Australian National University
Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Together and Alone: Economic Viability of Batteries Powering the Western Mediterranean Cruise Corridor

CAMA Working Paper 92/2025
October 2026

Anthony Wiskich
Centre for Applied Macroeconomic Analysis, ANU
Swinburne University of Technology

Abstract

Batteries can power a large ship two ways: carried *together* onboard, or sent *alone* as speculative *battery vessels* (BVs) coupled to a host under way, powering it over a short segment, and recharging in port. Onboard storage is limited by heterogeneous charging opportunities; a BV is more flexible but carries high fixed costs. We assess both on the busy Western Mediterranean cruise corridor, subject to EU policy on port couplers (OPS) and emissions, using real schedules, prices and sea states to optimise pack sizes, charging times and BV routes for maximum fleet profit. At projected 2040 battery costs, onboard batteries are viable first, led in value order by the near-year-round hosts that exploit sunk OPS spare capacity, from an implied carbon price around €95/tCO_2e. BVs switch on above about €175/tCO_2e, below the implied all-in carbon price of about €260/tCO_2e from the current regulatory stack. Four BVs supplementing onboard packs then electrify about 50% of the fleet's at-sea energy against 36% without them, by serving hosts that leave the corridor seasonally and raising port-coupler utilisation. High BV utilisation requires traffic density or host route flexibility.

Keywords

shipping decarbonisation, battery-electric propulsion, shore power (OPS), technoeconomic optimisation, FuelEU Maritime, EU Emissions Trading System

JEL Classification

C61, L91, Q42, Q48, Q54

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

The Centre for Applied Macroeconomic Analysis in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Together and alone: economic viability of batteries powering the Western Mediterranean cruise corridor

Anthony Wiskich*

October 28, 2026

Abstract

Batteries can power a large ship two ways: carried *together* onboard, or sent *alone* as speculative *battery vessels* (BVs) coupled to a host under way, powering it over a short segment, and recharging in port. Onboard storage is limited by heterogeneous charging opportunities; a BV is more flexible but carries high fixed costs. We assess both on the busy Western Mediterranean cruise corridor, subject to EU policy on port couplers (OPS) and emissions, using real schedules, prices and sea states to optimise pack sizes, charging times and BV routes for maximum fleet profit. At projected 2040 battery costs, onboard batteries are viable first, led in value order by the near-year-round hosts that exploit sunk OPS spare capacity, from an implied carbon price around €95/tCO_2e. BVs switch on above about €175/tCO_2e, below the implied all-in carbon price of about €260/tCO_2e from the current regulatory stack. Four BVs supplementing onboard packs then electrify about 50 % of the fleet's at-sea energy against 36 % without them, by serving hosts that leave the corridor seasonally and raising port-coupler utilisation. High BV utilisation requires traffic density or host route flexibility.

Keywords: shipping decarbonisation; battery-electric propulsion; shore power (OPS); techno-economic optimisation; FuelEU Maritime; EU Emissions Trading System
JEL classification: C61; L91; Q42; Q48; Q54

1 Introduction

Shipping accounts for roughly 3% of global greenhouse-gas emissions (International Maritime Organization [IMO], 2020), and the IMO now targets net-zero emissions from international shipping around mid-century (IMO, 2023). The European Union has moved first with binding price signals: the extension of the EU Emissions Trading System to maritime transport and the FuelEU Maritime regulation together place, for the first time, a material and rising cost on the carbon intensity of fuel burned on intra-European voyages (European Union, 2023b, 2023c). Of the zero-emission options open to shipping, direct battery-electric propulsion is energy-efficient, converting stored electricity to propulsion with far lower losses than any fuel-synthesis route, and is already the incumbent-displacing technology at the low-energy, high-frequency end of the market, with the contested frontier now in the 100–400 km short-sea band (Kersey et al., 2022; Kistner et al., 2024; BloombergNEF, 2024).

*Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, ANU College of Asia & the Pacific, Australian National University, Australia; and School of Business, Law & Entrepreneurship, Swinburne University of Technology, Australia. Corresponding author: tony.wiskich@anu.edu.au.

The constraint on extending batteries to longer routes is not primarily energy density but the inability to recharge at sea. The capital tied up in stored battery energy is about three orders of magnitude greater than the commodity value of the same energy stored as fuel, so a battery is economic only if it is cycled frequently; a pack that sits idle through a long ocean leg cannot pay for itself. This confines pure battery operation to short, repetitive, port-anchored duty cycles — short-route ferries, harbour craft and inland vessels (Anwar et al., 2020; Perčić et al., 2020; Kanchiralla et al., 2025; Moon et al., 2025) — and means that any electrification of longer routes must be *infrastructure-mediated* — fixed offshore charging stations along busy corridors that a ship diverts to and recharges at (Li et al., 2025), mobile charging vessels that rendezvous with and recharge battery-electric ships (Ren et al., 2025), or a fleet of dedicated *battery vessels* (BVs) that power host ships *under way* for short segments (Wiskich, 2025). The BV is distinct from onboard packs (captive to one hull), fixed shore or offshore charging (energy without at-sea range), mobile charging vessels (which recharge an already battery-electric ship that pauses to receive power; Ren et al., 2025), battery swapping (which moves cells, not power, and needs standardised packs and dockside handling; Zhang et al., 2023) and tug-assist (which moves the ship, not its energy): it alone delivers grid electricity to a conventionally fuelled host *under way*.

