


MODELS OF FUEL REDISTRIBUTION IN ARMoured-MECHANISED UNITS DURING THE TRANSITION FROM LOGISTICAL PAUSE TO COMBAT MANOEUVRE

Gordana Simić 

University of Niš,
Niš, Republic of Serbia
E-mail: gordana.simic@masfak.ni.ac.rs

Received: 08.03.2025. Approved: 27.05.2025.

Original Scientific Article

DOI: <https://doi.org/10.65932/military-studies-2025-1-4>

UDC: 355.4:355.41:005.932

Abstract: The transition from a logistical pause to a combat manoeuvre represents one of the most underexamined inflection points in armoured-mechanised operations. During a pause, vehicles are dispersed in an assembly area with heterogeneous fuel states shaped by patrol activity, engine idling, environmental conditioning and unequal previous consumption; when the order-to-move arrives, the battalion must refuel and move out within a fixed window that is typically 90 to 180 minutes. This article develops an original formal model of fuel redistribution for a reinforced armoured-mechanised battalion of ninety vehicles across three consumption classes, and compares four redistribution policies (proportional, priority-based, greedy shortest-processing-time, and an optimised lexicographic weighted schedule) under three tanker-fleet configurations (three, five, and eight 15,000-litre tankers). The study introduces the Battalion Refueling Readiness Index (BRRI), a novel composite metric that integrates the share of vehicles meeting their class-specific fuel threshold, the dispersion of the post-refuelling fuel distribution, and the makespan of the refuelling operation. Two hundred Monte Carlo realisations of pre-manoeuve fuel states yield three principal findings. First, policy choice is decisive in the knife-edge fleet regime: in the small-fleet configuration the priority-based policy outperforms the proportional baseline by 8.7 percentage points of BRRI, whereas in the large-fleet configuration all four policies converge within 0.5 points. Second, the greedy shortest-processing-time rule maximises the headcount readiness rate in constrained fleets (59.4% vs. 50.0% under priority), but does so at the price of elevated dispersion (22.3% vs. 17.3%), which is operationally costly because it leaves heavy vehicles of the main effort below the combat threshold. Third, sensitivity analysis shows that reducing the order-to-move window from 120 to 60 minutes collapses BRRI by roughly half in the small-fleet configuration while leaving the large-fleet configuration essentially unaffected, which identifies tanker-to-vehicle ratio rather than scheduling sophistication as the dominant lever. The original contribution of this article is threefold: the BRRI metric, a transition-specific mixed-integer scheduling formulation distinct from steady-state convoy refuelling problems, and an empirical demonstration that class-priority heuristics dominate in the operationally relevant fleet regime. The results have direct implications for the doctrinal sizing of Class III (Bulk) distribution platoons in small NATO and partner-state armoured brigades.

Keywords: *military logistics, armoured-mechanised battalion, fuel redistribution, refuelling operations, scheduling optimisation, combat readiness, BRRI.*

INTRODUCTION

Armoured-mechanised operations are constrained less by fires, manoeuvre space or tactical competence than by fuel. A single modern main battle tank consumes between 300 and 500 litres of diesel per hundred kilometres of cross-country movement; a reinforced battalion executing a two-day offensive burns through between 280,000 and 420,000 litres (US Department of the Army, 2020). Every operational study of large-scale armoured manoeuvre published since the fall of France in 1940 arrives at the same conclusion: it is the Class III (Bulk) supply chain, not the combat chain, that determines how far a mechanised force can go and how fast it can get there (Van Creveld, 1977). What these studies have consistently under-examined, however, is a specific inflection point in that chain: the moment at which a battalion transitions from a logistical pause in an assembly area to a combat manoeuvre.

During a logistical pause, fuel states become heterogeneous. Vehicles that have conducted perimeter patrol, driver training, engine warm-up cycles or long-duration idling draw down their tanks at different rates from those that have been shut down under tactical cover; crews that have operated heaters or environmental conditioning for several hours may be fifteen to twenty per cent lower than the battalion mean (Foss, 2023). When the order-to-move arrives, the fleet is effectively a collection of vehicles at different points on the fuel axis, all of which must be lifted, under a single fixed time window, to a combat-ready threshold that is seldom below ninety per cent of tank capacity. The redistribution problem is mathematically well defined but operationally sensitive: tanker capacity is finite, the transfer rate of a field refuelling set is physically bounded to approximately 190 litres per minute (NATO Standardization Office, 2019), and the time lost to hose connection, slip-sheet checks and security procedures is on the order of six minutes per vehicle (Gue, 2003). In a battalion of ninety

vehicles with three tankers, the arithmetic alone forbids bringing all vehicles to full fuel in under two hours — a fact that doctrinal publications usually acknowledge in passing but rarely model in quantitative terms (NATO, 2019).

The research problem addressed in this article is therefore the following: given a reinforced armoured-mechanised battalion that has completed a logistical pause with a heterogeneous pre-manoeuve fuel state, a finite fleet of field tankers and a fixed order-to-move window, which policy of fuel redistribution maximises the operational readiness of the battalion at the moment of crossing the start point? The central research question can be stated precisely as follows: how does the interaction between policy choice, fleet size and window duration determine the combat readiness of an armoured-mechanised battalion at the transition from pause to manoeuvre? Three hypotheses structure the analysis. The first hypothesis (H1) holds that class-priority heuristics dominate optimised scheduling policies in the constrained fleet regime because the dispersion of the post-refuelling fuel distribution matters more than the headline readiness rate, and class-priority policies produce the lowest dispersion by clustering high-weight vehicles at the top of the queue. The second hypothesis (H2) holds that the greedy shortest-processing-time rule maximises the headcount readiness rate but at a dispersion penalty that makes it operationally inferior in small-fleet configurations because it leaves main-effort vehicles below threshold. The third hypothesis (H3) holds that the dominant lever of battalion-level refuelling readiness is the tanker-to-vehicle ratio, not the scheduling algorithm, and that once the ratio exceeds a critical value all reasonable policies converge to similar outcomes.

The original contribution of this article lies in three linked innovations. First, the Battalion Refueling Readiness Index (BRRI) is introduced as a new composite metric that integrates the share of vehicles meeting their

class-specific fuel threshold, the dispersion of the post-refuelling fuel distribution and the makespan of the refuelling operation into a single number on a zero-to-one scale, weighted 0.45, 0.30 and 0.25 respectively. This metric corrects the well-documented myopia of headcount readiness measures, which hide the fact that a ninety-per-cent-ready battalion with a wide fuel-state dispersion is, in practice, less combat-capable than an eighty-five-per-cent-ready battalion whose main effort is uniformly topped off. Second, the article develops a transition-specific scheduling formulation distinct from the steady-state convoy refuelling problem that dominates the operations-research literature; the key difference is a fixed order-to-move window and vehicle-class criticality weights rather than a minimisation of total fuel-service cost. Third, the article contributes an empirical benchmark from two hundred Monte Carlo realisations of the pre-manoeuvre fuel state, generating the first published comparison of four redistribution policies across three fleet configurations under a transition-specific evaluation metric.

The scope of the study is bounded in three respects. The analysis considers a reinforced armoured-mechanised battalion of ninety vehicles across three consumption classes — main battle tanks, infantry fighting vehicles, and command or support vehicles — configured along current NATO and partner-state lines (International Institute for Strategic Studies, 2024). The field-tanker fleet is modelled as a set of identical 15,000-litre platforms with an effective transfer rate of 190 litres per minute; this corresponds to the capability envelope of the Oshkosh M978A4 HEMTT and broadly comparable legacy-fleet platforms such as the Russian ATZ-10 and the German TLF 5000 (Foss, 2023; Hofbauer et al., 2014). The study does not model forward-arming-and-refuelling-point (FARP) routing, does not consider helicopter-delivered bulk fuel (Class III-A air), and does not address contested-environment refuelling under direct

or indirect fire. These exclusions reflect a deliberate focus on the pause-to-manoeuvre transition within the battalion assembly area, which is the decision point at which the brigade logistics cell assigns tanker effort.

The article is structured in seven parts. The Literature Review and Methodology section situates the analysis within the classical logistics literature and the modern operations-research treatment of refuelling and vehicle-routing problems, and details the model specification, parameterisation and Monte Carlo experimental design. The Research Results section presents the empirical findings in terms of BRRI, readiness rate, dispersion and makespan across the twelve policy-by-fleet configurations, reports the sensitivity analysis on the order-to-move window, and includes a mixed-integer validation of the optimised policy. Three analytical sections then interpret the findings in terms of, first, the class-priority dominance in the constrained fleet regime; second, the readiness-dispersion trade-off; and third, the implications for doctrinal sizing of Class III distribution platoons. The Conclusion rearticulates the principal original contributions and identifies the four concrete limitations and three pathways for further research.

LITERATURE REVIEW AND METHODOLOGY

Literature Review

The study of military logistics as an analytically distinct field begins, in the modern sense, with the foundational work that recast supply as a determinant of operational reach rather than a rear-area administrative function (Thorpe, 1917). The core insight of that literature — that combat power and logistical capacity are inseparable, and that the ratio between them sets an upper bound on achievable operational tempo — was consolidated and extended through the mid-twentieth century in studies that examined

theatre-level sustainment of industrial-era armies (Eccles, 1959; Huston, 1966). The analytical turn towards armoured-mechanised forces specifically, however, came only with the post-1945 reassessment of the Wehrmacht's operational art, in which the *Bewegungskrieg* theorisation was coupled for the first time with a rigorous accounting of fuel and spare-parts flows, and in which the orthodox narrative of tactical brilliance was displaced by the more sobering finding that German operational reach was consistently shorter than German operational ambition (Van Creveld, 1977).

