th shortest route, whatever the assignment. For $t = 0$ the
 506 longest route is at least twice the largest depot cost, for $t = 1$ the second
 507 longest is at least twice the $(p_1 + 1)$ -th largest depot cost, and so on.

508 We compute p_t from the total capacity. With the demands sorted as
 509 $q_{[1]} \leq q_{[2]} \leq \dots$, we set $p_t = \max\{p : \sum_{j=1}^p q_{[j]} \leq tQ\}$. This gives $p_t \geq P^*(t)$,
 510 because any $P^*(t)$ customers served by t vehicles have a total demand of at
 511 most tQ , which is at least the sum of the $P^*(t)$ smallest demands. The rank
 512 bounds are not implied by (pair-env) at fractional points. At $x_i^k = 1/m$ for

all i and k , for example, (pair-env) only gives $\Phi_k \geq \max_{i \neq j} T_{ij}/m \leq 4D_{(1)}/m$,
 whereas (rank-lb) with $t = 0$ gives $\Phi_m \geq 2D_{(1)}$.

The pair-envelope construction extends to triples of customers, for which
 the TSP cost is still submodular, but not to larger sets (Hrerer and Penn, 1995;
 Herer, 1999). The resulting triple-envelope inequalities dominate (pair-env),
 but there are $m|V'|(|V'| - 1)(|V'| - 2)$ of them, too many to include in the
 initial model, and when we separated them as cutting planes they did not
 solve any of our hardest open instances. We therefore use (pair-env) and
 (rank-lb) as coupling inequalities.

5.4. Upper-bounding bilevel cuts

We now deal with the non-convex constraint $\Phi_k \leq \phi_k(x^k)$, that is, (7h) in
 the dense model and (8g) in the sparse one. We replace it by two families
 of valid inequalities. The first one is obtained from tour enumeration (Sec-
 tion 5.4.1) and the second one from the separated route, by charging insertion
 and removal penalties (Section 5.4.2).

Both families use the insertion penalty

$$M_i^1 = \max \{0, \max \{c_{ri} + c_{it} - c_{rt} : r, t \in V, r \neq t, i \notin \{r, t\}\}\},$$

defined for every node i , which bounds the cost of inserting i between two
 consecutive nodes of any tour. The second family also uses a removal penalty,
 defined in Section 5.4.2 with respect to the separated route, which plays the
 opposite role and makes it too expensive for a tour to pass through a customer
 that is not assigned to it.

The bounds that both penalties provide are proved in the appendix.
 The insertion bound of M_i^1 is shown in Appendix A.8, and Lemma 1 in
 Appendix A.9 combines both penalties to bound the length of a shortest tour.

5.4.1. No-good cuts

The TSP cost cannot decrease when customers are added, so $\text{TSP}(S) \leq$
 $\text{TSP}(S')$ for $S \subseteq S'$. If the customers of vehicle k are a subset of V_T , then
 $\phi_k(x^k) \leq \text{TSP}(V_T)$, which gives the no-good cuts

$$\Phi_k \leq \text{TSP}(V_T) + (\bar{\Phi} - \text{TSP}(V_T)) \sum_{i \in V \setminus V_T} x_i^k, \quad T \in \mathcal{T}, k \in K, \quad (\text{bilevel-UB})$$

where $\bar{\Phi}$ is an upper bound on the length of any route (Appendix B.1). If a
 customer outside V_T is assigned to k , the right-hand side is at least $\bar{\Phi}$ and
 the cut is inactive.

544 This cut charges the same penalty $\bar{\Phi} - \text{TSP}(V_T)$ for every customer outside
 545 V_T , so a single extra customer deactivates it. Charging instead each extra
 546 customer its insertion penalty M_i^1 gives the cut

$$\Phi_k \leq \text{TSP}(V_T) + \sum_{i \in V \setminus V_T} M_i^1 x_i^k, \quad T \in \mathcal{T}, \quad k \in K, \quad (\text{bilevel-Ins})$$

547 which is valid because the extra customers can be inserted one by one into an
 548 optimal tour over V_T , each at a cost of at most M_i^1 . The same M_i^1 appears
 549 again in Section 5.4.2, where it is the penalty that decouples the lower level
 550 from the assignment. Neither cut dominates the other, since M_i^1 can be
 551 smaller or larger than $\bar{\Phi} - \text{TSP}(V_T)$, but taking the smaller of the two for
 552 each customer gives a cut that is never weaker than either.

553 **Proposition 5.** *Let $T \in \mathcal{T}$ and $g = \bar{\Phi} - \text{TSP}(V_T)$. The capped insertion cut*

$$\Phi_k \leq \text{TSP}(V_T) + \sum_{i \in V \setminus V_T} \min\{g, M_i^1\} x_i^k, \quad T \in \mathcal{T}, \quad k \in K, \quad (\text{bilevel-Cap})$$

554 *is valid for every feasible assignment x^k , and it dominates (bilevel-UB) and*
 555 *(bilevel-Ins).*

556 The proof is given in Appendix A.8.

557 5.4.2. Cuts with insertion and removal penalties

558 The no-good cuts keep the full length $\text{TSP}(V_T)$ of the separated route and
 559 only charge the customers added to it. The second family uses the route T
 560 itself and also accounts for the nodes of T that are not assigned to vehicle
 561 k . The idea is to transform T into a tour over the nodes assigned to k in
 562 two steps, dropping the nodes of T that are not assigned and inserting the
 563 assigned nodes that T does not visit, and to bound the length of the resulting
 564 tour, which is at least $\phi_k(x^k)$. The length of T is charged arc by arc, and only
 565 for the arcs that leave an assigned node i , through the term $\sum_{(i,j) \in A_T} x_i^k c_{ij}$.
 566 Dropping a node i of T between its neighbours r and t replaces the path
 567 $r \rightarrow i \rightarrow t$ by the arc (r, t) ; the arc term still pays c_{ri} but no longer c_{it} ,
 568 so the missing amount $c_{rt} - c_{ri}$ is covered by a removal penalty of at least
 569 $c_{rt} - c_{ri} + c_{it}$. For $i \in V_T$, let

$$M_i^{2,T} = \max \left\{ 0, \max \{ c_{rt} - c_{ri} + c_{it} : r, t \in V_T, \quad r \neq t, \quad i \notin \{r, t\} \} \right\},$$

570 and let $M_d^{2,T} = c_{0d}$ when $V_T = \{0, d\}$. Inserting an assigned customer $u \notin V_T$
 571 costs at most M_u^1 , as for the no-good cuts. For a route $T \in \mathcal{T}$ this gives the
 572 cut

$$\Phi_k \leq \sum_{(i,j) \in A_T} x_i^k c_{ij} + \sum_{i \notin V_T} M_i^1 x_i^k + \sum_{i \in V_T} M_i^{2,T} (1 - x_i^k), \quad T \in \mathcal{T}, \quad k \in K.$$

(bilevel-L)

573 **Proposition 6.** *Let $T \in \mathcal{T}$ and $k \in K$. Inequality (bilevel-L) is valid for*
 574 *every feasible assignment x^k , and the inequality built from an optimal tour*
 575 *over the nodes assigned to vehicle k is tight at x^k .*

576 Tightness is immediate and the rest of the proof is given in Appendix A.9.

577 **Proposition 7.** *For a route $T \in \mathcal{T}$ and $k \in K$, neither (bilevel-Cap) nor*
 578 *(bilevel-L) dominates the other in general.*

579 The proof is given in Appendix A.10.

580 We observe that computing M^1 with respect to V_T as well can cut off
 581 feasible solutions (see Example 1 in Appendix A.11). On the other hand,
 582 replacing the arc-cost term $\sum_{(i,j) \in A_T} x_i^k c_{ij}$ by the constant C_T^* can only weaken
 583 the cut.

584 6. Computational study

585 This section evaluates the models and inequalities of Section 5. Section 6.1
 586 describes the experimental setting. Section 6.2 chooses between the dense and
 587 the sparse formulation and among the upper-bounding cuts, and Section 6.3
 588 measures the effect of the coupling inequalities. Section 6.4 compares the
 589 resulting method with the state of the art on the Range-CVRP-TSP, and
 590 Section 6.5 compares the fairness measures with each other.

591 6.1. Experimental setting

592 All configurations are branch-and-cut algorithms based on one of the two
 593 formulations of Section 5, in which the cuts on the route lengths are added as
 594 lazy cuts at integer solutions; the separation, the TSP subproblems and the
 595 caching of their optimal tours are described in Appendix B.2. The bounds
 596 $\underline{\Phi}_k$ and $\bar{\Phi}$ used by the cuts are computed once per instance (Appendix B.1).
 597 Every run is warm-started with the cost-efficient solution, which is feasible
 598 for every budget.

Table 4: Families of valid inequalities and the formulations in which they are used.

Family	Inequalities	Variant	Formulation	Role
No-good	(bilevel-UB)	uniform penalty	both	$\Phi_k \leq \phi_k(x^k)$
Insertion	(bilevel-Ins)	penalty M_i^1	both	$\Phi_k \leq \phi_k(x^k)$
Capped insertion	(bilevel-Cap)	penalty $\min\{g, M_i^1\}$	both	$\Phi_k \leq \phi_k(x^k)$
Insertion and removal penalties	(bilevel-L)	route-restricted $M^{2,T}$	both	$\Phi_k \leq \phi_k(x^k)$
Lower-bounding	(9), (10)	L-shaped or logic-based	sparse	$\Phi_k \geq \phi_k(x^k)$
Pair envelope	(pair-env)	convex envelope on pairs	both	coupling
Rank bound	(rank-lb)	depot costs and capacity	both	coupling
Subtour elimination	(5m)	per vehicle or aggregated	dense	route feasibility

Table 4 lists the inequalities of Section 5. TSP-optimality can be imposed with (bilevel-Cap) (which dominates the other two no-good cuts, namely (bilevel-UB) and (bilevel-Ins)), or with the penalty cut (bilevel-L). Since neither (bilevel-Cap) nor (bilevel-L) dominates the other (Proposition 7), we compare these two. The lower-bounding cuts (9) and (10) are only needed in the sparse formulation, where we always use both because neither dominates the other (Proposition 2). In the dense formulation, constraint (7g) already ensures that $\Phi_k \geq \phi_k(x^k)$.