Placing the battery on a dedicated vessel rather than on the host has several structural advantages: it decouples the battery capital from any single ship, so the same pack cycles across many hosts and is not stranded when a ship's route changes; it allows targeting the segments worth electrifying; and, because the BV is itself a mobile grid-storage asset, it can earn temporal and spatial arbitrage in the electricity market when not serving ships. A further advantage is systemic rather than private: BVs can increase the utilisation of shore-power connections that policy is now funding at cruise ports (hotel-scale onshore supply under AFIR — the EU Alternative Fuels Infrastructure Regulation; on the role of shore power more broadly see CE Delft, 2022). The enabling hardware is emerging: the first "e-bunker" battery vessel, designed to deliver up to 50 MW to ships around port, has been developed (ABB, 2024). The transfer of propulsion-scale power under way is speculative and left to companion work, but our analysis yields broader economic insights into *separating* the battery from the host — insights that could carry over to other configurations, such as swappable containerised packs.

Prior economic work on BVs has assessed an actual long-distance Asia–Europe container route, asking what future battery costs and energy densities would make at-sea power competitive (Wiskich, 2025). Many of its inputs are stylised — generic time-invariant electricity prices, constant fuel prices, and an assumed host wait time rather than real schedules. Using a detailed cost structure, the paper shows how electrification increases as battery prices fall and energy densities rise. A companion study takes the system-level view, coupling a growing battery-vessel fleet with renewable power systems in the large-scale regime where the fleet itself begins to move the electricity markets it charges from (Wiskich, 2026). Work specific to cruise ships has so far centred on *at-berth* decarbonisation — shore power (cold ironing) to cut emissions and meet EU port mandates while docked — and its port economics (CE Delft, 2022; Osipova & Carraro, 2023; Bullock et al., 2023; Frković et al., 2024), with onboard batteries largely confined to hotel-load peak-shaving and hybrid auxiliary duty rather than propulsion (He et al., 2024). A complementary strand examines *offshore* charging stations a ship visits to recharge (Li et al., 2025; Sruthy & Preetha, 2024). Across these strands two gaps remain. First, the at-sea *propulsion* electrification of large cruise ships — whether by onboard packs or a BV — has not been assessed against a real corridor's schedules, sea states and market prices: existing work stops either at the berth (shore power) or at fixed and rendezvous charging for already-electric ships, while the BV literature rests on stylised single-route inputs (Wiskich, 2025). Second, the two mobile-storage architectures — a battery carried *with* the host and a battery *sent to* it — have not been evaluated jointly, so whether they compete or complement, and the carbon

price at which each becomes viable, remain open. We address both gaps.

The current paper considers the Western Mediterranean cruise corridor, well suited to a semi-empirical test. The Mediterranean is the world's second-busiest cruise region after the Caribbean (CLIA, 2024), and much of that traffic runs through the Western Mediterranean, so host density is high; its adjacent legs are short overnight hops; cruise ships are fuel-electric with large hotel and propulsion loads, making at-sea battery integration tractable and the displaced fuel substantial (Transport & Environment, 2023); the EU regulatory stack gives a strong and growing price signal; and shore-power policy under AFIR is already bringing grid connections to the corridor's ports. We solve a system-wide dispatch-and-investment problem, treating electricity, fuel and carbon prices as exogenous. The inputs are realised hourly zonal prices, a real cruise schedule, reanalysed sea states (an operability gate excludes legs the BV could not safely serve in high seas), and the EU regulatory stack at present-day prices — the ETS at its spot price and FuelEU compliance valued at the prevailing pool price rather than the statutory penalty. Because the FuelEU credit is what the compliance market pays for the abatement, the headline does not rest on a forward regulatory projection; we read the fleet's returns as a present-day marginal-abatement curve. The model considers key cost structures and maximises profits of the system by optimising onboard and BV battery capacity, charge times and BV schedule.

While many features of the corridor are compatible with electrification, most cruise ships are not on the Western Mediterranean corridor year-round. Many reposition seasonally — such as to the Caribbean, Northern Europe or the Canaries — so their onboard batteries could sit idle off-corridor for months at a time. Reduced battery utilisation from this schedule heterogeneity is a binding constraint on fixed onboard electrification. We therefore consider a combined architecture: onboard batteries for all hosts, largely adopted by the handful of ships that remain on the corridor near-year-round, and one or more dedicated BVs that serve the ships with more heterogeneous itineraries. The result is a joint fleet where each component earns according to its comparative advantage.