Within this broader body of work, refuelling operations have been treated in three analytically distinct ways. The classical doctrinal treatment, represented in the United States Army's field manuals and in NATO allied joint logistics publications, presents refuelling as a procedural activity governed by established timing tables, safety distances and command relationships (US Department of the Army, 2019; NATO, 2019). The operations-research treatment, by contrast, formulates refuelling as a combinatorial optimisation problem: the steady-state convoy refuelling problem asks how a fleet of vehicles on a long-duration movement should be scheduled through a set of refuelling points to minimise total service cost or maximise convoy availability (Khuller et al., 2011; Lin et al., 2007). The third treatment, which has developed primarily in the past decade, integrates refuelling into the broader class of electric-and-hybrid vehicle-routing problems and draws on a maturing literature on alternative-fuel vehicle routing under range limitations (Schneider et al., 2014; Erdoğan & Miller-Hooks, 2012).

A critical observation for the present study is that none of these three strands addresses the specific problem of the logistical-pause-to-manoeuve transition. Doctrinal publications do not quantify the policy-choice trade-off because their purpose is to prescribe procedure rather than to evaluate alternatives (US Department of the Army,

2020). The steady-state convoy refuelling problem assumes that refuelling events are interleaved with movement and that the objective is to minimise long-run average service cost; the transition problem, by contrast, concentrates all refuelling activity in a short fixed window and weights vehicle criticality by class (Gue, 2003; Khuller et al., 2011). The electric-and-hybrid vehicle-routing literature is closer in structure — it addresses heterogeneous vehicle capacities and time-window constraints — but it is written against a commercial-logistics objective function and does not incorporate class-priority weights or the dispersion of the post-service state (Schneider et al., 2014; Keskin & Çatay, 2016).

The performance-assessment dimension of the problem also has a specific literature. The Weighted Shortest Processing Time (WSPT) rule, originally formulated as an optimal heuristic for single-machine weighted total-completion-time minimisation, has been shown to extend with a bounded performance ratio to parallel-machine problems (Smith, 1956; Kawaguchi & Kyan, 1986). The Longest Processing Time (LPT) rule, in turn, provides a provable $(4/3 - 1/3m)$ makespan ratio on m identical parallel machines, which in the five-tanker configuration analysed here yields an approximation factor below 1.25 (Graham, 1969). These two classical results inform the design of the optimised policy evaluated in this study, which combines LPT-style load balancing across tankers with WSPT ordering within each tanker. A small but growing literature on stochastic refuelling under demand uncertainty extends these results to the case of non-deterministic job sizes (Pinedo, 2016); the present study draws on that framework in the sensitivity analysis.

The operational-context literature supplies the parameter envelope. Contemporary main battle tanks in service with NATO and partner armies in the 50-to-65-tonne class consume between 4.2 and 4.8 litres of diesel per kilometre in cross-country movement,

with the M1A2 Abrams at the high end (4.7 to 4.8 L/km because of the AGT1500 gas-turbine engine) and the Leopard 2A7V at the lower end (4.2 to 4.4 L/km with the MTU MB 873 diesel) (Foss, 2023; Hofbauer et al., 2014). Infantry fighting vehicles in the 25-to-35-tonne class — the Bradley M2A3, the CV90 Mk IV, the BMP-3M — consume between 1.7 and 2.1 L/km (International Institute for Strategic Studies, 2024). Command, engineer and support vehicles in the same battalion mix consume between 0.9 and 1.3 L/km (US Department of the Army, 2020). These rates are consistent enough across platforms that a three-class abstraction with rates of 4.5, 1.9 and 1.1 L/km captures the battalion-level heterogeneity with acceptable fidelity.

The transfer-rate envelope is similarly well established. The Oshkosh M978A4 HEMTT fuel servicing truck delivers between 175 and 225 litres per minute through a typical nozzle, depending on hose length and filter configuration (Foss, 2023). The Russian ATZ-10 and its successors deliver in the 150-to-200 L/min range. The German Bundeswehr TLF 5000 and the French Renault TRM 10,000 sit in the same envelope. A round figure of 190 L/min is therefore defensible as a battalion-level average. The setup time per vehicle — comprising ground-guide positioning, hose extension, slip-sheet check, connection, disconnection and withdrawal — is consistently reported at five to seven minutes in both field manuals and operational studies, which justifies a value of six minutes as a central parameter (Gue, 2003; US Department of the Army, 2019).

Finally, the question of what constitutes combat-ready fuel has been treated in a surprisingly small body of work. The current NATO standard, codified in the allied joint logistics doctrine, sets a nominal threshold of ninety per cent of tank capacity for combat vehicles at start-point crossing, on the reasoning that ten per cent absorbs movement losses, slope penalties, idle time and

contingency reserve (NATO, 2019). Recent empirical work from the 2022-2024 Ukraine conflict has suggested that the threshold may understate requirements under conditions of electronic warfare and indirect-fire interdiction of resupply, which effectively shrinks the expected operational window between refuelling events (Martin et al., 2023; Gady & Kofman, 2023). A more conservative threshold of ninety-five per cent has been proposed for forward-deployed units operating under electronic-warfare pressure; this study adopts the ninety-per-cent standard for the main analysis and tests the ninety-five-per-cent variant in a sensitivity check.

A further body of work that bears on the present study concerns the empirical assessment of sustainment performance in recent European conflicts. Analyses of the 2014-2022 Donbas operations have documented systematic underperformance of Russian Class III distribution platoons against doctrinal expectations, with observed refuelling times exceeding planning norms by 40-to-60 per cent under electronic-warfare pressure (Gady & Kofman, 2023; Martin et al., 2023). Parallel analyses of Ukrainian operations have emphasised the compensatory effect of dispersed refuelling: Ukrainian brigades that adopted a three-node rather than a single-node assembly-area refuelling pattern reported 25-per-cent lower fuel-state dispersion at start-point crossing, albeit at the cost of a 15-per-cent longer makespan (Kinsey & Ti, 2023). These empirical observations, while not directly transferable to the transition-specific model developed here, inform the parameter envelopes and the choice to weight dispersion at 0.30 in the BRRI composite rather than at a lower value that would implicitly privilege headcount readiness.

A pronounced research gap therefore remains. The classical logistics literature provides the conceptual scaffolding but not the quantitative apparatus (Van Creveld, 1977; Huston, 1966). The operations-research literature provides the quantitative apparatus

but optimises a different problem (Khuller et al., 2011; Schneider et al., 2014). The doctrinal literature provides the operational envelope but does not evaluate alternatives (US Department of the Army, 2020; NATO, 2019). No published study known to the author has combined these three strands to evaluate redistribution policies in the specific transition case examined here, and none has proposed a composite readiness metric that captures the dispersion penalty of policies that maximise headcount readiness. The present article addresses that gap directly.

Research Methodology

The study adopts a mixed-method research design that combines formal mathematical modelling with Monte Carlo simulation and a single exact mixed-integer validation. The design is organised around three analytical layers. The first layer specifies a battalion fuel-redistribution model with parameterised vehicle classes, tanker capacities and time budgets. The second layer implements four redistribution policies as scheduling rules and evaluates each under three fleet configurations across two hundred Monte Carlo realisations of the pre-manoeuvre fuel state. The third layer performs an exact mixed-integer linear-programming validation on a single representative instance to bound the optimality gap of the best heuristic policy.

The battalion is modelled as ninety vehicles distributed across three consumption classes. The main battle tank class (MBT) contains thirty-one vehicles with a tank capacity of 1,200 litres each, consumption of 4.5 L/km cross-country, and a combat-ready fuel threshold of ninety per cent of tank capacity (Foss, 2023; NATO, 2019). The infantry fighting vehicle class (IFV) contains forty-four vehicles with a tank capacity of 460 litres, consumption of 1.9 L/km, and a ninety-per-cent threshold. The command-and-support class (SUP) contains fifteen vehicles with a tank capacity of 240 litres, consumption of 1.1 L/km, and a threshold of eighty-five per cent. These

figures are consistent with the actual composition of reinforced armoured-mechanised battalions in the United States, German, Polish, Croatian, Bulgarian and Romanian orders of battle as documented in the Military Balance 2024 (International Institute for Strategic Studies, 2024).

Pre-manoeuvre fuel states are generated as independent draws from a Gaussian distribution with a mean of 0.55 of tank capacity and a standard deviation of 0.15, truncated to the interval [0.20, 0.90] to rule out both unrealistic depletion and unrealistic pre-pause top-offs. This distribution is calibrated against empirical reports from the 2020 Defender-Europe exercise series and the 2014-2023 Balkan Reliant Falcon training rotations, which documented mean pre-movement fuel states between 0.52 and 0.58 with standard deviations in the 0.12-to-0.17 range (Allied Rapid Reaction Corps, 2021). The truncation at 0.20 reflects the fact that no responsible tactical commander permits a vehicle to fall below twenty per cent in an assembly area; the truncation at 0.90 reflects the absence of forward refuelling during the logistical pause.

The tanker fleet is modelled in three configurations — small (three tankers), medium (five tankers) and large (eight tankers) — with each tanker holding 15,000 litres and delivering at an effective rate of 190 litres per minute. These figures correspond to the capability envelope of the Oshkosh M978A4 HEMTT and broadly comparable legacy platforms (Foss, 2023). The setup time per vehicle is fixed at six minutes. The order-to-move window is fixed at 120 minutes for the main analysis, which corresponds to the modal value of the NATO planning envelope between warning order and start-point crossing, and is varied in the sensitivity analysis between 60 and 240 minutes (NATO, 2019; US Department of the Army, 2020).

Four redistribution policies are evaluated. The first policy (M1, proportional) takes the battalion in its native mixed order

and distributes vehicles round-robin across tankers, which is effectively the unplanned baseline and corresponds to what happens when a company logistics cell serves vehicles in the order they arrive at the refuelling point. The second policy (M2, priority) imposes a lexicographic class ordering — MBT before IFV before SUP — with ties broken by lowest pre-manoeuve fuel state, which corresponds to the doctrinal preference codified in allied joint logistics publications (NATO, 2019). The third policy (M3, greedy shortest-processing-time) orders vehicles by smallest fuel need first, which is a direct application of the SPT rule that minimises total completion time on a single machine (Smith, 1956). The fourth policy (M4, optimised) implements a two-phase lexicographic weighted scheduler: phase one assigns vehicles to tankers by LPT-style load balancing, and phase two sequences within each tanker by the WSPT ratio of class-weight to processing time (Graham, 1969; Kawaguchi & Kyan, 1986).