The default configuration is the sparse formulation with both lower-bounding cuts and, for the non-monotonic measures, the capped insertion cut (bilevel-Cap). For the monotonic measures MAX and MEDIAN, no upper-bounding cut is used, since these measures are monotone (Theorem 1) and the bilevel structure is then not needed (Section 4). The strengthened configuration adds the coupling inequalities (pair-env) and (rank-lb) to the default one. Sections 6.2 and 6.3 motivate these choices, and the rest of the study uses the strengthened configuration.

We use the benchmark of van Rossum et al. (2026a), see also van Rossum et al. (2026b). Each instance is obtained from the CVRPLIB instance X-n641-k35 (Uchoa et al., 2017) by taking $n + 1$ consecutive nodes, the first of which is the depot, with $m = 5$ vehicles, capacity $Q = \lfloor \sum_i q_i / 4 \rfloor$ and rounded-up Euclidean distances. There are 20 instances for each $n \in \{15, 20, 25, 30\}$. We generated 17 more instances with $n = 35$ by the same rule, which is the number of disjoint blocks of 36 nodes that the original instance contains. The budget B is 101, 105 or 110% of the minimum total length of a solution with five routes, which we computed to proven optimality for every instance. The budget therefore says how much extra distance the fleet may drive to obtain a fairer plan. The time limit is 10 minutes for $n = 15$ and one hour otherwise.

626 All runs were made on a MacBook Pro with an Apple M4 processor and
627 24 GB of RAM, using Gurobi 13.0 through its Python interface with default
628 parameters and one run at a time. In the figures, $|C|$ denotes the number of
629 customers n . Claude Fable 5.1 (Anthropic) assisted in writing the code and
630 running the computational experiments, and the authors verified all code and
631 results.

632 6.2. Dense versus sparse formulation and choice of the upper-bounding cut

633 Figure 2 compares the dense formulation, with and without max-flow
634 separation of the subtour cuts, and the sparse formulation on the Range-
635 CVRP-TSP with $n = 15$ and $B = 110\%$. All 60 runs end with a proof of
636 optimality. The sparse formulation is clearly faster, with a median time of
637 4.4 s against 70.1 s and 79.0 s for the two dense variants (16 to 18 times faster).
638 Removing the routing variables is therefore the main source of speed, well
639 ahead of any strengthening of the dense model, as the max-flow cuts illustrate.
640 With about 40 cuts per run they reduce the mean time of the dense model
641 by 10% (from 94.6 s to 85.0 s). About half of this saving comes from a single
642 instance whose time drops from 292 s to 190 s. However, they help on only 12
643 of the 20 instances, and the median time even rises, from 70.1 s to 79.0 s.

644 The right panel of Figure 2 shows that the sparse model is faster even
645 though it adds many more cuts, about 3 300 per run in the median against 40
646 to 80 for the dense variants. Each node of the dense model carries all routing
647 variables and is expensive to solve, whereas the linear programs of the sparse
648 model are small. Therefore, in the rest of the study we use the sparse model.

649 Next we choose the upper-bounding cut. Table 5 compares the capped
650 insertion cut (bilevel-Cap), the penalty cut (bilevel-L) and their combination
651 on the five non-monotonic measures, which are the only ones that need such a
652 cut. At $n = 15$ all runs finish in less than 20 seconds. The capped insertion cut
653 has the lowest median time, (bilevel-L) is slightly slower and the combination
654 lies in between. At $n = 20$, with a common limit of 1 000 s, the three are
655 close. The capped insertion cut is still the fastest in the median, while the
656 combination solves three more of the hardest runs within the limit. Since
657 both families are separated on the same tours, the combination adds roughly
658 the sum of their cuts, and its slightly better tail is paid for with slower typical
659 runs. We therefore keep (bilevel-Cap) in the default configuration.




Figure 2: Dense and sparse formulations on the Range-CVRP-TSP, $n = 15$, $B = 110\%$, 20 instances per configuration, all solved to optimality. Left: distribution of the solution time (log scale). Right: empirical distribution of the number of cuts added per run, summed over all active families (log scale).

Table 5: Upper-bounding cuts in the sparse formulation, five non-monotonic measures with 20 instances each, $B = 110\%$. Median solution time, average number of upper-bounding cuts per run and number of runs not solved within the time limit (600 s for $n = 15$ and 1 000 s for $n = 20$).

Cut	$n = 15$			$n = 20$		
	Median (s)	Cuts/run	Unsolved	Median (s)	Cuts/run	Unsolved
(bilevel-Cap)	4.51	218.7	0/100	196.0	800.9	12/100
(bilevel-L)	5.66	237.7	0/100	222.3	947.5	10/100
both	4.69	434.1	0/100	226.8	1 587.9	9/100

6.3. Effect of the coupling inequalities

Table 6 shows what happens when the rank bounds (rank-lb), the pair-envelope inequalities (pair-env), or both, are added to the default configuration, for every measure with $B = 110\%$ and $n \in \{15, 20\}$. All variants reach the same optimal values.

On their own, the rank bounds give a moderate improvement. At $n = 15$ they almost halve the search tree and the time, and at $n = 20$ they remove about a third of the nodes and of the time. Their effect depends on the measure. It is large for RANGE and GINI, whose time at $n = 20$ falls from 217 to 132 s and from 410 to 191 s. It disappears for MAX, which even becomes slower (166 to 206 s).

The pair-envelope inequalities have a much larger effect. At $n = 15$ the

Table 6: Effect of the coupling inequalities for each fairness measure at $B = 110\%$. Average solution time (s) and average number of nodes over 20 instances.

n	Measure	default		+rank		+pair-env		+pair-env+rank	
		Time (s)	Nodes	Time (s)	Nodes	Time (s)	Nodes	Time (s)	Nodes
15	MAX	2.39	36,504	1.81	29,584	0.15	862	0.08	821
	MEDIAN	6.30	17,053	1.36	25,096	0.20	237	0.11	386
	RANGE	5.16	97,817	2.55	42,657	0.21	1,599	0.15	1,382
	GINI	6.07	114,268	3.07	50,940	0.21	1,856	0.15	1,638
	MAD(MED)	5.70	105,688	3.64	57,350	0.23	2,230	0.18	2,239
	MAD(MIN)	5.89	121,007	3.59	62,002	0.23	2,273	0.15	1,754
	VARIANCE	5.87	124,876	3.82	70,716	0.19	1,798	0.16	1,421
20	MAX	166.0	395,987	205.8	493,849	0.83	4,401	0.39	4,312
	MEDIAN	278.9	730,993	223.0	585,261	1.02	4,229	0.74	5,174
	RANGE	216.9	625,032	131.7	380,425	0.76	7,078	1.37	7,174
	GINI	410.0	995,295	190.7	493,655	0.83	7,600	0.70	7,924
	MAD(MED)	402.5	860,700	267.2	646,143	1.10	9,510	0.98	9,484
	MAD(MIN)	457.6	1,279,098	284.3	767,674	1.09	9,086	0.83	8,367
	VARIANCE	454.4	1,342,094	295.6	814,052	1.04	10,556	0.98	10,061

search tree becomes 57 times smaller on average over the seven measures (42 to 72 times for each measure), and the average time drops from 5.3s to 0.2s. At $n = 20$ the tree becomes 119 times smaller (88 to 173 times for each measure), and the average time falls from about six minutes to about one second. The reason lies in the linear relaxation. Without (pair-env) nothing links Φ_k to the assignment until lazy cuts accumulate, so the root bound of the range is zero and the search has to enumerate assignments. With them, the root bound already reaches 50 to 70% of the optimal value, and the number of lazy cuts per run drops from thousands, or tens of thousands at $n = 20$, to a few hundred. The improvement is similar for all measures. The hardest ones for the default configuration, GINI, both MAD measures and VARIANCE, which need more than 400s at $n = 20$, now take about one second like the others. Adding the rank bounds on top of the pair-envelope inequalities changes the size of the search tree only moderately, by at most 23% except for MEDIAN at $n = 15$, which hardly matters at these times, and we keep them because they cost almost nothing.

6.4. Comparison with the state of the art on the Range-CVRP-TSP

The range is the only measure for which another exact method exists. van Rossum et al. (2026a) solve the Range-CVRP-TSP with a branch-and-price algorithm and handle TSP-optimality either by postprocessing or, exactly but at a prohibitive cost on larger instances, inside the pricing problem. van

Table 7: Comparison with the state of the art on the Range-CVRP-TSP (20 instances per size and 17 for $n = 35$, time limit of 10 minutes for $n = 15$ and one hour otherwise).

n	B (%)	Strengthened configuration			van Rossum et al. (2026b)		
		#Solved	Time (s)	Gap (%)	#Solved	Time (s)	Gap (%)
15	101	20/20	0.11	0.00	20/20	4	0.00
	105	20/20	0.14	0.00	20/20	7	0.00
	110	20/20	0.15	0.00	20/20	11	0.00
20	101	20/20	0.48	0.00	20/20	19	0.00
	105	20/20	0.71	0.00	20/20	56	0.00
	110	20/20	1.37	0.00	20/20	159	0.00
25	101	20/20	103.6	0.00	20/20	299	0.00
	105	20/20	7.6	0.00	20/20	468	0.00
	110	20/20	9.0	0.00	16/20	1 310	0.27
30	101	19/20	222.7	1.49	20/20	772	0.00
	105	20/20	33.3	0.00	17/20	1 419	0.07
	110	20/20	27.2	0.00	15/20	1 747	0.34
35*	101**	12/17	1 466	7.20	—	—	—
	105	11/17	1 608	6.92	—	—	—
	110	15/17	834	0.53	—	—	—

Unsolved runs enter the averages with the time limit and their final gap. Budgets are relative to the minimum cost of each instance. The right columns are the figures reported by van Rossum et al. (2026b). *New instances, for which no other results exist. **Better of two runs per instance, with and without the two-index variables and constraints described in the text.