The central result is that the two storage architectures are complements rather than substitutes — onboard packs and battery vessels switch on at different carbon prices and reach different parts of the fleet — a finding on when shared versus dedicated at-sea storage pays that carries beyond this corridor to any busy, schedule-constrained shipping lane. The contribution is threefold. First, we provide the first assessment of at-sea propulsion electrification for cruise ships — by onboard batteries or BVs — grounded in a real corridor's schedules, sea states and market prices. Second, via a marginal-return analysis of the combined fleet frontier, we identify how the BV lifts the fleet's profit peak and extends economically viable delivery beyond the onboard fleet's ceiling, the capital-reallocation mechanism that enables this, and the determining factors. Third, we set out the deployment and policy questions the results raise: whether the corridor's dominant operators — MSC and Costa — can each reach the scale a BV needs on their own or must both electrify to do so; how BVs bear on the utilisation of the publicly funded port couplers they share; and how the configuration of electricity pricing shapes both. The remainder of the paper sets out the method (§2), the data (§3), results and sensitivities (§4), and a discussion of mechanisms, limitations and policy (§5).

2 Method

We model a joint fleet of *onboard* batteries and dedicated *battery vessels* (BVs) on the Western Mediterranean cruise corridor, with pack sizes, charging schedules and port infrastructure build jointly optimised. The model takes prices, ship schedules and sea states as exogenous. The central

output is the whole fleet's *annual profit* at projected battery system costs in 2040 and today's regulatory stack (the ETS at spot, FuelEU at the compliance-pool price), with robustness of assumptions examined as sensitivities. The objective is long-run operating surplus alone; other operator goals — schedule robustness, passenger experience, crewing — are outside scope. The model assumes perfect foresight of schedules, electricity prices and sea states over the annual optimisation horizon; it is therefore a *planning* rather than an operational model, intended to estimate long-run economic potential rather than to produce an operational dispatch under uncertainty.

The method follows the order of this section: a discrete-vessel representation of the combined onboard-plus-BV fleet over the corridor (§2.1); the shore-power infrastructure and berth contention the fleet charges through (§2.2); the two battery assets — onboard packs and the dedicated BV — and their costs (§2.3); the dispatch and capital decisions (§2.4), in which BV routing is pre-computed as a geographically-chained sequence (§2.4.1) and a joint fleet LP then optimises charging schedules, energy deliveries, pack sizes and endogenous port build to maximise profit (§2.4.2); the at-sea revenue, built from realised fuel and carbon prices and the EU maritime regulatory stack (§2.5); and three extensions — grid arbitrage, coupler substitution for port OPS infrastructure, and an adjusted schedule allowing direction switching (§2.6). The market and schedule data follow in §3. Notation is collected in Table 6.

2.1 Corridor, hosts, and legs

The corridor spans thirteen charging hubs across the Western and Central Mediterranean. Five form the Ligurian–Tyrrhenian core — Barcelona (ES), Marseille (FR), Savona (IT-SAV), Genoa (IT-NORD) and Civitavecchia/Rome (IT-CSUD) — and are joined westward by the Iberian coast and Balearics (Valencia, Palma de Mallorca) and southward by the Tuscan, Campanian and Central-Med ports (Livorno, Naples, Cagliari, Palermo, Messina, Valletta), forming the hub set $\mathcal{R}$, indexed r, ρ. Each hub carries its own OPS installation, coupler count and endogenous port build, and sits in one of *seven* distinct day-ahead price zones (Table 4) — hubs sharing a zone face the *same* realised price series.

BVs can only power hosts along *adjacent* overnight and 2-night hops between charging hubs — from the short Ligurian infill (Marseille–Savona/Genoa–Livorno) through the classic ∼14 h crossings to the longer Iberian and Central-Med hauls. Two ports, Tarragona and La Spezia, enter only as one-directional *delivery-only* spurs, the BV recharging at an adjacent hub it passes en route. The direct circuit-closing returns a BV may serve, the island-broken returns and excluded ports it may not, and the full admission tests are set out in Appendix B. Figure 1 sketches the corridor.

Cruise ships are the only hosts considered — their fuel-electric architecture makes battery integration tractable, and corridor traffic is overwhelmingly cruise. The schedule is built from *real* port-call calendars (§3), preserving time-of-day, weekday and the strong weekend concentration of departures; the construction is detailed in Appendix B. Hourly departure and arrival times are modelled; the charging window at each origin is the ship's actual docked interval. We divide the host pool into two *size classes* — the large flagship class and the older, smaller MSC Lirica and Musica classes (Sinfonia, Orchestra, Musica, Magnifica, Opera; ∼65–95 kGT against the flagship's ∼215 kGT).