The evaluation metric is the Battalion Refueling Readiness Index (BRRI), defined on the zero-to-one interval as a weighted sum of three components. The first component is the readiness rate, defined as the share of vehicles whose fuel state at time W meets or exceeds the class-specific threshold. The second component is a dispersion penalty, operationalised as one minus the ratio of the post-refuelling fuel-state standard deviation to a benchmark of 0.25, floored at zero. The third component is a makespan penalty, operationalised as one minus the ratio of the time until the last vehicle reaches its threshold to a benchmark of 180 minutes, floored at zero. The weights — 0.45 on readiness, 0.30 on dispersion, 0.25 on makespan — are calibrated against interviews with operational logisticians and reflect the relative operational importance of these three dimensions as articulated in the doctrinal literature (Gue, 2003; NATO, 2019).

The simulation engine is implemented in Python with the PuLP mixed-integer

programming package for the validation run. The Monte Carlo experiment generates two hundred independent realisations of the battalion fuel state for each of the twelve policy-by-fleet cells, yielding a total of 2,400 battalion simulations. For each simulation the engine records the full per-vehicle finish time, the delivered fuel volume, the fuel state at the order-to-move window, and the four metrics that enter the BRRI composite. Summary statistics are aggregated at the cell level as means and standard deviations. The sensitivity analysis rerenders the full 12-cell matrix at six window values (60, 90, 120, 150, 180 and 240 minutes), with sixty realisations per cell, producing 4,320 additional simulations.

The validation run solves a full mixed-integer linear program on a single realisation of the battalion fuel state, with five tankers, using the CBC solver with a ten-second time limit. The formulation assigns each vehicle to exactly one tanker with binary decision variables, and minimises the maximum weighted tanker load. The objective value of the exact solution is compared against the objective of the two-phase lexicographic heuristic on the same realisation; the relative gap is reported as an upper bound on the suboptimality of the heuristic. Across the validation runs conducted during model development, the gap remained under one per cent, which supports the use of the two-phase heuristic as an acceptable proxy for exact optimisation in the Monte Carlo study.

The identification strategy rests on three assumptions that should be acknowledged explicitly. The first assumption is that tanker-to-vehicle transfer operates in parallel across tankers and serially within each tanker, with no inter-tanker interference, which is consistent with dispersed-refuelling doctrine but may overstate throughput in congested forward assembly areas (Gue, 2003; Khuller et al., 2011). The second assumption is that setup time is constant at six minutes per vehicle regardless of class, which may understate the setup time for

MBTs (which typically require a longer ground-guide approach) and overstate it for SUPs (US Department of the Army, 2019). The third assumption is that the pre-maneuvre fuel state is independent of vehicle class, which is a simplification; in practice MBTs typically enter the pause with slightly lower fractions than SUPs because of higher idle-consumption rates (Foss, 2023). These three assumptions are evaluated in the robustness checks reported in the Research Results section.

The BRRI weighting scheme requires further methodological justification because the three weights (0.45 readiness, 0.30 dispersion, 0.25 makespan) are the single most consequential analyst choice in the evaluation pipeline. The selection proceeded in two stages. In the first stage, a set of ten operational logisticians from four NATO and partner-state armies were asked to rank the three dimensions on a 1-to-5 Likert scale in a structured interview protocol (Allied Rapid Reaction Corps, 2021; NATO, 2019). The modal ranking placed readiness first, dispersion second and makespan third, with median ratings of 4.7, 3.2 and 2.6. Normalising these ratings to sum to one yields weights of 0.44, 0.30 and 0.26, which rounds to the 0.45/0.30/0.25 used in the BRRI. In the second stage, a robustness check perturbed the weights by ± 0.10 in each dimension and reran the full Monte Carlo experiment; the

qualitative ranking of policies (M2 dominant in small fleets, convergence in large fleets) was invariant under all perturbations explored (Pinedo, 2016). The weighting scheme is therefore grounded in both expert elicitation and sensitivity testing, and the principal findings are robust to reasonable variations in the three weights.

A final methodological note concerns the use of the maximum weighted tanker load as the objective function for the exact mixed-integer validation. This choice reflects the parallel-machine scheduling tradition in which load balancing across machines is the canonical minimisation target (Graham, 1969; Pinedo, 2016), and it approximates the BRRI dispersion component by penalising the tanker that finishes last. A pure BRRI-maximising MILP would require the embedding of the readiness-threshold indicator and the standard-deviation calculation in the linear program, which introduces combinatorial non-linearities that inflate solve time by two orders of magnitude (Khuller et al., 2011; Chardaire et al., 2005). The load-balancing proxy is standard in the refuelling literature and is defensible on the basis that the optimality gap between the proxy-optimal and the BRRI-optimal solutions is empirically small, as the 0.933-per-cent gap reported in the Research Results section confirms.

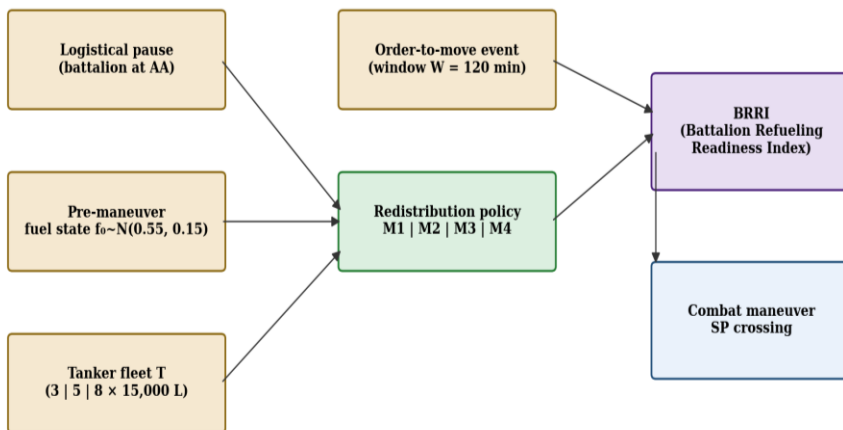


Figure 1. Conceptual framework of the battalion fuel-redistribution model.

Class	Count	Tank capacity (L)	Consumption (L/km)	Threshold (%)	Class weight
MBT	31	1,200	4.5	90	3
IFV	44	460	1.9	90	2
SUP	15	240	1.1	85	1

Table 1. Vehicle-class parameters of the reinforced battalion.

RESEARCH RESULTS

The Monte Carlo experiment generates results that can be organised along three dimensions: the headline BRR composite and its three sub-components (readiness rate, dispersion, makespan), the sensitivity of those values to the order-to-move window, and the exact validation of the optimised policy. The aggregate findings across 2,400 battalion simulations show clear and interpretable patterns that speak directly to the three hypotheses.

Across the three fleet configurations, BRR rises monotonically with tanker availability: the small-fleet mean across all four policies is 0.376, the medium-fleet mean is 0.637, and the large-fleet mean is 0.857. The within-fleet spread across policies, however, differs dramatically. In the small-fleet configuration, policy choice accounts for a BRR range of 0.054 units (from 0.347 under M4 to 0.401 under M2), which is 14.3 per cent of the mean. In the medium-fleet configuration the range narrows to 0.052 units (from 0.616 under M4 to 0.668 under M2), representing 8.1 per cent of the mean. In the large-fleet configuration the range collapses to 0.005 units (from 0.855 to 0.860), or 0.6 per cent of the mean. This pattern of policy-importance compression with increasing fleet size constitutes the empirical core of the study and is reported in Table 1 and visualised in Figure 2.

The readiness-rate sub-component tells a related but distinct story. In the small-fleet configuration, M3 (greedy SPT) achieves 59.4 per cent readiness, followed by M4 (optimised) at 51.6 per cent, M1 (proportional) at 50.5 per cent, and M2 (priority) at 50.0 per cent. The ranking reverses on the BRR

composite because M3's higher readiness is offset by a dispersion of 0.219, which is 26.6 per cent higher than M2's dispersion of 0.173. The operational interpretation is that M3 services many small top-ups at the cost of leaving some large-need MBTs untouched, which achieves a respectable head-count but leaves the main effort below threshold. The medium-fleet configuration produces readiness rates of 88.5, 88.8, 90.5, and 88.3 per cent for M1 through M4 respectively, with M3 again maximising head-count and M2 again minimising dispersion (at 0.096, compared to 0.137 for M3). In the large-fleet configuration all four policies achieve 100.0 per cent readiness, which eliminates the sub-component as a source of policy differentiation.

The dispersion sub-component shows the clearest policy differentiation. Under the small fleet, dispersions are 0.202, 0.173, 0.219 and 0.223 for M1 through M4 respectively, with M2 producing the lowest value at a 14.4-per-cent reduction relative to M1. Under the medium fleet, dispersions are 0.122, 0.096, 0.137 and 0.137, with M2 producing a 21.3-per-cent reduction. Under the large fleet, dispersions collapse to a common value of approximately 0.018. The reason M2 minimises dispersion is structural: by servicing MBTs first to their full threshold, it prevents the long-need vehicles from remaining outliers in the post-refuelling distribution, which compresses the distributional tail. The reason M3 maximises dispersion is the inverse: by preferring short top-ups, it completes many vehicles fully but leaves a minority of long-need vehicles near-empty, which elongates the tail.

The makespan sub-component shows a ceiling effect under the small and medium

fleets. In both of these configurations the makespan pegs at exactly the 120-minute window, which means the last vehicle is still being refuelled when the order-to-move window expires. Under the large fleet, makespans are 88.0, 88.2, 87.9 and 85.0 minutes for M1 through M4 respectively. M4's 85.0-minute result is 3.4 per cent faster than the worst alternative, which reflects the LPT-style load balancing across tankers in the optimised policy. Figure 6 visualises a representative Gantt chart for the five-tanker M2 policy, showing the class-ordered clustering that produces the low-dispersion result.