693 Rossum et al. (2026b) then impose it exactly with lifted TSP-optimality
694 cuts in a branch-price-and-cut algorithm, which is the current state of the
695 art. Table 7 compares our strengthened configuration with their best setting
696 (“Cuts + Lifting”, CPLEX 22.1, one hour per instance on an AMD Rome
697 7H12 processor with 16 GB of RAM) on the same instances and budgets, and
698 adds our results for $n = 35$.

699 Up to $n = 20$ both methods solve every instance, and ours is 36 to 116
700 times faster on average (0.1 to 1.4 s against 4 to 159 s). The difference grows
701 with the size of the instances. At $n = 25$ we solve all 60 runs in 8 to 104 s
702 on average, whereas the branch-price-and-cut needs 299 to 1 310 s and leaves
703 four instances open at $B = 110\%$. Our average of 104 s at $B = 101\%$ comes
704 from a single instance that takes 2 011 s. At $n = 30$ we solve 59 of the 60
705 runs in 27 to 223 s on average, against 772 to 1 747 s and 52 solved runs for
706 van Rossum et al. (2026b). They do solve the run we leave open, instance
707 14 at $B = 101\%$, which ends with a gap of 30% after one hour in our case.
708 The two methods ran on different machines, but differences of one to two

orders of magnitude cannot be explained by hardware alone. The gaps of 80 to 90% at $n = 25$ that van Rossum et al. (2026b) report for an earlier version of this work were obtained with a preliminary implementation of the default configuration.

Instances with 35 customers are at the limit of what the method solves in one hour. With $B = 110\%$, 15 of the 17 instances are solved in 17 to 1 465 s, and the other two end with gaps of 0.2% and 8.7%. With $B = 105\%$, 11 instances are solved and the remaining six end with gaps between 3% and 45%. Compared with the cost-efficient plan, the best fair plans found reduce the range by 39% on average at $B = 105\%$ and by 62% at $B = 110\%$. The tight budget $B = 101\%$ is the hardest case, with only 8 instances solved. For this budget we also tried adding to the sparse model two-index edge variables y_e linked to the assignment, together with the budget written on the edges, the inequality $\sum_{k \in K} \Phi_k \geq \sum_e c_e y_e$ and the usual capacity and connectivity cuts. The two-index variables make each node slower but the relaxation stronger. At $B = 105\%$ this extended model is 5 to 18 times slower on the four instances that both versions solve. At $B = 101\%$, however, it solves 10 instances, four of which the base version cannot solve, while it loses two that the base version solves within 12 minutes. Keeping the better of the two runs, as in Table 7, gives 12 solved instances out of 17, and the five open ones end with gaps above 2% in both versions.

Table 7 only covers the range because no exact method exists for the other measures, and the method of van Rossum et al. (2026b) is built around the range. Our method does not depend on the measure. The results for the other six measures in Section 6.5 are obtained with the same code, where only λ and the loss change and the upper-bounding cuts are switched off for the monotonic measures. With $B = 110\%$, the strengthened configuration solves every instance with up to 30 customers for all seven measures. To our knowledge, it is the first exact method for the X-CVRP-TSP beyond the range.

6.5. Comparison of fairness measures

Figure 3 shows the solution times of the seven measures with the strengthened configuration and $B = 110\%$. All 560 runs end with a proof of optimality, so the figure compares solution times only, resulting in three observations.

First, the size of the instance matters much more than the measure. The median time over all measures grows by a factor of 3 to 5 with every five additional customers, from 0.1–0.2 s at $n = 15$ to 7–24 s at $n = 30$, or 34 s for

746 MAX. Up to $n = 25$, the medians of the seven measures differ by at most a
 747 factor of 2.5 for a given size.

748 Second, the measures differ in their worst cases rather than in their typical
 749 behavior. At $n = 25$ the slowest run takes between 23 and 68 s for MAD(MIN),
 750 MAX, MAD(MED) and GINI, but 128 s for RANGE, 186 s for MEDIAN and 387 s
 751 for VARIANCE. At $n = 30$ the slowest run takes at most 11 minutes for six
 752 measures, while MAD(MED) needs up to 29 minutes. MAX and MAD(MED)
 753 also have the heaviest tails, with a tenth of their runs above 9 minutes, against
 754 less than 2 minutes for the other measures.

755 Third, some instances are hard for every measure. Instance 9 is the slowest
 756 instance with $n = 20$ for five of the seven measures, and instance 14 is the
 757 slowest with $n = 25$ for five measures and with $n = 30$ for three. Instance 14
 758 is also by far the hardest instance with $n = 25$ for the minimum-cost problem,
 759 and instance 9 one of the hardest with $n = 20$. The difficulty therefore lies
 760 mainly in proving that the assignment is optimal under the budget, which is
 761 common to all measures.

762 The monotonic measures MAX and MEDIAN do not need upper-bounding
 763 cuts, yet they are not easier than the others. Up to $n = 25$ they behave like
 764 the non-monotonic measures, and at $n = 30$ MAX has the largest median
 765 time (34 s against 7 to 24 s) and the largest search trees (122 000 nodes in
 766 the median, against 41 000 to 111 000). Skipping the upper-bounding cuts
 767 does help. Separating (bilevel-Cap) anyway, which is valid but unnecessary,
 768 multiplies the median time at $n = 30$ by 2.7 for MAX and by 1.5 for MEDIAN,
 769 with the same optimal values. This behavior also suggests that, when only
 770 one of these measures is of interest, a more tailored formulation may be
 771 preferable. Since their weight vectors contain only a single nonzero entry,
 772 the model unnecessarily incurs the burden of sorting elements that do not
 773 contribute to the objective.

774 7. Managerial insights

775 This paper shows that fair plans can now be computed exactly at realistic
 776 sizes and for many measures. This section uses them to answer three questions
 777 that a planner faces before adopting such a plan. Throughout the section, the
 778 cost-efficient plan is the minimum-cost solution with five routes that defines
 779 the budgets (Section 6.1). Its routes are TSP-optimal, and we also refer to
 780 it as the plan with $B = 100\%$. A fair plan is an optimal solution of the
 781 X-CVRP-TSP for a given budget B . Since the vehicles are identical, a driver




Figure 3: Solution time per fairness measure and instance size (strengthened configuration, $B = 110\%$, 20 instances per box, logarithmic scale). The dashed line is the time limit.

782 has no identity across plans. We therefore compare routes by their position in
 783 the sorted order and say that the driver in position k is worse off if the k -th
 784 shortest route of the fair plan is longer than that of the cost-efficient plan. In
 785 all our instances the five drivers are on duty, since every vehicle must be used
 786 (constraints (5d)).

787 The first question is how much fairness a given budget buys. Over the
 788 instances with up to 25 customers, allowing 1, 5 and 10% more distance
 789 than the cost-efficient plan reduces its range by 5–11%, 25–37% and 47–59%
 790 on average. The effect varies a lot between instances. At $B = 110\%$ the
 791 reduction goes from 19% to 91%, and at $B = 101\%$ the fair plan is often the
 792 cost-efficient plan itself. The second question is whether TSP-optimality may
 793 be left out of the model for the sake of simplicity; Section 7.1 shows what this
 794 costs. The third question is who pays for this fairness. For a non-monotonic
 795 measure, the measure can improve while most drivers drive more than in the
 796 cost-efficient plan. In its extreme form this is the workload inconsistency of
 797 Definition 2, which TSP-optimality does not prevent (Section 3). Section 7.2
 798 shows that the plan can become much worse for the drivers without imposing
 799 the budget constraint.

800 7.1. The price of ignoring TSP-optimality

801 Matl et al. (2018) ask “to what extent workload inconsistency and non-
 802 TSP-optimality are a relevant concern in practice for nonmonotonic measures”,

Table 8: The model without TSP-optimality enforcement against the exact one (strengthened configuration, 20 instances per row). Route inconsistent: runs whose relaxed optimum is strictly below the true optimum. Overstatement: relative gap between the two optimal values. Post-processing: runs on which the repaired plan is not optimal, and its relative gap to the optimum (average and maximum, in %).

n	measure	B (%)	relaxed model			post-processing		
			route inc.	avg	max	subopt.	avg	max
15	MAX	110	0/20	0.0	0.0	0/20	0.0	0.0
	MEDIAN	110	0/20	0.0	0.0	0/20	0.0	0.0
	RANGE	101	20/20	3.6	7.9	0/20	0.0	0.0
	RANGE	105	20/20	8.6	19.8	8/20	1.9	16.4
	RANGE	110	20/20	12.9	41.4	16/20	13.5	56.6
	GINI	110	20/20	9.5	36.6	14/20	7.1	45.1
	MAD(MED)	110	20/20	13.4	39.7	14/20	10.4	56.6
	MAD(MIN)	110	20/20	22.0	60.4	14/20	20.4	144.6
	VARIANCE	110	20/20	15.1	52.0	16/20	14.6	64.1
20	MAX	110	0/20	0.0	0.0	0/20	0.0	0.0
	MEDIAN	110	0/20	0.0	0.0	0/20	0.0	0.0
	RANGE	101	20/20	3.5	8.3	0/20	0.0	0.0
	RANGE	105	20/20	6.5	17.7	12/20	4.6	38.6
	RANGE	110	20/20	12.0	90.9	16/20	12.6	118.6
	GINI	110	20/20	10.9	93.3	18/20	14.5	104.9
	MAD(MED)	110	20/20	13.2	94.4	14/20	24.7	121.1
	MAD(MIN)	110	20/20	20.4	95.4	16/20	24.0	174.3
	VARIANCE	110	20/20	15.0	99.1	16/20	27.5	273.6

and their enumeration study on instances with up to 14 customers found both to be frequent among Pareto-optimal solutions of the bi-objective problem. The first part of the question concerns workload inconsistency, to which we return in Section 7.2; here we address the second part. We solved every run of the $n \in \{15, 20\}$ grid a second time with the same strengthened configuration but without any TSP-optimality enforcement, that is, with constraint (7h) dropped for every vehicle. The optimal value of this relaxed model is a lower bound on the true optimum. We then repaired the relaxed plan by replacing every route by the shortest tour over its customers and evaluated the measure on the repaired plan. Table 8 reports, per size, measure and budget, how often the relaxed plan is route inconsistent, by how much the relaxed model overstates the achievable fairness, and how far the repaired plan is from the true optimum.