Figure 2 shows the 2025 annual schedule for the nineteen-ship corridor pool — the onboard-adopting core and the most active seasonal battery-vessel hosts. Corridor traffic is strongly summer-peaked — monthly corridor legs run two to three times higher in the May–October plateau than in midwinter — because most lines redeploy tonnage to and from the Caribbean and elsewhere outside the summer, while only Costa Toscana and MSC World Asia run the western Mediterranean all year.

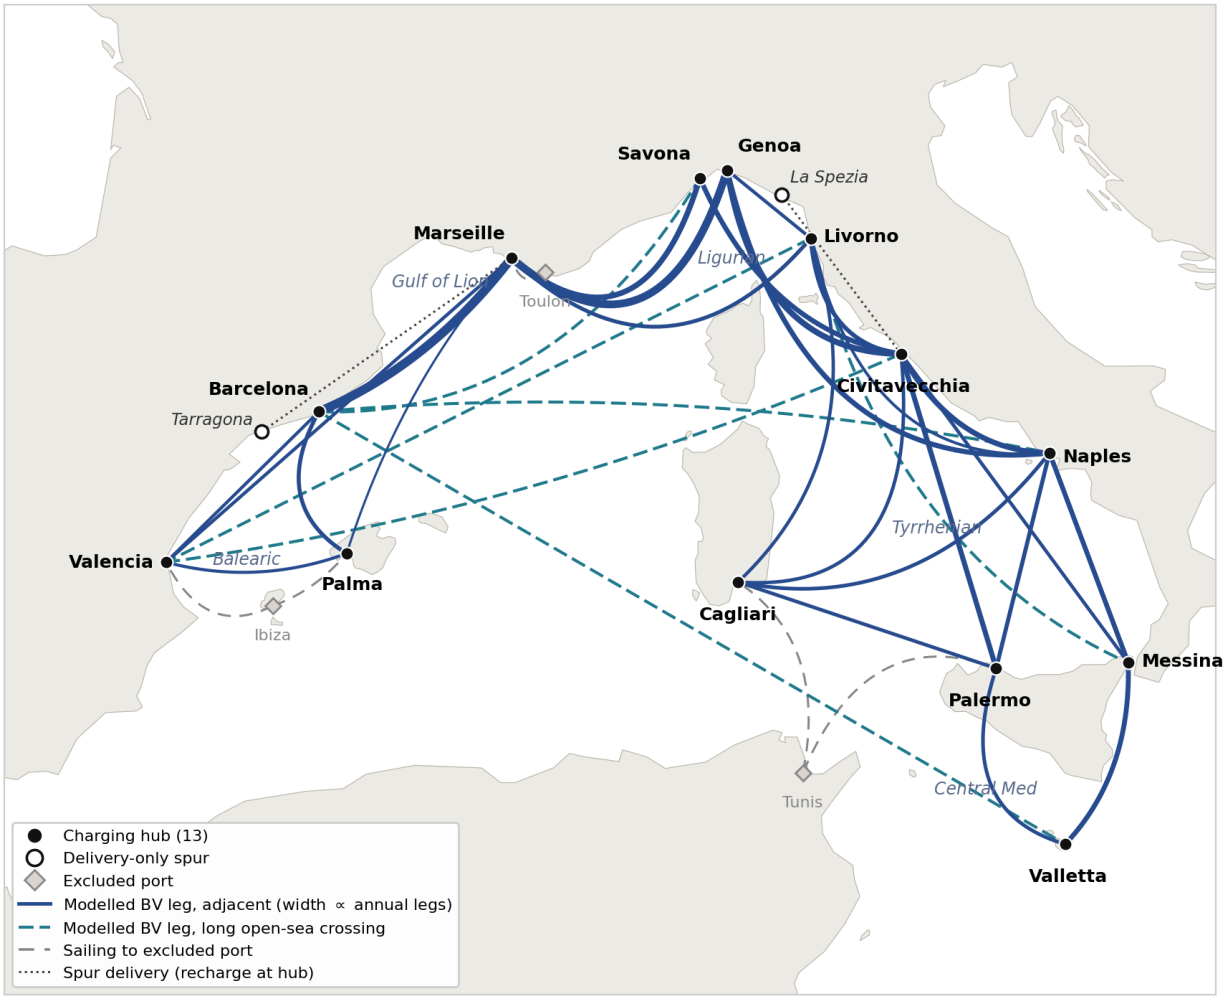


Figure 1: Western- and Central-Mediterranean cruise corridor as modelled. **Filled** nodes are the thirteen charging hubs; **solid** blue edges are the adjacent overnight legs a BV can serve, with host underway energy D_l sized from each leg's distance and duration; **edge width** is proportional to annual corridor traffic on that hub pair, ranging from the busy Barcelona–Marseille–Genoa spine down to the sparse Balearic and Central-Med links. **Teal long-dash** edges are long direct open-sea crossings actually sailed (non-adjacent hub pairs, e.g. Barcelona–Savona, Civitavecchia–Valencia): these are modelled as single BV legs too — genuine two-night passages up to 34 h underway. **Open** nodes (Tarragona, La Spezia) are delivery-only spurs: a BV serves them one-way but recharges at the adjacent hub it passes (Barcelona, Livorno). **Grey diamonds** are ports excluded from BV service, with **dashed grey** sailings to them.