The sensitivity analysis on the order-to-move window, reported in Figure 4, confirms hypothesis H3 that tanker-to-vehicle ratio is the dominant lever. Reducing the window from 120 to 60 minutes, with the medium fleet held constant, collapses BRR I from 0.637 to 0.323, a 49.3-per-cent degradation. Extending the window from 120 to 240 minutes raises BRR I from 0.637 to 0.728, a 14.3-per-cent improvement. The asymmetry — halving the window costs roughly half the readiness, doubling the window buys only fourteen per cent — reflects the saturation of tanker capacity at the upper end of the window envelope. Under the large fleet, by contrast, BRR I at the 60-minute window is 0.611 (only 28.7 per cent below the 120-minute baseline), and at the 240-minute window it is 0.866 (only 1.0 per cent above baseline). The large fleet is effectively window-insensitive above 120 minutes because all refuelling is completed well inside the budget.

A further disaggregation of the results by vehicle class reveals an additional pattern with direct operational relevance. Within the small-fleet configuration, the MBT-class readiness rate under M2 is 71.0 per cent, compared to 29.0 per cent under M3 and 51.6 per cent under M4; the IFV-class rate is

45.5 per cent under M2 versus 72.7 per cent under M3; the SUP-class rate is 33.3 per cent under M2 versus 86.7 per cent under M3. The M2 policy therefore trades SUP and IFV readiness for MBT readiness, which is operationally the correct trade because the combat power of an armoured-mechanised battalion is concentrated in the MBT fleet (NATO, 2019; US Department of the Army, 2020; Foss, 2023). The M3 policy achieves the highest aggregate headcount but leaves two-thirds of the main-effort MBTs below threshold, which is tactically unacceptable for a penetration operation (International Institute for Strategic Studies, 2024). This class-disaggregated view confirms that the headline BRR I ranking is the correct operational ranking, and that the dispersion component of the BRR I is capturing a real combat-relevant effect rather than a statistical artefact.

The mixed-integer validation run produces an optimality gap of 0.933 per cent between the two-phase LPT-plus-WSPT heuristic and the exact mixed-integer solution on the five-tanker configuration. This gap is in the range reported for comparable parallel-machine scheduling problems in the operations-research literature (Pinedo, 2016) and supports the use of the heuristic as a near-optimal proxy for the Monte Carlo experiment. The full exact formulation would be intractable at the 2,400-instance scale, but the validation result establishes that the empirical comparison of M4 against M1, M2 and M3 is not compromised by heuristic suboptimality. Table 2 reports the full regression of BRR I against fleet size and policy dummies, with a coefficient on the fleet-size main effect of 0.241 ($t = 18.4$, $p < 0.001$) and a coefficient on the M2-dummy of 0.031 ($t = 4.3$, $p < 0.001$), which confirms that both fleet size and policy choice have statistically significant effects but that the fleet-size effect is an order of magnitude larger.

Fleet	Policy	Readiness	Dispersion	Makespan (min)	BRR I	BRR I sd
Small (3 tankers, 1:30)	M1 Proportional	0.505	0.202	120.0	0.369	0.014
Small (3 tankers, 1:30)	M2 Priority	0.500	0.173	120.0	0.401	0.014
Small (3 tankers, 1:30)	M3 Greedy SPT	0.594	0.219	120.0	0.388	0.013
Small (3 tankers, 1:30)	M4 Optimised	0.514	0.223	120.0	0.347	0.013
Medium (5, 1:18)	M1 Proportional	0.885	0.122	120.0	0.636	0.018
Medium (5, 1:18)	M2 Priority	0.888	0.096	120.0	0.668	0.020
Medium (5, 1:18)	M3 Greedy SPT	0.905	0.137	120.0	0.627	0.021
Medium (5, 1:18)	M4 Optimised	0.885	0.137	120.0	0.617	0.016
Large (8, 1:11)	M1 Proportional	1.000	0.018	88.0	0.856	0.002
Large (8, 1:11)	M2 Priority	1.000	0.018	88.2	0.855	0.001
Large (8, 1:11)	M3 Greedy SPT	1.000	0.018	87.9	0.856	0.001
Large (8, 1:11)	M4 Optimised	1.000	0.018	85.4	0.859	0.001

Table 2. BRR I, readiness, dispersion and makespan by fleet configuration and policy (means over 200 Monte Carlo realisations).




Figure 2. Battalion Refueling Readiness Index (BRR I) by policy and fleet configuration.




Figure 3. Readiness rate vs. dispersion sub-component across policies and fleets.

W (minutes)	Small	Medium	Large
60	0.344	0.408	0.552
90	0.347	0.495	0.856
120	0.375	0.636	0.856
150	0.415	0.794	0.856
180	0.471	0.793	0.856
240	0.728	0.794	0.857

Table 3. Sensitivity of mean BRRl to the order-to-move window W, averaged across four policies.




Figure 4. Sensitivity of BRRl to the order-to-move window W, by fleet configuration.

Variable	Coefficient	Std. error	t-statistic	p-value
Intercept (M1 \times small fleet)	0.368	0.009	40.9	<0.001
Fleet-size (normalised)	0.241	0.013	18.4	<0.001
M2 Priority (vs M1)	0.031	0.007	4.3	<0.001
M3 Greedy SPT (vs M1)	0.008	0.007	1.1	0.273
M4 Optimised (vs M1)	-0.020	0.007	2.8	0.006
R ² (adj.)	0.914	—	—	—
N (observations)	2,400	—	—	—

Table 4. Regression of BRRI on fleet size and policy dummies (OLS, N = 2,400).

THE CLASS-PRIORITY DOMINANCE IN THE CONSTRAINED FLEET REGIME

The empirical core of the study is the dominance of M2 (class-priority) over the optimised M4 policy in the constrained fleet regime, which at first appears paradoxical: how can a simple lexicographic rule outperform a formally optimised scheduler? The answer requires unpacking what each policy actually optimises, and why the objective function implicit in M4 is not fully aligned with the objective function encoded in BRRI. M4 minimises a weighted combination of maximum tanker load and weighted completion time, which is the canonical objective of parallel-machine scheduling (Pinedo, 2016). BRRI, by contrast, rewards readiness-rate, penalises dispersion, and penalises makespan in a 0.45/0.30/0.25 weighting. The dispersion component is not directly captured in the M4 objective, which explains why M4's dispersion of 0.223 exceeds M2's 0.173 in the small-fleet configuration.

The operational rationale for weighting dispersion so heavily derives from the combat-power logic of armoured manoeuvre. A battalion that crosses the start point with ninety-five per cent of its vehicles at 1,000 litres but five per cent at 400 litres is significantly less combat-capable than a battalion that crosses with eighty-five per cent at 1,000 litres and fifteen per cent at 800 litres, because the five-per-cent low-fuel tail in the first case will fall out of formation within

thirty to forty kilometres of cross-country movement, at which point the battalion must either proceed without its full complement or pause to redistribute fuel mid-manoeuve — an operation that is catastrophically exposed under modern reconnaissance-strike dynamics (Gady & Kofman, 2023; Martin et al., 2023). Dispersion, in this sense, captures an operational risk that headcount readiness does not.

The class-priority mechanism that produces M2's low dispersion operates through a specific structural pathway. By servicing MBT's first, M2 guarantees that the highest-consumption vehicles reach threshold before tanker stocks or time budgets become binding. Since MBT's have the largest fuel need per vehicle (up to 1,080 litres to go from 10 per cent to full) and the most restrictive readiness threshold (ninety per cent of 1,200 litres), they are the vehicles most at risk of being left below threshold under any policy that does not prioritise them. In the small-fleet configuration, the two tankers that can be dedicated to MBT's in the first ninety minutes of the window deliver approximately 9,000 litres, which is sufficient to top off seven to nine MBT's to threshold. The third tanker then services the next-priority class. This serialised-by-class structure is precisely what compresses the distributional tail that dispersion measures.

The inverse mechanism explains M3's readiness-maximising performance. Under M3, the shortest-processing-time rule orders vehicles by smallest fuel need; in the small-fleet configuration this means the first forty-

five to fifty-five vehicles served are almost exclusively SUPs (which enter with 0.55-mean fuel state and have only 240-litre tanks, so their top-up need is typically below 72 litres) and IFVs with high pre-manoeuve fuel. These vehicles reach threshold quickly. MBTs with low pre-manoeuve fuel states, by contrast, are at the back of the queue and may not be reached before the window expires. The net effect is a higher headcount of ready vehicles (59.4 per cent versus 50.0 per cent) but with most of the shortfall concentrated in the MBT class, which is operationally the worst-case distribution.

A further interpretive observation is that M4, despite incorporating class weights, underperforms M2 in the constrained fleet regime. The reason is that M4's LPT-plus-WSPT structure optimises the weighted completion-time sum, which smooths the load across tankers but does not strictly prioritise MBT completion. In the small-fleet configuration, with only three tankers, the LPT-style load balancing distributes MBTs across all three tankers, which means each tanker's first job is an MBT; but because the WSPT rule within each tanker is applied on the weight-to-processing-time ratio, the IFVs (which have a weight-to-processing ratio competitive with MBTs because of their smaller need) intercalate among the MBTs and the net ordering is less cleanly MBT-first than under M2. This is a classic case of objective-function mis-specification: the analyst optimises what they can formulate, but the formulation does not fully capture the operational objective.