For the five non-monotonic measures, on every one of the 280 runs the relaxed plan contains at least one route that is not a shortest tour. The fairness that the relaxed model claims is overestimated by 12% on average,

by 3–4% for the RANGE at $B = 101\%$ and by 20–22% for MAD(MIN), and by up to 99% on single runs. Repairing the routes afterwards does not recover the optimum either. At $B = 110\%$ the post-processed plan is suboptimal on 70–90% of the runs, by 7–28% on average and by up to 274%, because the assignment chosen by the relaxed model was only good together with its inflated routes. At $B = 105\%$ the post-processed RANGE plan is still suboptimal on 8 of the 20 instances with $n = 15$ and on 12 of the 20 with $n = 20$, by up to 16% and 39% respectively. Only at $B = 101\%$, where the slack leaves little room for detours, did post-processing return the optimal plan on all of our runs (the 40 RANGE runs with 15 and 20 nodes).

7.2. Who pays for fairness

With a non-monotonic measure, a plan can be Pareto-efficient for cost and fairness and yet workload inconsistent: every driver drives more than in some other feasible plan. Matl et al. (2018) prove that monotonic measures exclude the phenomenon and that TSP-optimality does not (their Observation 3). In this section we take the cost-efficient plan as the reference and compare the individual route lengths of the optimal fair plans of the benchmark, all of them with TSP-optimal routes, with those of the reference plan, driver by driver.

Over the 40 instances with $n \in \{15, 20\}$ and the range as measure, four of the 120 optimal fair plans are dominated route by route by the cost-efficient plan itself, at $B = 105\%$ and 110% , and are therefore workload inconsistent although every route is a shortest tour. On instance 18 with $n = 20$ and $B = 110\%$, the sorted route lengths of the cost-efficient plan are (221, 780, 1332, 1431, 1577) with range 1356, and those of the fair plan are (732, 786, 1332, 1431, 1577) with range 845. No driver is better off, the longest routes are kept and two shorter routes become longer. Definition 2 compares a plan with every cheaper feasible plan and not only with the cost-efficient one, so this count is only a lower bound. Relative to the cost-efficient plan, at least one driver improves in all other runs in which the fair plan differs from the cost-efficient one, but most drivers can still be worse off, as instance 15 with $n = 20$ shows (Figure 4). Sorted from shortest to longest, the routes of the cost-efficient plan ($B = 100\%$) have lengths (732, 1318, 1713, 1761, 2213), with range 1481 and total cost 7737. At $B = 110\%$ the range drops to 612, with lengths (1464, 1478, 1724, 1767, 2076), but 4 of the 5 drivers are worse off. The shortest route doubles, and only the driver of the longest route benefits. This optimal plan is not unique. Another plan with the same

range, found by the strengthened configuration, leaves only two drivers worse off, so the choice among optimal fair plans also matters for the drivers. At $B = 105\%$ the range is 967, with lengths (1116, 1324, 1761, 1838, 2083), and again 4 drivers are worse off. We also solved this instance with $B = 120\%$, which is outside the grid of Section 6.1. The range then falls to 132, with lengths (1796, 1818, 1868, 1873, 1928) and optimality proven in 179 s. The fleet is then almost perfectly balanced, but still 4 drivers are worse off. Only at $B = 101\%$ is the picture different. The fair plan keeps three of the five cost-efficient routes unchanged, only one driver is worse off, and the range is 1351. Without any budget, that is, without constraint (5g), the plan of instance 15 becomes workload inconsistent in the most extreme way. Within one hour the solver finds a solution with range 18 and lengths (2281, 2281, 2291, 2298, 2299), without proving optimality, in which every driver is worse off than in the cost-efficient plan and the total cost is 11 450, 48% above the minimum. For a planner, the lesson is that with a non-monotonic measure a moderate increase of the budget can make almost the whole fleet worse off than in the cost-efficient plan, or even produce a workload inconsistent plan, while the fairness measure improves considerably. A tight budget limits the redistribution to one or two drivers, whereas removing the budget can make every driver worse off at once.

This pattern is not specific to one instance. Table 9 reports, over the 20 instances of each size, the number of drivers who are worse off than in the cost-efficient plan and the relative change of their route lengths, for the optimal plans found by the strengthened configuration. Both the typical and the worst-case number of drivers worse off grow with the budget. At $B = 110\%$ the median is 2 to 3 drivers out of 5 and the worst case 4. The median relative change is zero at every budget, because many routes are kept unchanged. The mean is much higher, pulled up by drivers whose cost-efficient route served a single nearby customer, for whom a small absolute increase becomes a very large relative one.

8. Conclusions

We introduced the X-CVRP-TSP, a capacitated vehicle routing problem in which every route must be a shortest tour for its customers and the fairness of the plan is measured by an ordered function of the route lengths. The ordered framework covers the usual fairness measures in a single model, from the mean and the maximum to the range, the Gini mean difference, absolute




Figure 4: Cost-efficient plan and fair (RANGE) plans with $B = 101, 105, 110$ and 120% and without budget, instance 15 with $n = 20$. Customer coordinates are recovered from the distance matrix by classical multidimensional scaling. Legends give the route lengths from shortest (v_1) to longest (v_5), and panel titles give the total cost and the range. The plan without budget is the best solution found within one hour.

892 deviations and the variance. The Range-CVRP-TSP had been studied before,
893 and the other measures are addressed here for the first time.

894 A bilevel model, in which a planner assigns customers to vehicles and each
895 driver solves a TSP, guarantees that all routes are TSP-optimal. It rules out
896 the needless detours that non-monotonic measures would otherwise reward,
897 and it keeps the routing part independent of the measure, so that one method
898 serves all of them. Our experiments reveal that, without TSP-optimality
899 enforcement, every benchmark run with a non-monotonic measure returns a
900 route inconsistent plan whose fairness is overstated by 12% on average, and
901 post-processing the routes does not recover the optimum. This quantifies, on
902 the benchmark, the concern about non-TSP-optimality raised by Matl et al.
903 (2018).

904 Imposing TSP-optimality does not prevent workload inconsistency. On
905 our benchmark, with a budget 10% above the minimum cost, two or three

Table 9: Fair (RANGE) plans of the strengthened configuration against the cost-efficient plan ($B = 100\%$), 20 instances per size. Number of drivers worse off (out of 5, one value per instance) and relative change of the sorted route lengths (over all drivers and instances).

n	B (%)	Drivers worse off (of 5)				Relative change vs. $B = 100\%$ (%)			
		Min	Median	Mean	Max	Min	Median	Mean	Max
15	101	0	0.0	0.35	2	-29.9	0.0	4.1	304.0
	105	0	1.5	1.55	3	-29.5	0.0	49.5	2607.1
	110	1	2.0	1.95	4	-23.6	0.0	97.8	3785.7
20	101	0	0.0	0.65	2	-19.0	0.0	5.9	318.5
	105	1	2.0	1.95	4	-29.8	0.0	38.0	703.7
	110	1	3.0	2.45	4	-22.9	0.0	66.5	1144.0

of five drivers are worse off than in the cost-efficient plan, some optimal RANGE plans are dominated route by route by the cost-efficient plan itself, and without a budget every driver can be worse off. Tight budgets limit this effect.

The sparse reformulation and the coupling inequalities make the approach practical. The resulting method solves the Range-CVRP-TSP up to two orders of magnitude faster than the state of the art, handles instances with 35 customers, and solves instances with up to 30 customers for seven fairness measures.

Future research includes decomposition methods based on column generation, heuristics for larger instances, and fair routing with time windows, uncertain demands or several depots. In all these cases the bilevel structure allows the routing part to change without touching the fairness part. The exact detection of workload inconsistency, which requires comparing a plan with every cheaper feasible plan, and the separation of constraints that prevent it are the subject of future work.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Claude Fable 5.1 (Anthropic) in order to check grammar, spelling and references, to review the correctness of the manuscript and to assist in drafting proofs. The use of this tool for the code and the computational experiments is described in Section 6.1. After using this tool, the authors carefully reviewed, verified and edited the content and take full responsibility for it.

Acknowledgements

This work was supported by the Agencia Estatal de Investigación, Ministerio de Ciencia, Innovación y Universidades (MICIU/AEI/10.13039/501100011033) [grant numbers PID2024-156594NB-C21, PCI2024-155024-2 and CEX2024-001517-M].

References

- Adams, J.S., 1963. Towards an understanding of inequity. *The Journal of Abnormal and Social Psychology* 67, 422–436.
- Applegate, D., Cook, W., Dash, S., Rohe, A., 2002. Solution of a min-max vehicle routing problem. *INFORMS Journal on Computing* 14, 132–143.
- Archetti, C., Coelho, L.C., Speranza, M.G., Vansteenwegen, P., 2025. Beyond fifty years of vehicle routing: Insights into the history and the future. *European Journal of Operational Research* 330, 355–372.
- Artigues, C., Jozefowiez, N., Sarpong, B.M., 2018. Column generation algorithms for bi-objective combinatorial optimization problems with a min-max objective. *EURO Journal on Computational Optimization* 6, 117–142.
- Balinski, M.L., Quandt, R.E., 1964. On an integer program for a delivery problem. *Operations Research* 12, 300–304.
- Bektaş, T., Letchford, A.N., 2020. Using ℓ_p -norms for fairness in combinatorial optimisation. *Computers & Operations Research* 120, 104975.
- Bertazzi, L., Golden, B., Wang, X., 2015. Min-max vs. min-sum vehicle routing: A worst-case analysis. *European Journal of Operational Research* 240, 372–381.
- Bertsimas, D., Farias, V.F., Trichakis, N., 2011. The price of fairness. *Operations Research* 59, 17–31.
- Bertsimas, D., Farias, V.F., Trichakis, N., 2012. On the efficiency-fairness trade-off. *Management Science* 58, 2234–2250.
- Blanco, V., Gázquez, R., 2023. Fairness in maximal covering location problems. *Computers & Operations Research* 157, 106287.
- Blanco, V., Marín, A., Puerto, J., 2024. Intra-facility equity in discrete and continuous p-facility location problems. *Computers & Operations Research* 162, 106487.
- Blanco, V., Pozo, M.A., Puerto, J., Torrejón, A., 2026. Optimization-based computation and integration of ordered measures. *Annals of Operations Research* 364, 2799–2834.