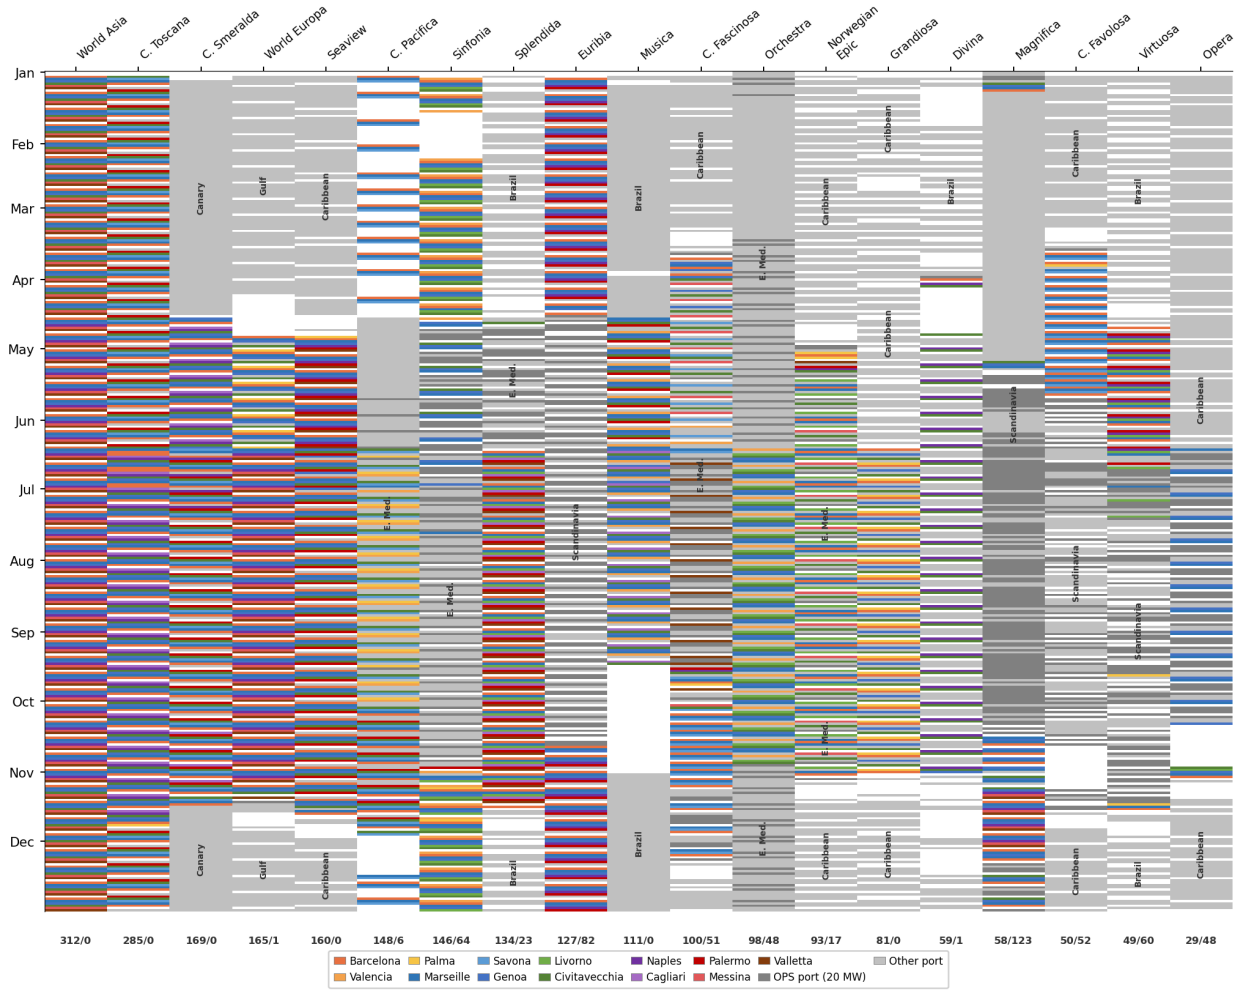


Figure 2: Annual port-call schedule for the 19-ship corridor pool, 2025. Each row is one calendar day (January at top); each column one ship. **Coloured cells**: stop at one of the thirteen charging hubs (legend, west to east). **Grey**: the ship is away from the corridor. *Darker grey* marks a call at an out-of-corridor port with an assumed 20 MW port coupler; *lighter grey* marks other non-hub calls and extended out-of-corridor sailings. Column totals (bottom) are annual corridor hub calls / out-of-corridor 20 MW OPS calls. The two year-round anchors (World Asia, Toscana) hold near-continuous corridor presence; the other ships are seasonal.