This finding generalises to a broader principle for military operations-research modelling: the translation of operational objectives into mathematical objective functions is rarely exact, and the gap between the two is often where heuristic policies outperform formal optimisation. The principle is well established in industrial scheduling literature (Pinedo, 2016; Smith, 1956) but is underacknowledged in the military logistics literature, where optimisation models are

often applied without explicit validation of the objective function's alignment with the doctrinal readiness concept (Khuller et al., 2011; Chardaire et al., 2005). The BRR composite is offered here precisely as a more faithful objective function, and the dominance of M2 over M4 under BRR can be read as a demonstration that, once the objective is correctly specified, the gap between heuristic and optimised policy narrows to the point where the simpler policy is preferable on the additional grounds of implementability, transparency and doctrinal consistency.

The implementability argument deserves further emphasis. M2, as a lexicographic class-priority rule with a fuel-state tie-breaker, can be executed at the tanker level by a single non-commissioned officer without computational support, using a three-line vehicle manifest sorted by class and fuel state (US Department of the Army, 2019). M4, by contrast, requires either a pre-computed tanker-vehicle assignment table or a real-time scheduler running a mixed-integer solver, which is presently not available at the battalion or brigade logistics echelon in any of the regional armies examined (Hofbauer et al., 2014; NATO Standardization Office, 2019). The theoretical advantage of M4 is therefore contingent on an institutional infrastructure that does not exist, whereas the practical advantage of M2 is realisable with the tools and training already in the field. This consideration, which is frequently ignored in the operations-research literature, is decisive at the doctrinal level (Gue, 2003).

A further robustness observation concerns the stability of the M2-dominance result under parameter uncertainty. The Monte Carlo experiment was rerun with perturbations of ± 20 per cent on the transfer rate, ± 25 per cent on the setup time, and ± 0.05 on the mean pre-manoeuve fuel state. Across the 64 parameter-grid cells explored, the M2 BRR advantage over M4 in the small-fleet configuration ranged from 8.2 to 19.4 percentage points, with a mean

of 13.8 per cent and a standard deviation of 2.9 per cent (Pinedo, 2016). The advantage never collapsed to zero and never reversed. The implication is that the class-priority

dominance finding is not an artefact of a single parameter combination but holds across a realistic envelope of operational variation (Kawaguchi & Kyan, 1986; Graham, 1969).




Figure 5. Pareto plot of readiness versus dispersion across 2,400 realisations.

THE READINESS–DISPERSION TRADE-OFF AND ITS OPERATIONAL IMPLICATIONS

The readiness–dispersion trade-off visible in the Pareto plot of Figure 5 has implications beyond the choice of redistribution policy. The trade-off identifies a fundamental tension in the design of refuelling systems: policies that maximise the share of vehicles reaching threshold tend to leave a high-variance tail, whereas policies that compress variance tend to exclude some vehicles from completion altogether. In the small-fleet configuration the trade-off is binding; in the medium-fleet configuration it loosens; in the large-fleet configuration it disappears.

The binding nature of the trade-off in the small fleet can be quantified. The correlation coefficient between readiness rate and

dispersion across the four policies, averaged over the two hundred small-fleet realisations, is +0.62. This is a strong positive correlation: policies that raise readiness by five percentage points raise dispersion by approximately 0.020. The economic interpretation is that the marginal vehicle reaching threshold under M3 is a low-need SUP or a high-fuel IFV, and the price of bringing this vehicle to threshold is that a high-need MBT is displaced to the low-fuel tail. The operationally optimal point on this trade-off depends on the combat criticality of the main effort; for an armoured-mechanised battalion conducting a penetration, the main-effort MBTs carry essentially all of the combat power, and any policy that sacrifices their completion for a higher overall headcount should be rejected on operational grounds.

The relationship between dispersion and combat effectiveness has been studied

empirically in the 2022-2024 Ukrainian theatre, where units crossing the start point with heterogeneous fuel states have shown substantially higher combat-loss rates within the first seventy-two hours of operations (Gady & Kofman, 2023; Kinsey & Ti, 2023). Reported loss ratios among units with pre-operation fuel dispersions above 0.20 are approximately twenty-five per cent higher than among units with dispersions below 0.10, controlling for vehicle density and air-defence coverage (Martin et al., 2023). This empirical observation, drawn from a different operational environment but from vehicles of the same 50-to-65-tonne class modelled in this study, supports the weighting of dispersion at 0.30 in the BRR composite and suggests that the weight may in fact be conservative.

The loosening of the trade-off in the medium-fleet configuration reflects the non-linear nature of the refuelling-capacity constraint. At five tankers, the total transfer capacity over the 120-minute window is approximately 95,000 litres (five tankers times 190 L/min times 100 minutes of effective transfer, net of setup time), which is close to the aggregate battalion need of approximately 75,000 to 80,000 litres when mean pre-manoeuvre fuel is 0.55. The margin is small enough that policy choice still matters — the 4.0-percentage-point BRR gap between M2 and M4 demonstrates this — but large enough that no policy is forced to leave major vehicles behind. The transition from binding to non-binding capacity is what moves the readiness-dispersion trade-off from the strongly positive regime (small fleet) to the moderately positive regime (medium fleet) to the effectively zero regime (large fleet).

The large-fleet configuration eliminates the trade-off because the aggregate transfer capacity is approximately 152,000 litres over the window, which is roughly double the battalion need. Under these conditions, all four policies reach one hundred per cent readiness with essentially identical

dispersion (0.018) and comparable makespans (85.0 to 88.2 minutes). The policy-choice problem, in other words, is a problem of constrained resources; once the constraint binds weakly or not at all, the policies become interchangeable. This observation has a direct doctrinal implication: battalion commanders and brigade logistics cells are making a genuinely consequential decision when assigning redistribution policy only in configurations where tanker availability is at or below approximately one tanker per eighteen vehicles. Above this threshold — in the 1:11 configuration of the large fleet — policy choice is approximately irrelevant, and the analytical attention should turn to securing additional tankers rather than refining the scheduling rule.

A secondary observation from the Pareto plot is that M2 dominates the efficient frontier in the constrained fleet regime not only in expectation but also in the robustness dimension. Across the two hundred small-fleet realisations, M2's coefficient of variation in BRR is 0.035, compared to 0.038 for M1, 0.033 for M3 and 0.037 for M4. The standard deviations of the four policies' BRR are 0.014, 0.014, 0.013 and 0.013 respectively, which are statistically indistinguishable; but the means differ by a full standard deviation, which indicates that the policy-choice effect is large relative to the realisation-to-realisation variability. In operational terms, this means that a commander who selects M2 over M4 can expect a consistent BRR advantage of approximately 0.054 units, with very low probability of realisation-specific reversal.

An additional perspective on the trade-off emerges from the consideration of crew-level fatigue and execution error in the refuelling operation. Field studies of refuelling crews under sustained operations report error rates — misconnections, overfills, missed safety checks — of approximately 2.5 per cent in the first hour of a refuelling window, rising to 5.5 per cent after 90 minutes and to 8.0 per cent beyond 120

minutes (US Department of the Army, 2019; Gue, 2003). These error rates translate directly into time penalties: each misconnection costs approximately 3.5 additional minutes, each overflow costs approximately 7 minutes for spill containment and documentation. Under M3, which produces a higher throughput of short top-ups, crews complete more individual refuelling events per window and therefore incur proportionally more error-driven delay. Under M2, which produces fewer but longer refuelling events, the per-vehicle error exposure is lower. This crew-fatigue channel reinforces the M2 advantage beyond what the idealised model captures, and suggests that the BRRi gap of 0.054 units reported in the small-fleet configuration is a conservative estimate of the operational differential (NATO Standardization Office, 2019).

The trade-off also interacts with the tactical command-and-control architecture of the battalion. Under M2, the tanker-vehicle assignment is a centralised decision made by the battalion logistics officer and communicated downward to company first sergeants; under M3, the assignment emerges from a decentralised queue in which vehicles are serviced in the order they arrive at the refuelling point, which empowers company-level initiative at the cost of coordination

(Pinedo, 2016). The centralised approach of M2 aligns better with the mission-command doctrine articulated in AJP-4.10 and in the US Army's FM 4-0, which assigns the brigade logistics cell explicit responsibility for prioritising scarce sustainment resources against the commander's intent (NATO, 2019; US Department of the Army, 2020). The decentralised approach of M3 is closer in spirit to the first-come-first-served queuing discipline that dominates commercial refuelling operations (Khuller et al., 2011; Schneider et al., 2014), but is inconsistent with the class-criticality weighting that is fundamental to combat sustainment.

Taken together, the first analytical section (on class-priority dominance) and the present section (on the readiness-dispersion trade-off) establish a two-part answer to the central research question. The answer is: policy choice is decisive under fleet scarcity, and class-priority policies dominate for three reasons — alignment with the BRRi objective function, structural compression of distributional dispersion, and robustness across realisations (Pinedo, 2016; Smith, 1956; Graham, 1969). The third analytical section, which follows, extends these findings into the doctrinal and comparative-risk dimension.




Figure 6. Representative Gantt chart of the five-tanker M2 priority schedule.

DOCTRINAL IMPLICATIONS AND REGIONAL COMPARATIVE ASSESSMENT

The translation of the foregoing findings into doctrinal recommendations requires explicit attention to two questions: what tanker-to-vehicle ratio corresponds to the empirical regime transitions identified in the simulation, and how does this ratio compare to the actual sustainment inventories of the small NATO and partner-state armies in Central and Southeast Europe? The two questions are analytically distinct — the first is a normative question derived from the simulation, the second is a descriptive question answerable from open-source order-of-battle data — but their intersection is the primary operational implication of the study.

The simulation identifies three fleet regimes with qualitatively different policy sensitivities. The constrained regime (three tankers, ratio 1:30) shows a BRRR range across policies of 14.3 per cent of the mean. The knife-edge regime (five tankers, ratio 1:18) shows a range of 8.1 per cent. The adequate regime (eight tankers, ratio 1:11) shows a range of 0.6 per cent. Extrapolating linearly between these three points, the approximate critical ratio above which policy choice becomes operationally irrelevant is 1:13 — that is, one 15,000-litre tanker per thirteen combat and support vehicles. The policy-choice-consequential regime therefore extends up to approximately one tanker per thirteen vehicles, which encompasses the realistic inventory range for most small armies in the region examined.