- 966 Boland, N., Domínguez-Marín, P., Nickel, S., Puerto, J., 2006. Exact pro-
 967 cedures for solving the discrete ordered median problem. *Computers &*
 968 *Operations Research* 33, 3270–3300.
- 969 Campbell, A.M., Vandenbussche, D., Hermann, W., 2008. Routing for relief
 970 efforts. *Transportation Science* 42, 127–145.
- 971 Dantzig, G.B., Ramser, J.H., 1959. The truck dispatching problem. *Manage-*
 972 *ment Science* 6, 80–91.
- 973 Deleplanque, S., Labbé, M., Ponce, D., Puerto, J., 2020. A branch-price-and-
 974 cut procedure for the discrete ordered median problem. *INFORMS Journal*
 975 *on Computing* 32, 582–599.
- 976 Dempe, S., 2002. Foundations of Bilevel Programming. volume 61 of *Non-*
 977 *convex Optimization and Its Applications*. Kluwer Academic Publishers,
 978 Dordrecht.
- 979 Dempe, S., Zemkoho, A.B., 2013. The bilevel programming problem: reformu-
 980 lations, constraint qualifications and optimality conditions. *Mathematical*
 981 *Programming* 138, 447–473.
- 982 Edmonds, J., 1970. Submodular functions, matroids, and certain polyhedra,
 983 in: Guy, R., Hanani, H., Sauer, N., Schönheim, J. (Eds.), *Combinatorial*
 984 *Structures and Their Applications (Proceedings of the Calgary International*
 985 *Conference, 1969)*. Gordon and Breach, New York, pp. 69–87.
- 986 Errami, N., Queiroga, E., Sadykov, R., Uchoa, E., 2024. VRPSolverEasy:
 987 a Python library for the exact solution of a rich vehicle routing problem.
 988 *INFORMS Journal on Computing* 36, 956–965.
- 989 Fachini, R.F., Armentano, V.A., 2020. Logic-based Benders decomposition
 990 for the heterogeneous fixed fleet vehicle routing problem with time windows.
 991 *Computers & Industrial Engineering* 148, 106641.
- 992 Halvorsen-Weare, E.E., Savelsbergh, M.W.P., 2016. The bi-objective mixed
 993 capacitated general routing problem with different route balance criteria.
 994 *European Journal of Operational Research* 251, 451–465.
- 995 Harmati, I.Á., Coroianu, L., Fullér, R., 2024. The median under orness. *Fuzzy*
 996 *Sets and Systems* 481, 108901.
- 997 Held, M., Karp, R.M., 1962. A dynamic programming approach to sequencing
 998 problems. *Journal of the Society for Industrial and Applied Mathematics*
 999 10, 196–210.
- 1000 Herer, Y.T., 1999. Submodularity and the traveling salesman problem. *Euro-*
 1001 *pean Journal of Operational Research* 114, 489–508.
- 1002 Herer, Y.T., Penn, M., 1995. Characterizations of naturally submodular

1003 graphs: A polynomially solvable class of the TSP. *Proceedings of the*
1004 *American Mathematical Society* 123, 673–679.

1005 Hooker, J.N., Williams, H.P., 2012. Combining equity and utilitarianism in a
1006 mathematical programming model. *Management Science* 58, 1682–1693.

1007 Huang, M., Smilowitz, K., Balcik, B., 2011. Models for relief routing: Equity,
1008 efficiency and efficacy. *Procedia - Social and Behavioral Sciences* 17, 416–
1009 437.

1010 Jozefowicz, N., Semet, F., Talbi, E.G., 2002. Parallel and hybrid models for
1011 multi-objective optimization: Application to the vehicle routing problem,
1012 in: *Parallel Problem Solving from Nature – PPSN VII*, Springer, Berlin.
1013 pp. 271–280.

1014 Karsu, Ö., Morton, A., 2015. Inequity averse optimization in operational
1015 research. *European Journal of Operational Research* 245, 343–359.

1016 Kleinert, T., Labbé, M., Ljubić, I., Schmidt, M., 2021. A survey on mixed-
1017 integer programming techniques in bilevel optimization. *EURO Journal on*
1018 *Computational Optimization* 9, 100007.

1019 Labbé, M., Ponce, D., Puerto, J., 2017. A comparative study of formulations
1020 and solution methods for the discrete ordered p-median problem. *Computers*
1021 *& Operations Research* 78, 230–242.

1022 Laporte, G., 2009. Fifty years of vehicle routing. *Transportation Science* 43,
1023 408–416.

1024 Laporte, G., Louveaux, F.V., 1998. Solving stochastic routing problems
1025 with the integer L-shaped method, in: Crainic, T.G., Laporte, G. (Eds.),
1026 *Fleet Management and Logistics*. Kluwer Academic Publishers, Boston, pp.
1027 159–167.

1028 Ljubić, I., Marín, A., Puerto, J., Temprano, F., 2026. Ordered median
1029 traveling salesman problem. *Networks* 88, 59–84.

1030 Ljubić, I., Pozo, M.A., Puerto, J., Torrejón, A., 2024. Benders decomposition
1031 for the discrete ordered median problem. *European Journal of Operational*
1032 *Research* 317, 858–874.

1033 López Sánchez, A., Lujak, M., Semet, F., Billhardt, H., 2024. How to achieve
1034 fair and efficient cooperative vehicle routing? *AI Communications* 37,
1035 223–245.

1036 Lovász, L., 1983. Submodular functions and convexity, in: Bachem, A.,
1037 Grötschel, M., Korte, B. (Eds.), *Mathematical Programming: The State of*
1038 *the Art (Bonn 1982)*. Springer, Berlin, pp. 235–257.

1039 Lozano, J., González-Gurrola, L.C., Rodríguez-Tello, E., Lacomme, P., 2016.
1040 A statistical comparison of objective functions for the vehicle routing

1041 problem with route balancing, in: 2016 Fifteenth Mexican International
 1042 Conference on Artificial Intelligence (MICAI), IEEE. pp. 130–135.
 1043 Mandell, M.B., 1991. Modelling effectiveness-equity trade-offs in public service
 1044 delivery systems. *Management Science* 37, 467–482.
 1045 Marín, A., Ponce, D., Puerto, J., 2020. A fresh view on the discrete ordered
 1046 median problem based on partial monotonicity. *European Journal of*
 1047 *Operational Research* 286, 839–848.
 1048 Marsh, M.T., Schilling, D.A., 1994. Equity measurement in facility location
 1049 analysis: A review and framework. *European Journal of Operational*
 1050 *Research* 74, 1–17.
 1051 Matl, P., Hartl, R.F., Vidal, T., 2018. Workload equity in vehicle routing
 1052 problems: A survey and analysis. *Transportation Science* 52, 239–260.
 1053 Mesa, J.A., Puerto, J., Tamir, A., 2003. Improved algorithms for several
 1054 network location problems with equality measures. *Discrete Applied Math-*
 1055 *ematics* 130, 437–448.
 1056 Miller, C.E., Tucker, A.W., Zemlin, R.A., 1960. Integer programming formu-
 1057 lation of traveling salesman problems. *Journal of the ACM* 7, 326–329.
 1058 Nickel, S., Puerto, J., 2006. *Location Theory: A Unified Approach*. Springer,
 1059 Berlin.
 1060 Pessoa, A., Sadykov, R., Uchoa, E., Vanderbeck, F., 2020. A generic exact
 1061 solver for vehicle routing and related problems. *Mathematical Programming*
 1062 183, 483–523.
 1063 Pozo, M.A., Puerto, J., Torrejón, A., 2024. The ordered median tree location
 1064 problem. *Computers & Operations Research* 169, 106746.
 1065 Puerto, J., Rodríguez-Chía, A.M., Tamir, A., 2017. Revisiting k-sum opti-
 1066 mization. *Mathematical Programming* 165, 579–604.
 1067 Rawls, J., 1971. *A Theory of Justice*. Harvard University Press, Cambridge,
 1068 MA.
 1069 Reiter, P., Gutjahr, W.J., 2012. Exact hybrid algorithms for solving a bi-
 1070 objective vehicle routing problem. *Central European Journal of Operations*
 1071 *Research* 20, 19–43.
 1072 van Rossum, B., Chen, R., Lodi, A., 2026a. Efficient branching rules for
 1073 optimizing range and order-based objective functions. *Mathematical Pro-*
 1074 *gramming* 216, 539–572.
 1075 van Rossum, B., Chen, R., Lodi, A., 2026b. Enforcing TSP-optimality in fair
 1076 vehicle routing by cutting planes. *Operations Research Letters* 70, 107539.
 1077 Savas, E.S., 1978. On equity in providing public services. *Management Science*
 1078 24, 800–808.

1079 Svestka, J.A., Huckfeldt, V.E., 1973. Computational experience with an m-
1080 salesman traveling salesman algorithm. *Management Science* 19, 790–799.
1081 Toth, P., Vigo, D., 2014. *Vehicle Routing: Problems, Methods, and Applica-*
1082 *tions*. MOS-SIAM Series on Optimization. second ed., SIAM, Philadelphia.
1083 Tsang, M.Y., Shehadeh, K.S., 2025. A unified framework for analyzing and
1084 optimizing a class of convex fairness measures. *Operations Research* 74,
1085 1087–1103.
1086 Uchoa, E., Pecin, D., Pessoa, A., Poggi, M., Vidal, T., Subramanian, A.,
1087 2017. New benchmark instances for the capacitated vehicle routing problem.
1088 *European Journal of Operational Research* 257, 845–858.