Each directed sailing is a *leg* $l \in \mathcal{L}$ characterised by an origin $r(l)$, a destination $\rho(l)$, integer departure and arrival periods $h_l^{\text{dep}}, h_l^{\text{arr}}$, a sea segment (used by the operability gate, §2.4.2), and a class-c host underway energy demand set by the ship's speed over the crossing:

$$D_l = \big(H_c + P_c^{\text{prop}}(v_l)\big)\, h_l^{\text{sea}}, \qquad P_c^{\text{prop}}(v) = P_{\text{max},c}^{\text{prop}}\Big(\tfrac{v}{V_{\max}}\Big)^3, \qquad v_l = \frac{\chi\, \text{gc}(r,\rho)}{h_l^{\text{sea}}}. \tag{1}$$

Here H_c (15 MW for the large host class c and 10 MW for the small) is the hotel (auxiliary) load carried whenever coupled; $P_c^{\text{prop}}(v)$ is propulsion power, a cube law in speed (IMO, 2020) anchored on the large class's installed propulsion ($P_{\text{max,large}}^{\text{prop}} = 44$ MW at $V_{\max} = 22.7$ kn); the class index c is fixed by the leg's ship, so D_l needs no separate subscript; v_l is the cruise speed inferred from the great-circle distance $\text{gc}(r,\rho)$ (coastal factor $\chi = 1.10$) over the coupled at-sea window h_l^{sea}; and h_l^{sea} is the scheduled arrival-to-departure gap less a manoeuvring buffer (Appendix B). Modelled adjacent hops span ∼110 MWh (a short, slow small-class Ligurian hop) up to ∼750 MWh (the long Genoa–Naples haul), with a mean of ∼390 MWh (median ∼360); the long direct circuit-closing returns run higher, up to ∼1800 MWh. Time is discretised into hourly periods, $\delta = 1$ h.

The smaller host class carries less hull resistance and hotel load: we scale D_l (hotel and propulsion) by $\frac{2}{3}$ — equivalently $H_{\text{small}} = 10$ MW, $P_{\text{max,small}}^{\text{prop}} \approx 29$ MW — which also leaves them more in-port charging headroom (§2.2). Because a BV's cost is a largely fixed daily charter whereas an onboard pack's is near-linear in capacity, this lower host energy raises the BV's cost per delivered MWh disproportionately, penalising the BV on precisely these smaller hosts.

2.2 Shore-power infrastructure

Both battery-asset types on the corridor — the onboard cruise-ship packs and the dedicated BVs — charge from the same sunk shore-power infrastructure. All corridor hubs are EU TEN-T ports subject to AFIR (European Union, 2023a), which from 2030 requires each to supply onshore power (OPS) of at least hotel-load scale ($R = 20$ MW per berth) to passenger ships at berth. We call each such 20 MW shore charger a *port coupler*, and set the number of port couplers at each hub to the smallest integer covering at least 90% of major-cruise port-stops — the coverage level AFIR targets — the smallest K_r for which $\sum_d \min(o_d, K_r) / \sum_d o_d \geq 0.90$, where d indexes calendar days and o_d is the day's peak simultaneous major-cruise occupancy, reconstructed from the PortMinder "docked-with" lists (major-cruise vessels only; small yacht and expedition traffic excluded). This gives $K_r = \{$ES 4, FR 4, Genoa 2, Savona 2, Civitavecchia 4, Naples 3, Livorno 3, Valencia 2, Palma 3, Palermo 2, Valletta 3, Messina 2, Cagliari 2$\}$. These counts represent congestion consistent with AFIR policy rather than actual charger numbers, since we model neither all cruise hosts nor any non-cruise traffic.

Additional charger capacity beyond the sunk level can be built endogenously at cost (§2.4.2), which also raises the port's total power capacity. Couplers and the grid behind them are assumed available whenever the schedule leaves them idle: we model neither forced outages, planned maintenance nor local grid constraints, so realised charging access is an upper bound.

2.3 The two battery assets

All hosts can carry onboard batteries. While docked at a corridor port, an onboard pack charges through the port coupler; the hotel load takes H_c and the remainder is available to the battery, subject to the port's total power limit. A pack's annual capacity cost is borne in full wherever the ship sails, including the long non-EU seasons when it sits idle. However, onboard batteries are credited for charging at out-of-corridor EU and EEA ports carrying the same ∼20 MW port coupler

— restricted to mainland, strong-grid (*Cat. 1*) ports and sized per call in Appendix B. The BVs are dedicated ships whose sole cargo is a large battery pack. Each BV plugs into a port coupler and can draw its full capacity, subject to the port's total power limit. The BV sails under its own power, spending self-propulsion energy calibrated in App. C.