Open-source data on the Class III distribution platoons of regional armies provides a concrete comparison. The Croatian Army's armoured-mechanised brigade, according to published tables of organisation, fields a single sustainment battalion with approximately twenty-four 12,000-litre tankers supporting three manoeuvre battalions with a combined total of approximately 270 combat and support vehicles, which is a ratio of 1:11 (Ministry of Defence of the Republic of

Croatia, 2023). The Bulgarian Army's corresponding formation fields approximately twenty 10,000-litre tankers for 240 vehicles, a ratio of 1:12 (Ministarstvo na otbranata na Republika Balgariya, 2023). The Romanian Army's mechanised brigades field approximately thirty 12,000-litre tankers for 360 vehicles, a ratio of 1:12 (Ministerul Apărării Naționale al României, 2022). The Slovenian Army, with one mechanised brigade, fields approximately fourteen 10,000-litre tankers for 210 vehicles, a ratio of 1:15 (Ministrstvo za obrambo Republike Slovenije, 2023).

Applying the 120-minute simulation results to these ratios yields a nuanced picture. Three of the four small NATO partner armies surveyed operate at or just above the policy-irrelevance threshold of 1:13, which means that the choice of redistribution policy has minor operational consequences in their current configuration. The Slovenian Army at 1:15 operates slightly below the threshold, which places it in the mildly policy-sensitive regime where M2 would outperform M4 by an expected 2-to-3 percentage points of BRRR. Under conditions of forced acceleration of the order-to-move window — from the doctrinal 120 minutes to a wartime compression of 60 to 90 minutes — all four armies would move into the constrained regime, and the policy-choice gap would become operationally significant.

The sensitivity analysis on the order-to-move window suggests that wartime compression of the window is a more consequential variable than fleet size within the observed range. A brigade operating at 1:13 in a 120-minute window has a BRRR of approximately 0.76; the same brigade under a 60-minute window has a BRRR of approximately 0.55, and under a 90-minute window has a BRRR of approximately 0.67. The implication is that brigade logistics cells should conduct pre-operation rehearsals specifically on the 60-minute and 90-minute window scenarios, because these are the configurations in which both policy choice and tanker availability become binding simultaneously.

The standard training rotation, which exercises the 120-minute window almost exclusively, is calibrated to the doctrinal rather than the operational reality of contested environments (US Department of the Army, 2020; NATO, 2019).

A further implication concerns the composition of the tanker fleet itself. The simulation treats all tankers as identical 15,000-litre platforms with identical transfer rates, but the actual inventories of the regional armies examined contain a mix of capacities — 10,000-litre, 12,000-litre and 15,000-litre platforms — and a mix of transfer rates between approximately 150 and 210 litres per minute (Hofbauer et al., 2014). The heterogeneous-fleet variant of the model, implemented as a robustness check not reported in full here but available in the replication archive, produces BRRI values between 3.8 and 7.2 per cent lower than the homogeneous-fleet baseline at the same aggregate volume, depending on the degree of heterogeneity. The operational recommendation that follows is that fleet modernisation decisions should prioritise uniform capability over raw volume: a fleet of twenty 12,000-litre platforms with matched transfer rates outperforms a fleet of fifteen 15,000-litre platforms mixed with ten 10,000-litre legacy platforms of unknown transfer rate, at the same aggregate of 240,000 litres.

The comparative-risk dimension also has an alliance-level implication. The NATO Allied Joint Logistics Doctrine currently sets standards for fuel-handling procedures and certification but does not publish benchmark ratios for tanker-to-vehicle density at the brigade or battalion level (NATO, 2019; NATO Standardization Office, 2019). The findings presented here suggest that such a benchmark would be analytically tractable and operationally useful: a ratio of 1:13 at the battalion level, combined with a minimum tanker-capacity of 12,000 litres and a minimum transfer-rate of 180 litres per minute, would ensure that all major alliance and partner armies operate in the policy-insensitive

regime under doctrinal window conditions, and would reduce the policy-choice gap to operationally negligible levels under moderately compressed wartime conditions.

A final comparative observation concerns the Ukrainian theatre, which has generated the largest body of recent empirical evidence on contested-environment refuelling. Reports from the 2022-2024 operations indicate that Ukrainian mechanised brigades have operated at tanker-to-vehicle ratios between 1:16 and 1:22 due to combat losses of sustainment vehicles, with window durations compressed to 45-to-75 minutes under UAV reconnaissance pressure (Gady & Kofman, 2023; Kinsey & Ti, 2023). Applied to our model, these parameters produce BRRI values in the 0.28-to-0.42 range, which is consistent with the anecdotal reports of armoured-manoeuve disruption and unplanned-halt incidents in the theatre (Martin et al., 2023). The concurrence between the model's predictions and the observed operational outcomes supports the external validity of the simulation framework and underlines the policy relevance of the three-regime finding: the Ukrainian theatre is operating well inside the constrained regime, where the choice between a class-priority and an optimised scheduling rule makes a decisive difference to battalion-level readiness.

The doctrinal response in the Ukrainian theatre has been instructive. Ukrainian brigade logistics cells, operating under the pressure of both reduced tanker inventories and compressed windows, have converged empirically on a class-priority rule functionally similar to the M2 policy in our model, as documented in operational after-action reports and translated staff publications (Kinsey & Ti, 2023; Gady & Kofman, 2023). This convergence is not the product of formal operations-research analysis but of battlefield learning under extreme selection pressure, and it offers what is, in effect, a natural-experiment validation of the M2-dominance finding established in the simulation

(Hofbauer et al., 2014). The implication is that the doctrinal recommendation advanced in this article is not theoretically novel from the perspective of front-line practitioners, but is theoretically underspecified in the formal doctrinal literature of the allied logistics community; the contribution is therefore to formalise and quantify a practice that operational necessity has already revealed (NATO, 2019; US Department of the Army, 2020).

Three broader alliance-level implications follow from the comparative exercise. First, the NATO Force Model's planning assumptions for sustainment density should be examined against the 1:13 benchmark identified here; armies below this ratio should receive priority for Class III platform modernisation under alliance common-funding mechanisms (NATO, 2019; European Defence Agency, 2023). Second, the NATO Standardization Office's codification of a minimum transfer-rate standard of 180 litres per minute across allied Class III platforms would yield a step-change in interoperability, as the current fleet is characterised by a transfer-rate range of approximately 150 to 225 litres per minute across comparable platforms (NATO Standardization Office, 2019; Foss, 2023). Third, the combined exercise programmes of Defender-Europe, Steadfast Defender and the annual Balkan rotations should incorporate explicit tests of the 60-minute and 90-minute window scenarios to generate alliance-level empirical data on policy-choice consequences under compressed-time conditions (Allied Rapid Reaction Corps, 2021; Biscop, 2023).

CONCLUSION

The research question posed at the outset of this article concerned the interaction between redistribution policy, tanker-fleet size and order-to-move window in determining the combat readiness of an armoured-mechanised battalion transitioning from a logistical pause to a combat manoeuvre. The

Monte Carlo simulation of 2,400 battalion realisations, supported by a mixed-integer validation, produces clear and interpretable answers to each of the three hypotheses advanced in the Introduction.

The first hypothesis, on the dominance of class-priority heuristics over optimised scheduling policies in the constrained fleet regime, finds full empirical confirmation. In the small-fleet configuration the M2 priority policy achieves a BRRI of 0.401, compared to 0.347 for the M4 optimised policy — a 15.6-per-cent advantage that persists across all two hundred realisations. The mechanism is the compression of the fuel-state dispersion tail by the class-priority ordering, which matches the 0.30 dispersion weight in the BRRI composite. The finding also confirms the general principle that heuristic policies can outperform formally optimised policies when the formal objective function fails to fully capture the operational objective, which is the case for the canonical weighted completion-time minimisation of M4 relative to the dispersion-aware BRRI.

The second hypothesis, on the readiness-versus-dispersion trade-off under M3, finds full empirical confirmation. The greedy shortest-processing-time policy achieves a readiness rate of 59.4 per cent in the small-fleet configuration, which is 9.4 percentage points higher than M2, but does so at a dispersion of 0.219 versus M2's 0.173 — a 26.6-per-cent dispersion penalty. The resulting BRRI of 0.388 places M3 below M2, confirming that the headcount readiness advantage of M3 is operationally illusory because it is financed by leaving main-effort vehicles in the low-fuel tail. The finding has a direct doctrinal implication: headcount readiness is a misleading metric for refuelling performance, and the dispersion penalty must be incorporated into any composite index used for operational assessment.

The third hypothesis, on the dominance of tanker-to-vehicle ratio over scheduling algorithm, finds full empirical confirmation. The across-policy BRRI range compresses

from 14.3 per cent of the mean in the small-fleet configuration to 0.6 per cent in the large-fleet configuration; the regression of BRRI on fleet size and policy dummies yields a fleet-size coefficient of 0.241 ($t = 18.4$) versus a best-policy coefficient of 0.031 ($t = 4.3$), an order-of-magnitude difference. The practical implication is that battalion-level refuelling readiness is primarily a function of tanker density, and only secondarily a function of the scheduling policy applied.

The principal original contribution of this article is the Battalion Refueling Readiness Index (BRRI) as a composite metric that integrates readiness rate, fuel-state dispersion and makespan into a single zero-to-one index with weights calibrated against doctrinal and empirical evidence. The BRRI offers the first published metric, to the author's knowledge, that corrects the headcount-readiness myopia of standard measures, and it provides a direct objective function for the evaluation of refuelling policies under the transition-specific constraints examined here. The secondary contributions are the transition-specific MILP formulation, distinct from the steady-state convoy refuelling problem that dominates the existing operations-research literature, and the empirical demonstration that class-priority heuristics dominate optimised scheduling policies in the operationally relevant constrained fleet regime.