1089 Appendix A. Proofs

1090 Throughout, costs are symmetric and satisfy the triangle inequality, $S =$
1091 $\{i \in V : x_i^k = 1\}$ denotes the set of nodes assigned to vehicle k , which contains
1092 the depot and at least one customer, and $\text{TSP}(S) = \phi_k(x^k)$. For a route T
1093 we write $W = V_T \cap S$ for the nodes of T that are still assigned, $D = V_T \setminus S$
1094 for the dropped ones and $U = S \setminus V_T$ for the assigned nodes outside T .

1095 Appendix A.1. Proof of Theorem 1

1096 ($\Leftarrow$) Assume $\lambda_\ell \geq 0$ for all $\ell \in K$, and let $\Phi' = \Phi + \delta$ with $\delta \geq 0$.
1097 The order statistics are monotone, that is, $\Phi_{(\ell)} \leq \Phi'_{(\ell)}$ for every ℓ . Indeed,
1098 $\Phi_{(\ell)} = \min_{I \subseteq K, |I|=\ell} \max_{i \in I} \Phi_i$, and $\max_{i \in I} \Phi_i \leq \max_{i \in I} \Phi'_i$ for every set I .
1099 Hence

$$l_{\text{ord}}(\Phi', \boldsymbol{\lambda}) - l_{\text{ord}}(\Phi, \boldsymbol{\lambda}) = \sum_{\ell \in K} \lambda_\ell (\Phi'_{(\ell)} - \Phi_{(\ell)}) \geq 0,$$

1100 and the measure is monotone.

1101 ($\Rightarrow$) We prove the contrapositive. Let $\lambda_j < 0$ for some $j \in K$, choose
1102 $\Phi \in \mathbb{R}_{\geq 0}^m$ with pairwise distinct entries, and let i_j be the index of the entry in
1103 sorted position j , so that $\Phi_{i_j} = \Phi_{(j)}$. Since the entries are distinct, there is
1104 $\delta > 0$ such that $\Phi_{(j)} + \delta$ is still smaller than $\Phi_{(j+1)}$ (no condition is needed if
1105 $j = m$). Define Φ' by $\Phi'_{i_j} = \Phi_{i_j} + \delta$ and $\Phi'_i = \Phi_i$ for $i \neq i_j$. Then $\Phi' \geq \Phi$, the
1106 sorted order does not change, and only the j -th order statistic increases, by δ .
1107 Therefore

$$l_{\text{ord}}(\Phi', \boldsymbol{\lambda}) - l_{\text{ord}}(\Phi, \boldsymbol{\lambda}) = \lambda_j \delta < 0,$$

1108 so l_{ord} is not monotone.

1109 *Appendix A.2. Justification of Remark 1*

1110 Consider the absolute deviation from the q -th smallest route length,

$$\sum_{k \in K} |\Phi_k - \Phi_{(q)}|, \quad (\text{A.1})$$

1111 which is the loss (2) with $\boldsymbol{\lambda}$ equal to the q -th unit vector. It follows the
 1112 logic of the mean absolute deviation, but replaces the mean by a more robust
 1113 central value, the q -th order statistic (Harmati et al., 2024; Blanco et al.,
 1114 2026). Since $\Phi_{(k)} \leq \Phi_{(q)}$ for $k < q$ and $\Phi_{(k)} \geq \Phi_{(q)}$ for $k > q$, measure (A.1)
 1115 is the linear ordered measure with weight vector

$$(-1, \dots, -1, \underbrace{2q - m - 1}_{q\text{-th}}, 1, \dots, 1). \quad (\text{A.2})$$

1116 Its first entry is -1 if $q \geq 2$ and $1 - m$ if $q = 1$, so for $m \geq 2$ the vector
 1117 always has a negative entry, and the measure is not monotone by Theorem 1.

1118 *Appendix A.3. Justification of Remark 2*

Consider the squared deviation (3) from an ordered reference value $l_{\text{ord}}(\Phi, \boldsymbol{\lambda})$. As shown by Blanco et al. (2026), it can be written as a quadratic ordered measure

$$\mathcal{Q}(\Phi, \mathbf{A}) = \sum_{i, j \in K} a_{ij} \Phi_{(i)} \Phi_{(j)},$$

1119 where $\mathbf{A} = (a_{ij})$ is a symmetric matrix determined by $\boldsymbol{\lambda}$. For the variance,
 1120 that is, the squared deviation from the mean,

$$\sum_{i \in K} \left(\Phi_i - \frac{1}{m} \sum_{j \in K} \Phi_j \right)^2, \quad (\text{A.3})$$

1121 the matrix is

$$\mathbf{A} = \begin{pmatrix} 1 - \frac{1}{m} & -\frac{1}{m} & \cdots & -\frac{1}{m} \\ -\frac{1}{m} & 1 - \frac{1}{m} & \ddots & \vdots \\ \vdots & \ddots & \ddots & -\frac{1}{m} \\ -\frac{1}{m} & \cdots & -\frac{1}{m} & 1 - \frac{1}{m} \end{pmatrix},$$

1122 that is, the centering matrix, for which $\mathcal{Q}(\Phi, \mathbf{A})$ is the sum of squared devia-
 1123 tions from the mean. This form is not monotone, although $\boldsymbol{\lambda} = \frac{1}{m}(1, \dots, 1) \geq$
 1124 0 . For instance, $\Phi' = (1, 3, 5)$ and $\Phi'' = (2.5, 3, 5)$ satisfy $\Phi' \leq \Phi''$ component-
 1125 wise, but $\mathcal{Q}(\Phi', \mathbf{A}) = 8$ and $\mathcal{Q}(\Phi'', \mathbf{A}) = 3.5$. Lengthening the shortest route
 1126 thus reduces the loss.

1127 *Appendix A.4. Proof of Proposition 1*

1128 (a) Constraints (7f) together with the arc-cost identity (7g) force z^k to be
 1129 a feasible tour over $\{i \in V : x_i^k = 1\}$ of length exactly Φ_k . Since $\phi_k(x^k)$ is by
 1130 definition the length of a shortest such tour, $\Phi_k \geq \phi_k(x^k)$ holds automatically
 1131 for every feasible z^k , independently of (7h). Imposing (7h), $\Phi_k \leq \phi_k(x^k)$,
 1132 for every $k \in K$ therefore pins $\Phi_k = \phi_k(x^k)$ exactly. The route lengths of
 1133 every feasible assignment are then those of the optimal tours of the followers,
 1134 as in the bilevel problem (5), and since the objective, the budget and the
 1135 ordering constraints are the same in both models, they have the same feasible
 1136 assignments and the same optimal solutions.

1137 (b) Dropping (7h) at a single position k only removes a constraint from the
 1138 exact reformulation of part (a), so the resulting model is always a relaxation.
 1139 If $k < k_0$, the relaxation is not strict, because every feasible point of the
 1140 relaxed model can be turned into a feasible point of the exact one with an
 1141 objective value that is not larger. Indeed, replacing the inflated value Φ_k by
 1142 the true length $\phi_k(x^k)$ keeps the budget satisfied and, after re-sorting, moves
 1143 routes only into positions $j \leq k < k_0$, whose weights are non-negative, so
 1144 the objective cannot increase. If $k \geq k_0$ and $\lambda_k < 0$, the relaxation is strict
 1145 because the vehicle in position k may lengthen its route, up to Φ_{k+1} or the
 1146 budget, and every unit of extra length is rewarded by $|\lambda_k|$. The remaining
 1147 case, $k > k_0$ with $\lambda_k \geq 0$, is the non-obvious one, since an inflated Φ_k is never
 1148 directly rewarded by the objective. Let $i < k$ be a position with $\lambda_i < 0$, for
 1149 instance $i = k_0$. The reason it still matters is that the ordering constraints
 1150 (7e) couple the positions. A vehicle with a short optimal tour may drive a
 1151 longer tour and move up to position k , which pushes longer routes down into
 1152 position i , where the negative weight λ_i rewards them.

1153 *Appendix A.5. Proof of Proposition 2*

1154 Let $D = \{i \in V_T : x_i^k = 0\}$ be the set of dropped nodes at an integer point.
 1155 Then $\sum_{i \in V_T} x_i^k = |V_T| - |D|$, and the right-hand sides of the two cuts are

$$\underline{\Phi}_k + \underline{g}_k(1 - |D|) = \text{TSP}(V_T) - |D| \underline{g}_k \quad \text{and} \quad \text{TSP}(V_T) - \sum_{i \in D} (c_i^{[1]} + c_i^{[2]}),$$

1156 with $\underline{g}_k = \text{TSP}(V_T) - \underline{\Phi}_k$. Hence (10) is the stronger cut exactly when
 1157 $\sum_{i \in D} (c_i^{[1]} + c_i^{[2]}) < |D| \underline{g}_k$, and (9) is the stronger one when the reverse
 1158 inequality holds. Both cases occur. Consider an instance in which all costs
 1159 between distinct nodes are equal to 1, which is a metric. Then $c_i^{[1]} + c_i^{[2]} = 2$

for every node and every route has length at least 2, so we can take $\Phi_k = 2$. If V_T contains a single customer and x^k does not include it, then $\text{TSP}(V_T) = 2$, $\underline{g}_k = 0$, and the right-hand sides are 2 and 0, so (9) is stronger. If V_T contains $t \geq 4$ customers and $D \neq \emptyset$, then $\text{TSP}(V_T) = t+1$ and $\underline{g}_k = t-1 > 2 = c_i^{[1]} + c_i^{[2]}$ for every $i \in D$, so (10) is stronger.