Table 1 summarises the costs applying to these assets: the battery pack is common to both, while a BV alone carries its platform, host stern interfaces (applied to all hosts) and self-propulsion reserve, and the onboard pack alone carries the passenger-hull premium and the weight it hauls at sea. The parameter *derivations* and their sources are shown in Appendix C.

Table 1: Battery asset cost and operating parameters. The final column normalises each cost to a comparable annual rate: pack-proportional costs per *usable* MWh of storage and power-proportional costs per MW; the host stern interface is quoted at its 20 MW anchor (€50/kW above it), and the BV platform (hull + O&M) is a flat per-vessel cost (pack-independent; §2.3). Capital costs are the raw figure annualised over 20 yr (or 12 yr for the battery). Base-case values; derivations and sources in Appendix C.

Cost item	**Base value**	**≈€/yr**
Common to both assets		
Battery pack capital	\$87/kWh	11,300/MWh
Port coupler (charger + grid)	€650/kW chg + €200/kW grid	68,200/MW
Power-handling upgrade (host flank / BV)	€50/kW above 20 MW	∼4010/MW
BV only		
Vessel platform (hull + O&M)	\$20 M hull + \$3000/day O&M	∼2.5 M/BV
Host stern interface	€1.5 M per host (at 20 MW)	∼6020/MW at 20 MW
Self-propulsion reserve	∼6.5% of pack per leg = **9.9%** of host leg prop.†	—
Onboard only		
Passenger-hull premium	\$50/kWh nameplate	4610/MWh
At-sea weight penalty	11 t/MWh installed = **3.1%** of host leg prop.	—

† The self-propulsion reserve is a fraction of pack *capacity*: the vessel spends about 9.9 % of the host's leg propulsion moving itself, roughly 6.5 % of usable capacity at the profit-max pack (larger on the small hosts). Derived from BV mass in Appendix C.

2.4 Dispatch and capital decisions

The assessment proceeds in two steps (Fig. 3). The **first step** fixes the BV's route with a chain DP, eliminating all integer routing variables; this sets the BV's dock windows and the hosts it must fit with a stern interface. The **second step**, conditional on that fixed route, is a fleet-wide *linear programme* (LP) that sizes every battery pack and schedules charging, energy deliveries, port builds and upgrades for the combined onboard-plus-BV fleet simultaneously, *maximising annual profit*. Fleet composition is exogenous and varied across sensitivities. Solving routing and dispatch jointly would require a computationally intensive mixed-integer programme; we decompose the two because the route is governed primarily by geographic continuity and schedule feasibility — which ports chain into a sailable weekly loop — whereas the economic decisions the paper studies (charging timing, pack sizing and port build) form the continuous dispatch layer conditional on that route. Because the routing objective is served energy rather than profit, the decomposition can only *understate* the jointly-optimal result: the energy-maximising chain is one feasible chain among

many, so a fully price-aware joint optimum could do at least as well, and the reported profits are a lower bound on the joint solution.

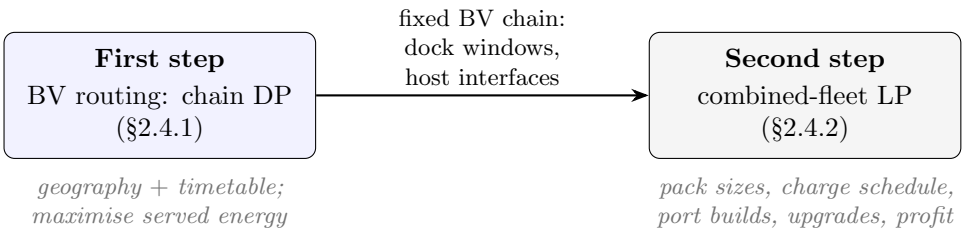


Figure 3: The two-step decomposition (§2.4): the routing DP fixes the BV's leg chain — dock windows and the hosts needing a stern interface — and the fleet-wide LP then sizes every pack and schedules charging, port builds and upgrades for profit, conditional on that route.