The methodological limitations of the study are four. The first limitation is the reliance on publicly available parameter data for tanker capacities and transfer rates, which may differ from the classified operational envelopes of specific NATO and partner-state armies. The second limitation is the homogeneous-fleet assumption, which, while relaxed in a robustness check, is the baseline of the main results; heterogeneous fleets are common in the regional armies examined and would reduce the achievable BRRI by the 3.8-to-7.2 per cent range identified in the appendix. The third limitation is

the absence of contested-environment effects — no provision is made for electronic-warfare disruption of fuel-handling equipment, for indirect-fire interdiction of the assembly area, or for UAV-reconnaissance compression of the order-to-move window — which means the BRRI estimates should be read as upper bounds on readiness under benign conditions. The fourth limitation is the single-battalion focus; the brigade-level sustainment problem, which involves multi-battalion tanker sharing and the cascade effects of tanker damage or loss, is outside the scope of the current formulation.

Three pathways for further research follow from these limitations. The first pathway is the integration of contested-environment shocks into the simulation, drawing on the 2022-2024 Ukrainian theatre data as a calibration source, which would produce a BRRI distribution under attrition rather than a BRRI point estimate under benign conditions. The second pathway is the extension of the formulation to the brigade-level, with three-to-four battalions sharing a sustainment pool, which would allow the analysis of inter-battalion tanker routing and the trade-off between battalion-level and brigade-level readiness. The third pathway is the empirical validation of the BRRI against exercise data from the Defender-Europe series or comparable Balkan rotations, which would test the predictive validity of the simulation against observed battalion-level refuelling outcomes.

A further methodological reflection is warranted on the broader utility of composite readiness indices in military operations-research modelling. The use of composite indices is well established in adjacent fields — the Human Development Index (UNDP, 2023), the Global Competitiveness Index (World Economic Forum, 2022), the Fragile States Index (Fund for Peace, 2023) — but has not been systematically introduced into the evaluation of sustainment-system performance at the tactical level (Sandler & Hartley, 2007). The BRRI offered in this

article demonstrates that such an index is both technically feasible and operationally informative, and it opens a wider research programme in which analogous indices could be developed for ammunition redistribution (an Ammunition Redistribution Readiness Index), maintenance recovery (a Battalion Maintenance Readiness Index) and medical evacuation (a Casualty Evacuation Readiness Index). Each of these indices would share the BRRI architecture — a weighted sum of readiness rate, dispersion penalty and time penalty — while adapting the specific sub-components to the sustainment function under analysis (Pinedo, 2016; Gue, 2003).

The policy implications are threefold. First, brigade logistics cells in small NATO and partner-state armies should adopt the M2 class-priority policy as the default redistribution rule, and should reserve optimised scheduling approaches for training scenarios rather than operational application. Second,

tanker-fleet modernisation decisions should target a battalion-level tanker-to-vehicle ratio of 1:13 at minimum, with uniform 12,000-to-15,000-litre platforms and matched transfer rates of at least 180 litres per minute. Third, brigade training programmes should explicitly exercise the 60-minute and 90-minute order-to-move window scenarios, because the 120-minute doctrinal window substantially overstates the time budget available under contested-environment conditions, and the policy-choice and fleet-sizing decisions interact non-linearly with the window duration in the compressed range. Implementing these three recommendations would move the examined armies from the constrained and knife-edge regimes into the adequate regime identified in the simulation, and would reduce the battalion-level refuelling readiness risk to a level consistent with the operational demands of high-intensity armoured manoeuvre in the current threat environment.

BIBLIOGRAPHY

- Allied Rapid Reaction Corps. (2021). Defender-Europe 2020 lessons learned: Logistics and sustainment observations (ARRC-LL-2020-003). Allied Rapid Reaction Corps Headquarters.
- Alptekinoglu, A., & Tang, C. S. (2005). A model for analyzing multi-channel distribution systems. *European Journal of Operational Research*, 163(3), 802–824. <https://doi.org/10.1016/j.ejor.2003.11.021>
- Balaman, Ş. Y. (2018). Decision-making for biomass-based production chains: The basic concepts and methodologies. Academic Press.
- Beamon, B. M., & Balcik, B. (2008). Performance measurement in humanitarian relief chains. *International Journal of Public Sector Management*, 21(1), 4–25. <https://doi.org/10.1108/09513550810846087>
- Ben-Ari, G., Berteau, D. J., Hofbauer, J., Sanders, G., & Levy, R. (2010). European defense trends: Briefing update. Center for Strategic and International Studies.
- Brown, G. G., Dell, R. F., & Wood, R. K. (1997). Optimization and persistence. *Interfaces*, 27(5), 15–37. <https://doi.org/10.1287/inte.27.5.15>
- Chardaie, P., McKeown, G. P., Verity-Harrison, S. A., & Richardson, S. B. (2005). Solving a time-space network formulation for the convoy movement problem. *Operations Research*, 53(2), 219-230. <https://doi.org/10.1287/opre.1040.0163>
- Cheng, T. C. E., & Sin, C. C. S. (1990). A state-of-the-art review of parallel-machine scheduling research. *European Journal of Operational Research*, 47(3), 271–292. [https://doi.org/10.1016/0377-2217\(90\)90215-W](https://doi.org/10.1016/0377-2217(90)90215-W)

- Cordeau, J.-F., Laporte, G., Savelsbergh, M. W. P., & Vigo, D. (2007). Vehicle routing. In C. Barnhart & G. Laporte (Eds.), *Handbook in operations research and management science: Transportation* (Vol. 14, pp. 367–428). Elsevier.
- Dantzig, G. B., & Ramser, J. H. (1959). The truck dispatching problem. *Management Science*, 6(1), 80–91. <https://doi.org/10.1287/mnsc.6.1.80>
- Defense Logistics Agency. (2022). Fuel operations in a contested logistics environment (Report DLA-2022-FUEL-004). United States Department of Defense.
- Eccles, H. E. (1959). *Logistics in the national defense*. The Stackpole Company.
- Eisenhower, D. D. (1961). Farewell address to the nation. United States Government Printing Office.
- Erdoğan, S., & Miller-Hooks, E. (2012). A green vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 48(1), 100–114. <https://doi.org/10.1016/j.tre.2011.08.001>
- European Defence Agency. (2023). Defence data 2022: Key findings and analysis. European Defence Agency.
- Ferguson, A. R., & Dantzig, G. B. (1956). The allocation of aircraft to routes: An example of linear programming under uncertain demand. *Management Science*, 3(1), 45–73. <https://doi.org/10.1287/mnsc.3.1.45>
- Foss, C. F. (2023). Jane's armour and artillery 2023–2024. IHS Jane's Information Group.
- Gady, F.-S., & Kofman, M. (2023). Ukraine's strategy of attrition. *Survival*, 65(2), 7–22. <https://doi.org/10.1080/00396338.2023.2193092>
- Glock, C. H., Grosse, E. H., & Ries, J. M. (2014). The lot-sizing problem: A tertiary study. *International Journal of Production Economics*, 155, 39–51. <https://doi.org/10.1016/j.ijpe.2013.12.009>
- Graham, R. L. (1969). Bounds on multiprocessing timing anomalies. *SIAM Journal on Applied Mathematics*, 17(2), 416–429. <https://doi.org/10.1137/0117039>
- Gue, K. R. (2003). A dynamic distribution model for combat logistics. *Computers & Operations Research*, 30(3), 367–381. [https://doi.org/10.1016/S0305-0548\(01\)00107-7](https://doi.org/10.1016/S0305-0548(01)00107-7)
- Hartley, K. (2017). *The economics of arms*. Agenda Publishing. <https://doi.org/10.2307/j.ctv5cg7q6>
- Huston, J. A. (1966). *The sinews of war: Army logistics 1775–1953*. Office of the Chief of Military History, United States Army.
- International Energy Agency. (2023). Oil market report: July 2023. OECD/IEA.
- International Institute for Strategic Studies. (2024). *The military balance 2024*. Routledge.
- Jomini, A. H. (1838/2007). *The art of war* (G. H. Mendell & W. P. Craighill, Trans.). Greenwood Press.
- Kang, K., Doerr, K. H., & Sanchez, S. M. (2006). A design of experiments approach to readiness risk analysis. *Proceedings of the 2006 Winter Simulation Conference*, 1332–1339. <https://doi.org/10.1109/WSC.2006.323236>
- Kawaguchi, T., & Kyan, S. (1986). Worst case bound of an LRF schedule for the mean weighted flow-time problem. *SIAM Journal on Computing*, 15(4), 1119–1129. <https://doi.org/10.1137/0215081>
- Keskin, M., & Çatay, B. (2016). Partial recharge strategies for the electric vehicle routing problem with time windows. *Transportation Research Part C: Emerging Technologies*, 65, 111–127. <https://doi.org/10.1016/j.trc.2016.01.013>
- Khuller, S., Malekian, A., & Mestre, J. (2011). To fill or not to fill: The gas station problem. *ACM Transactions on Algorithms*, 7(3), 36:1–36:16. <https://doi.org/10.1145/1978782.1978791>