Appendix A.6. Proof of Proposition 3

Let $g : \{0, 1\}^2 \rightarrow \mathbb{R}$ be the length of the shortest tour through the depot and the customers present among i, j , so that $g(0, 0) = 0$, $g(1, 0) = 2c_{0i}$, $g(0, 1) = 2c_{0j}$ and $g(1, 1) = T_{ij}$.

Validity. Since $\ell_{ij}(x^k)$ depends only on x_i^k and x_j^k , we write $\ell_{ij}(x_i^k, x_j^k)$ for its value. As the TSP cost cannot decrease when customers are added, it suffices to check $\ell_{ij}(x_i^k, x_j^k) \leq g(x_i^k, x_j^k)$ at the four points of $\{0, 1\}^2$, taking S equal to the customers present among $\{i, j\}$. First, $\ell_{ij}(0, 0) = 0 = g(0, 0)$. Second, $\ell_{ij}(1, 0) = 2c_{0i} = g(1, 0)$. Any tour through i contains two paths between the depot and i , each of length at least c_{0i} by the triangle inequality. Third, $\ell_{ij}(1, 1) = T_{ij} = g(1, 1)$. Finally, $\ell_{ij}(0, 1) = c_{ij} + c_{j0} - c_{0i} \leq (c_{i0} + c_{0j}) + c_{j0} - c_{0i} = 2c_{0j} = g(0, 1)$, again by the triangle inequality.

Envelope. Let $h^*(x_i^k, x_j^k) = \max\{\ell_{ij}(x_i^k, x_j^k), \ell_{ji}(x_j^k, x_i^k)\}$. A direct computation gives

$$\ell_{ij}(x_i^k, x_j^k) - \ell_{ji}(x_j^k, x_i^k) = (x_i^k - x_j^k) (2c_{0i} + 2c_{0j} - T_{ij}),$$

and $2c_{0i} + 2c_{0j} \geq T_{ij}$ because $c_{ij} \leq c_{i0} + c_{0j}$. Hence $h^*(x_i^k, x_j^k) = \ell_{ij}(x_i^k, x_j^k)$ on the triangle $\{x_i^k \geq x_j^k\}$, with vertices $(0, 0)$, $(1, 0)$ and $(1, 1)$, and $h^*(x_i^k, x_j^k) = \ell_{ji}(x_j^k, x_i^k)$ on the triangle $\{x_j^k \geq x_i^k\}$, with vertices $(0, 0)$, $(0, 1)$ and $(1, 1)$. In particular h^* is convex and equals g at the four vertices of the square. So h^* is the lower convex envelope of g on $[0, 1]^2$.

Equivalently, the inequality $g(0, 0) + g(1, 1) \leq g(1, 0) + g(0, 1)$ says that g is submodular on the two-element ground set $\{i, j\}$, and h^* is its Lovász extension, which for a submodular function coincides with the lower convex envelope on the unit cube (Lovász, 1983). The inequalities $\Phi_k \geq \ell_{ij}(x_i^k, x_j^k)$ and $\Phi_k \geq \ell_{ji}(x_j^k, x_i^k)$ are the two extended polymatroid inequalities of g in the sense of Edmonds (1970).

Finally, if h is affine in (x_i^k, x_j^k) and $\Phi_k \geq h$ is valid, then $h \leq g$ at the four vertices, which are incidence vectors of subsets, so $h \leq h^*$ on the square, and h^* is the strongest such bound. In particular, the radius inequality $\Phi_k \geq 2c_{0i}x_i^k$, the pair inequality $\Phi_k \geq T_{ij}(x_i^k + x_j^k - 1)$ and the cut (10) for

1194 $V_T = \{0, i, j\}$ are affine in (x_i^k, x_j^k) and valid, so each of them is dominated
 1195 at every point of the square by (pair-env) or by its counterpart with i and j
 1196 exchanged.

1197 *Appendix A.7. Proof of Proposition 4*

1198 Fix t and a feasible solution, and let R be the set of customers served by
 1199 the t longest routes, those in positions $m - t + 1, \dots, m$. These routes respect
 1200 the capacity, so $|R| \leq P^*(t) \leq p_t$. Let U_t be a set of $p_t + 1$ customers with
 1201 the largest depot costs. Since $|U_t| > |R|$, some customer $u \in U_t$ is served by
 1202 one of the other $m - t$ routes, those in positions $1, \dots, m - t$. The route that
 1203 serves u contains two paths between the depot and u , so its length is at least
 1204 $2c_{0u} \geq 2D_{(p_t+1)}$, and Φ_{m-t} , the longest of these $m - t$ routes, is at least as
 1205 long. If $p_t \geq |V'|$, the right-hand side is 0 by convention and there is nothing
 1206 to prove.

1207 *Appendix A.8. Proof of Proposition 5*

1208 Fix a feasible assignment x^k , let $S = \{i \in V : x_i^k = 1\}$ be the set of nodes
 1209 assigned to vehicle k , and let $U = S \setminus V_T$ be the assigned customers outside
 1210 V_T . For (bilevel-Ins), start from an optimal tour over V_T , which contains the
 1211 depot and at least one customer, and insert the customers of U one at a
 1212 time between two consecutive nodes $r \neq t$ of the current tour. Each insertion
 1213 of a customer u increases the length by $c_{ru} + c_{ut} - c_{rt} \leq M_u^1$ (Figure A.1),
 1214 whatever the tour built so far. This gives a tour over $V_T \cup U$ of length at most
 1215 $\text{TSP}(V_T) + \sum_{u \in U} M_u^1$, and $\phi_k(x^k) = \text{TSP}(S) \leq \text{TSP}(V_T \cup U)$ since $S \subseteq V_T \cup U$.

1216 For (bilevel-Cap), split U into $U_{\text{small}} = \{u \in U : M_u^1 \leq g\}$ and $U_{\text{large}} =$
 1217 $\{u \in U : M_u^1 > g\}$. If $U_{\text{large}} = \emptyset$, then $\min\{g, M_u^1\} = M_u^1$ for every $u \in U$,
 1218 and the right-hand side of (bilevel-Cap) is that of (bilevel-Ins), which is valid.
 1219 If $U_{\text{large}} \neq \emptyset$, the sum is at least g , because one term equals g and all terms
 1220 are nonnegative, since $g \geq 0$ for every separated route (Section 5.2), so

$$\text{TSP}(V_T) + \sum_{u \in U} \min\{g, M_u^1\} \geq \text{TSP}(V_T) + g = \bar{\Phi} \geq \phi_k(x^k).$$

1221 The dominance claim follows directly from the definition of the coefficients as
 1222 pairwise minima.

Appendix A.9. Proof of Proposition 6

The insertion penalty M_i^1 makes skipping an assigned customer i at least as expensive as visiting it (Figure A.1). The removal penalty $M_s^{2,T}$ makes visiting a customer s that is not assigned at least as expensive as skipping it (Figure A.2). The following lemma uses both and gives the proposition. Throughout, fix a feasible assignment x^k and let $S = \{i \in V : x_i^k = 1\}$ be the set of nodes assigned to vehicle k , $W = S \cap V_T$ the nodes of T that are assigned, which always contains the depot, $D = V_T \setminus W$ the dropped nodes of T and $U = S \setminus V_T$ the assigned nodes outside T . We write $R_L(T, x^k)$ for the right-hand side of (bilevel-L).




Figure A.1: Insertion penalty. Skipping the assigned customer i between two assigned nodes r and t (dashed arc) costs $c_{rt} + M_i^1$, which is at least the cost $c_{ri} + c_{it}$ of visiting it.




Figure A.2: Removal penalty. Visiting the customer s , which is not assigned, between the assigned node u and the node v (dashed arcs) counts c_{us} plus $M_s^{2,T}$, while the arc leaving s costs 0, since only arcs leaving assigned nodes are counted. This is at least the cost c_{uv} of going directly from u to v .

Lemma 1. For every route $T \in \mathcal{T}$, every $k \in K$ and every feasible assignment x^k , $R_L(T, x^k) \geq \phi_k(x^k)$.

Proof. Write $R_L(T, x^k) = \alpha + \beta + \gamma$, where $\alpha = \sum_{(i,j) \in A_T} x_i^k c_{ij}$ is the arc term, $\beta = \sum_{i \in D} M_i^{2,T}$ collects the removal penalties of the dropped nodes and $\gamma = \sum_{u \in U} M_u^1$ the insertion penalties of the assigned nodes outside T . The arc term counts exactly the arcs of T whose tail is in W , and all penalties

are nonnegative. We show that a tour over S of length at most $\alpha + \beta + \gamma$ exists. We distinguish two cases.

Case 1: W contains a customer. Visiting the nodes of W in the order in which T visits them gives a tour T_W over W , with at least two distinct nodes. Each arc (r, q) of T_W replaces a path $r \rightarrow d_1 \rightarrow \dots \rightarrow d_h \rightarrow q$ of T , where $h \geq 0$, $r \neq q$ belong to W and $d_1, \dots, d_h \in D$. If $h = 0$, the arc (r, q) belongs to T , and its cost is counted in α . If $h \geq 1$, the arc term counts c_{rd_1} , because $r \in W$, and none of the other arcs of the path, whose tails are dropped nodes. Since $r \neq q$ belong to V_T and $d_1 \notin \{r, q\}$, the pair (r, q) is admissible in the definition of $M_{d_1}^{2,T}$, so $M_{d_1}^{2,T} \geq c_{rq} - c_{rd_1} + c_{d_1q}$ (Figure A.2) and

$$c_{rd_1} + M_{d_1}^{2,T} \geq c_{rq} + c_{d_1q} \geq c_{rq}.$$

Every dropped node lies on exactly one such path, and the penalties of $d_2, \dots, d_h$ are nonnegative, so summing over the arcs of T_W gives $\alpha + \beta \geq c(T_W)$, the length of T_W . Next, insert the customers of U into T_W one at a time, as in Appendix A.8, each between two consecutive and distinct nodes and at a cost of at most M_u^1 . The result is a tour over $W \cup U = S$ of length at most $c(T_W) + \gamma \leq R_L(T, x^k)$.