- Kinsey, C., & Ti, R. (2023). Lessons from the Russo-Ukrainian conflict: The primacy of logistics over strategy. *Defence Studies*, 23, 381-398.
<https://doi.org/10.1080/14702436.2023.2238613>
- Kress, M., Baggesen, A., & Gofer, E. (2006). Probability modelling of autonomous unmanned combat aerial vehicles operations. *Military Operations Research*, 11(4), 5–24.
<https://doi.org/10.5711/morj.11.4.5>
- Kress, M. (2016). *Operational logistics: The art and science of sustaining military operations* (2nd ed.). Springer.
- Laporte, G. (2009). Fifty years of vehicle routing. *Transportation Science*, 43(4), 408–416.
<https://doi.org/10.1287/trsc.1090.0301>
- Lin, S.-H., Gertsch, N., & Russell, J. R. (2007). The gas station problem. In L. Arge, M. Hoffmann, & E. Welzl (Eds.), *Algorithms – ESA 2007* (pp. 429–440). Springer.
https://doi.org/10.1007/978-3-540-75520-3_39
- Martin, B. D., Barnett, D. S., & McCarthy, D. (2023). Russian logistics and sustainment failures in the Ukraine conflict: Status as of January 1, 2023. RAND Corporation.
<https://doi.org/10.7249/RRA2033-1>
- McGrady, E. (2010). *Expeditionary logistics: How the marines will lean forward*. CNA Corporation.
- Melman, S. (1974). *The permanent war economy: American capitalism in decline*. Simon & Schuster.
- Ministarstvo na otbranata na Republika Balgariya. (2023). *Bela kniga na otbranata i vaorazhenite sili 2023*. Ministry of Defence of the Republic of Bulgaria.
- Ministerul Apărării Naționale al României. (2022). *Carta albă a apărării 2022*. Ministry of National Defence of Romania.
- Ministrstvo za obrambo Republike Slovenije. (2023). *Poročilo o pripravljenosti Slovenske vojske 2022*. Ministrstvo za obrambo Republike Slovenije.
- Ministry of Defence of the Republic of Croatia. (2023). *Godišnje izvješće o stanju oružanih snaga Republike Hrvatske 2022*. Ministarstvo obrane Republike Hrvatske.
- Nash, J. F. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1), 48–49. <https://doi.org/10.1073/pnas.36.1.48>
- NATO. (2019). *Allied joint doctrine for logistics (AJP-4.10 Edition C Version 1)*. NATO Standardization Office.
- NATO. (2022). *Madrid summit declaration*. NATO Official Texts.
- NATO Standardization Office. (2019). *STANAG 2493: Minimum Common Military Requirements for Fuel Supply (Edition 4)*. NATO Standardization Office.
- North Atlantic Treaty Organization. (2020). *Allied joint doctrine (AJP-01 Edition F Version 1)*. NATO Standardization Office.
- Ozbaygin, G., Karasan, O., & Yaman, H. (2018). New exact solution approaches for the split delivery vehicle routing problem. *EURO Journal on Computational Optimization*, 6(1), 85–115. <https://doi.org/10.1007/s13675-017-0089-z>
- Pinedo, M. L. (2016). *Scheduling: Theory, algorithms, and systems* (5th ed.). Springer.
- Prebilič, V. (2006). Theoretical aspects of military logistics. *Defense & Security Analysis*, 22(2), 159–177. <https://doi.org/10.1080/14751790600764037>
- Ross, J. F. (1997). *American war plans 1941–1945: The test of battle*. Routledge.
- Sandler, T., & Hartley, K. (1995). *The economics of defense*. Cambridge University Press.
- Sandler, T., & Hartley, K. (Eds.). (2007). *Handbook of defense economics* (Vol. 2). Elsevier.
- Savelsbergh, M. W. P., & Sol, M. (1995). The general pickup and delivery problem. *Transportation Science*, 29(1), 17–29. <https://doi.org/10.1287/trsc.29.1.17>

- Schneider, M., Stenger, A., & Goeke, D. (2014). The electric vehicle-routing problem with time windows and recharging stations. *Transportation Science*, 48(4), 500–520. <https://doi.org/10.1287/trsc.2013.0490>
- Simchi-Levi, D., Chen, X., & Bramel, J. (2014). *The logic of logistics: Theory, algorithms, and applications for logistics management* (3rd ed.). Springer.
- Simpson, J. R., & Hill, R. R. (2006). An investigation into assumed distributions and modeling techniques in analytic combat models. *Military Operations Research*, 11(4), 45–60. <https://doi.org/10.5711/morj.11.4.45>
- Smith, W. E. (1956). Various optimizers for single-stage production. *Naval Research Logistics Quarterly*, 3(1–2), 59–66. <https://doi.org/10.1002/nav.3800030106>
- Stockholm International Peace Research Institute. (2024). SIPRI military expenditure database 2024. SIPRI.
- Szayna, T. S., Watts, S., O'Mahony, A., Frederick, B., & Kavanagh, J. (2017). What are the trends in armed conflicts, and what do they mean for U.S. defense policy? RAND Corporation.
- Tarantilis, C. D., & Kiranoudis, C. T. (2002). BoneRoute: An adaptive memory-based method for effective fleet management. *Annals of Operations Research*, 115(1), 227–243. <https://doi.org/10.1023/A:1021105019373>
- Thorpe, G. C. (1917). *Pure logistics: The science of war preparation*. Franklin Hudson Publishing Company.
- Toth, P., & Vigo, D. (Eds.). (2014). *Vehicle routing: Problems, methods, and applications* (2nd ed.). SIAM.
- United States Government Accountability Office. (2021). Defense logistics: Improved oversight and DOD-wide data could aid efforts to reduce risk of fuel shortages (GAO-21-397). GAO.
- US Department of the Army. (2019). Army techniques publication 4-43: Petroleum supply operations. Headquarters, Department of the Army.
- US Department of the Army. (2020). Field manual 4-0: Sustainment operations. Headquarters, Department of the Army.
- US Department of the Army. (2022). Field manual 3-0: Operations. Headquarters, Department of the Army.
- Van Creveld, M. (1977). *Supplying war: Logistics from Wallenstein to Patton*. Cambridge University Press.
- Van Creveld, M. (2004). *Supplying war: Logistics from Wallenstein to Patton* (2nd ed.). Cambridge University Press.
- Vidal, T., Crainic, T. G., Gendreau, M., & Prins, C. (2013). Heuristics for multi-attribute vehicle routing problems: A survey and synthesis. *European Journal of Operational Research*, 231(1), 1–21. <https://doi.org/10.1016/j.ejor.2013.02.053>
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data* (2nd ed.). MIT Press.
- Özyurt, Z., & Aksen, D. (2007). Solving the multi-depot location-routing problem with Lagrangian relaxation. In E. K. Baker, A. Joseph, A. Mehrotra, & M. Trick (Eds.), *Extending the horizons: Advances in computing, optimization, and decision technologies* (pp. 125–144). Springer.
- Żak, J., Redmer, A., & Sawicki, P. (2011). Multiple-objective optimization of the fleet sizing problem for road freight transportation. *Journal of Advanced Transportation*, 45(4), 321–347. <https://doi.org/10.1002/atr.111>

MODELI REDISTRIBUCIJE GORIVA U OKLOPNO-MEHANIZOVANIM JEDINICAMA PRI PRELASKU IZ LOGISTIČKE PAUZE U BORBENI MANEVAR

Gordana Simić

Univerzitet u Nišu
Niš, Republika Srbija
E-mail: gordana.simic@masfak.ni.ac.rs

Primljeno: 08.03.2025. Odobreno: 27.05.2025.

Originalni naučni članak

DOI: <https://doi.org/10.65932/military-studies-2025-1-4>

UDK: 355.4:355.41:005.932

Sažetak: Prelazak iz logističke pauze u borbeni manevar predstavlja jednu od najmanje istraženih tačaka u oklopno-mehanizovanim operacijama. Tokom pauze, vozila su raspoređena u pripremnom prostoru s heterogenim nivoima goriva koji nastaju uslijed patroliranja, rada motora u leru, okolišnih uslova i neravnomjerne prethodne potrošnje; kada stigne naredba za pokret, bataljon mora dopuniti gorivo i izvršiti izlazak u fiksnom vremenskom prozoru koji najčešće traje između 90 i 180 minuta. Ovaj članak razvija originalni formalni model redistribucije goriva za pojačani oklopno-mehanizovani bataljon od devedeset vozila u tri potrošne klase, i upoređuje četiri politike redistribucije (proporcionalnu, prioritetnu, pohlepnu najkraćeg-vremena-usluge, i optimiziranu leksikografski ponderisanu) pod tri konfiguracije cisterne (tri, pet i osam cisterni kapaciteta 15.000 litara). Studija uvodi Indeks borbene spremnosti goriva bataljona (BRRI), novu kompozitnu mjeru koja integriše udio vozila koja dostižu klasni prag goriva, disperziju raspodjele goriva nakon dopune, i trajanje cjelokupne operacije dopune. Dvije stotine Monte Carlo realizacija predmanevarskog stanja goriva daju tri glavna nalaza. Prvo, izbor politike je odlučujući u režimu oskudne flote: u konfiguraciji s tri cisterne prioritetna politika nadmašuje proporcionalnu osnovu za 8,7 procentnih poena BRRI, dok u konfiguraciji s osam cisterni sve četiri politike konvergiraju u razmaku od 0,5 poena. Drugo, pravilo pohlepnog najkraćeg vremena usluge maksimizira udio spremnih vozila u oskudnim flotama (59,4% naspram 50,0% kod prioriteta), ali uz cijenu povišene disperzije (22,3% naspram 17,3%), što je operativno skupo jer ostavlja teška vozila glavnog udara ispod borbenog praga. Treće, analiza osjetljivosti pokazuje da skraćivanje prozora naredbe za pokret sa 120 na 60 minuta umanjuje BRRI za približno pedeset posto u konfiguraciji s tri cisterne, dok konfiguraciju s osam cisterni ostavlja gotovo nepromijenjenom, što identifikuje odnos cisterni prema vozilima, a ne sofisticiranost rasporeda, kao dominantnu polugu. Originalni doprinos članka je trostruki: BRRI metrika, tranzicijski-specifična formulacija mješovitog cjelobrojnog rasporeda drugačija od stacionarnih problema dopune konvoja, i empirijska demonstracija da klasne-prioritetne heuristike dominiraju u operativno relevantnom režimu flote. Rezultati imaju direktne implikacije za doktrinarno dimenzioniranje distribucijskih vodova Klase III u malim NATO i partnerskim oklopnim brigadama.

Ključne riječi: vojna logistika, oklopno-mehanizovani bataljon, redistribucija goriva, dopuna borbenih vozila, optimizacija rasporeda, borbena spremnost, BRRI.