Case 2: $W = \{0\}$. Then T is $0 \rightarrow d_1 \rightarrow \dots \rightarrow d_h \rightarrow 0$ with $d_1, \dots, d_h \in D$, the arc term is $\alpha = c_{0d_1}$, and $U \neq \emptyset$ because vehicle k serves at least one customer. If $h \geq 2$, the pair $(d_2, 0)$ is admissible for $M_{d_1}^{2,T}$ and the pair $(0, d_1)$ for $M_{d_2}^{2,T}$, so

$$\begin{aligned} \beta &\geq M_{d_1}^{2,T} + M_{d_2}^{2,T} \geq (c_{d_20} - c_{d_2d_1} + c_{d_10}) + (c_{0d_1} - c_{0d_2} + c_{d_2d_1}) \\ &= 2c_{0d_1}, \end{aligned}$$

and $R_L(T, x^k) \geq 2c_{0d_1} + \gamma$. If $h = 1$, then $V_T = \{0, d_1\}$ and $M_{d_1}^{2,T} = c_{0d_1}$ by definition, so $\beta = c_{0d_1}$ and again $R_L(T, x^k) \geq 2c_{0d_1} + \gamma$. In both cases, insert the customers of U into the closed walk $0 \rightarrow d_1 \rightarrow 0$, of length $2c_{0d_1}$, each at a cost of at most M_u^1 , and then skip d_1 , which does not increase the length by the triangle inequality. The result is a tour over $S = \{0\} \cup U$ of length at most $R_L(T, x^k)$.

In all cases $\phi_k(x^k) = \text{TSP}(S) \leq R_L(T, x^k)$. $\square$

Proof of Proposition 6. Validity of (bilevel-L) is Lemma 1, and the tightness is immediate.

1268 *Appendix A.10. Proof of Proposition 7*

1269 With the notation of Appendix A.9, we first compare the two cuts in
 1270 general, for a route T that is an optimal tour over V_T , as in separation, so
 1271 that $C_T^* = \text{TSP}(V_T)$. If $D = \emptyset$, every node of T is assigned, so the arc term of
 1272 (bilevel-L) equals C_T^* and its removal penalties vanish. The cut becomes $\Phi_k \leq$
 1273 $C_T^* + \sum_{u \in U} M_u^1$, whereas (bilevel-Cap) reads $\Phi_k \leq C_T^* + \sum_{u \in U} \min\{g, M_u^1\}$,
 1274 so the capped insertion cut is at least as strong, and strictly stronger when
 1275 $M_u^1 > g$ for some $u \in U$. If $D \neq \emptyset$, the arc term of (bilevel-L) loses the arcs of
 1276 T that leave the dropped nodes, and only their penalties $M_i^{2,T}$ are added back,
 1277 whereas (bilevel-Cap) keeps the full length C_T^* . Depending on the instance,
 1278 the saving can be larger or smaller than the penalties, so (bilevel-L) can be
 1279 stronger.

1280 *Appendix A.11. Route-restricted insertion penalties are not valid*

1281 Section 5.4.2 restricts the removal penalty to the pairs of the route T
 1282 (Proposition 6). The same restriction is not valid for the insertion penalty, as
 1283 the following example shows.

1284 **Example 1.** Take Euclidean costs with the depot $0 = (0, 0)$ and customers
 1285 $v = (12, 0)$, $p = (6, 0)$, $a = (6, 8)$ and $b = (6, -8)$, so that $c_{0v} = 12$,
 1286 $c_{0p} = c_{pv} = 6$, $c_{0a} = c_{0b} = c_{av} = c_{bv} = 10$, $c_{pa} = c_{pb} = 8$ and $c_{ab} = 16$. Let
 1287 $T = 0 \rightarrow v \rightarrow 0$, of length 24. If M_i^1 is computed with $r, t \in V_T = \{0, v\}$, the
 1288 only pair is $(0, v)$, which gives $c_{0p} + c_{pv} - c_{0v} = 0$ for p and $10 + 10 - 12 = 8$
 1289 for a and for b . Consider the assignment $S = \{0, v, p, a, b\}$, in which vehicle k
 1290 serves these four customers and other vehicles serve further customers of the
 1291 instance. Then $D = \emptyset$, and the cut would give

$$\Phi_k \leq 24 + 0 + 8 + 8 = 40.$$

1292 However, $\text{TSP}(S) = 44$, attained for instance by $0 \rightarrow p \rightarrow a \rightarrow v \rightarrow b \rightarrow 0$,
 1293 so the cut would remove the true route length.

1294 **Appendix B. Algorithmic details**

1295 *Appendix B.1. Bounds on the route lengths*

1296 Lower bound. Let obj^* be the largest number of customers that $m - 1$
 1297 vehicles can serve, obtained from the assignment model

$$\max \sum_{i \in V'} \alpha_i$$

$$\begin{aligned}
\text{s.t. } & \sum_{j=1}^{m-1} \beta_{ij} = \alpha_i, & i \in V', \\
& \sum_{i \in V'} q_i \beta_{ij} \leq Q, & j = 1, \dots, m-1, \\
& \alpha, \beta \text{ binary.}
\end{aligned}$$

1298 Every route serves at least $a = n - \text{obj}^*$ customers, because the other $m - 1$
 1299 vehicles serve at most obj^* . If $a \geq 2$, every route is therefore at least as long as
 1300 a shortest tour through the depot and a customers with total demand at most
 1301 Q , which we compute with a cardinality-constrained TSP. Both problems are
 1302 solved under a time limit and only their dual bounds are used, which keeps
 1303 the bound valid when they are not solved to optimality. We set Φ_k to the
 1304 larger of $\min_{i \in V'} c_{0i}$ and the dual bound of the cardinality-constrained TSP.

1305 Upper bound. Since every vehicle serves at least one customer, no vehicle
 1306 serves all of V' , and each route is contained in $V' \setminus \{i\}$ for some i . Hence
 1307 $\max_{i \in V'} \text{TSP}(V' \setminus \{i\})$ is an upper bound on every route length, and we set
 1308 $\bar{\Phi}$ to the smaller of this value and the length of a tour through all customers.
 1309 This requires n TSP solves, done once per instance. For $n \geq 30$ these TSPs
 1310 are not always solved to optimality within the time limit, so we replace each
 1311 $\text{TSP}(V' \setminus \{i\})$ by the length of the best tour through the same customers
 1312 found by a multi-start 2-opt heuristic, which keeps the bound valid.

1313 *Appendix B.2. TSP subproblems and separation*

1314 To separate the cuts for vehicle k we need an optimal tour over the current
 1315 set S . For sets of at most 15 nodes, the depot included, we use the Held–Karp
 1316 dynamic program (Held and Karp, 1962), which runs in $O(2^{|S|}|S|^2)$ time and
 1317 takes at most a few milliseconds for the routes of our benchmark. For larger
 1318 sets we solve the TSP with Gurobi using the formulation of Miller et al. (1960).
 1319 The same set S appears many times during the search, assigned to different
 1320 vehicles or violating the bound in the other direction, so we store solved sets
 1321 and reuse their optimal tours.

1322 Algorithm B.1, our default separation algorithm, summarizes the separa-
 1323 tion at an integer solution $(\hat{x}, \hat{z}, \hat{\Phi})$. In the dense formulation, each cycle C of
 1324 z^k that does not contain the depot gives the cutset inequality (5m) with $S = C$
 1325 and, optionally, the aggregated inequality $\sum_{k' \in K} \sum_{i \in C, j \notin C} (z_{ij}^{k'} + z_{ji}^{k'}) \geq 2$. Op-
 1326 tionally, the aggregated inequality is also separated at the fractional points
 1327 of the root node by a max-flow procedure, which is the max-flow variant

1328 compared in Section 6.2. Since $\hat{\Phi}_k$ cannot be both smaller and larger than
 1329 C_T^* , at most one direction is separated for each vehicle.

Algorithm B.1 Separation at integer solutions.

Require: Integer solution $(\hat{x}, \hat{z}, \hat{\Phi})$, tolerance $\varepsilon = 10^{-4}$, flag **dense**, flag **ub** (true if and only if the measure is non-monotonic), cache of solved sets

- 1: $P \leftarrow \emptyset$; $T_k \leftarrow \emptyset$ for all $k \in K$ { P : sets whose TSP has to be solved; $T_k = \emptyset$: no bounding cut for k in this round}
- 2: **for all** $k \in K$ **do**
- 3: **if dense and** $\hat{z}^k$ contains cycles $C_1, \dots, C_r$ that avoid the depot **then**
- 4: separate constraints (5m)
- 5: **continue** { $\hat{\Phi}_k$ is not compared with a tour length: vehicle k gets connectivity cuts only; the other vehicles are still processed}
- 6: **end if**
- 7: $S_k \leftarrow \{i \in V : \hat{x}_i^k = 1\}$ {contains the depot}
- 8: **if dense and** $|S_k| \leq 3$ **then**
- 9: **continue** {the tour is fixed by (5k)–(5m) and $\Phi_k = \sum c_{ij} z_{ij}^k$ is exact}
- 10: **end if**
- 11: **if** S_k is in the cache **then**
- 12: $(T_k, C_k^*) \leftarrow \text{cache}[S_k]$
- 13: **else**
- 14: $P \leftarrow P \cup \{k\}$
- 15: **end if**
- 16: **end for**
- 17: solve the TSP over S_k for all $k \in P$. Store (T_k, C_k^*) in the cache
- 18: **for all** $k \in K$ with $T_k \neq \emptyset$ **do**
- 19: **if not dense and** $\hat{\Phi}_k < C_k^* - \varepsilon$ **then**
- 20: add (9) and (10) for (T_k, k)
- 21: **end if**
- 22: **if ub and** $\hat{\Phi}_k > C_k^* + \varepsilon$ **then**
- 23: add the upper-bounding cut (bilevel-Cap) for (T_k, k)
- 24: **end if**
- 25: **end for**