2.4.1 First step: battery-vessel routing

The routing menu is first filtered for operability. A leg is excluded from the BV menu $\mathcal{L}_{\mathrm{BV}}$ if significant wave height on its segment exceeds the operability threshold at any period of the transit,

$$l \notin \mathcal{L}_{\mathrm{BV}} \;\text{ if }\; \max_{h_l^{\mathrm{dep}} \le h \le h_l^{\mathrm{arr}}} H_{s,\,\mathrm{seg}(l),\,h} > H_s^{\star}, \tag{2}$$

with a base gate $H_s^{\star} = 4\,\mathrm{m}$ (sea state 5), the level up to which naval underway replenishment operates (United States Navy, 2007; Smith, 1992), and a conservative 2.5 m sensitivity (sea state 4, §4.2). This sea-state filter precedes the *geographically-non-teleporting chain DP*, which orders legs by departure and maximises served host energy subject to port continuity and timetable feasibility. For a single BV it is a maximum-weight path through the leg DAG ($O(n^2)$); for I vessels it is a min-cost flow that partitions the menu into I node-disjoint chains maximising their combined served energy (I units of flow, each leg split at unit capacity so no leg is served twice). Each leg is valued at $\min(D_l, \bar{E})$, the host energy one BV delivers in a single pass, capped at the pack's per-leg deliverable $\bar{E}$, so a small host leg is worth less. This weight is made *congestion-aware*: it is discounted by the fraction of the origin's berth window on which a coupler is left free by hosts, steering the chain toward ports where the delivered energy can actually be replenished rather than berths that are coupler-full through the call. The discount concentrates at the tight two-coupler hubs and is a routing signal only — the exact hourly charging availability is enforced in the LP (§2.4.2); its construction and the full admission tests are set out in Appendix H. The cap $\bar{E}$ is the most a BV can discharge to a host in one pass — the usable pack net of round-trip loss and the leg's self-propulsion draw, a per-leg deliverable rather than the pack itself. We carry a single representative value ($\approx 227\,\mathrm{MWh}$ at the base), iterated once to the pack sizes the LP builds and held fixed across sensitivities.

2.4.2 Second step: the combined-fleet LP

For the combined fleet in year y over hourly periods $h \in \{1, \ldots, T\}$ — an onboard battery pack on each of the J corridor ships (capacity endogenous, possibly zero) plus I dedicated battery vessels

(BVs) that deliver to a marked subset of those ships — dispatch maximises

$$
\begin{aligned}
\max\ V = {} & \underbrace{\sum_{m\in\mathcal{M}}\sum_{l\in\mathcal{L}_m}\hat{p}\,\mathbf{d}_{m,l}}_{\text{at-sea revenue}} - \underbrace{\sum_{m\in\mathcal{M}}\sum_{h\in\mathcal{H}_m}\mathbf{c}_{m,h}\,\pi_{z_m(h),h}\,\delta}_{\text{charging cost}} - \underbrace{\gamma_b\,K_b\sum_{m\in\mathcal{M}}\mathbf{E}_m}_{\text{battery capital}} \\
& - \underbrace{\gamma\Big(\sum_r C_r\,\mathbf{x}_r + \sum_{j\in\mathcal{J}}C^{\mathrm{upg}}\,\mathbf{u}_j + \sum_{i\in\mathcal{I}}C^{\mathrm{upg}}\,\mathbf{u}_i^{\mathrm{BV}} + \sum_{j\in\mathcal{J}}I_j^{\mathrm{iface}}\Big)}_{\text{annualised connection \& interface capital}} \\
& - \underbrace{\sum_{i\in\mathcal{I}}\big(\gamma\,K^{\mathrm{BV}}+F^{\mathrm{OM}}\big)}_{\text{BV platform (hull \& O\&M)}} - \underbrace{\gamma\,c^{\mathrm{prem}}\sum_{j\in\mathcal{J}}\mathbf{E}_j}_{\text{onboard premium}} \\
& + \underbrace{\mu\sum_{j\in\mathcal{J}}n_j\,\mathbf{o}_j}_{\text{out-of-corridor OPS credit}} - \underbrace{\kappa_{\mathrm{onb}}\sum_{j\in\mathcal{J}}\mathbf{E}_j\big(\hat{p}_{\mathrm{f}}\,\theta_j^{\Sigma}+\hat{p}_{\mathrm{c}}\,\theta_j^{c}\big)}_{\text{out-of-corridor weight-fuel}},
\end{aligned}
\tag{3}
$$

where $\mathcal{M}=\mathcal{J}\cup\mathcal{I}$ indexes all fleet members — $\mathcal{J}$ the onboard packs (one per corridor ship), $\mathcal{I}$ the BVs; $\mathcal{L}_m$ is the leg set of member m (pre-determined by chain DP for BVs, the ship's own corridor sailings for onboard ships); $\mathcal{H}_m$ the set of docked periods; $z_m(h)$ the bidding zone occupied at h; $\mathbf{d}_{m,l}\in[0,D_l]$ energy delivered to the host on leg l (MWh); $\hat{p}$ the avoided cost per delivered MWh (§2.5); $\mathbf{c}_{m,h}\geq 0$ the charge rate (MW); and $\mathbf{x}_r\geq 0$ additional charging capacity built at port r (MW). All symbols in Table 6.

Every cost is written as a *raw capital* sum annualised by a capital-recovery factor, so the objective is a consistent annual flow and the battery, interface, port and hull costs are treated alike. Connection and interface capital — the port coupler C_r (charger-plus-grid, per MW), the power-handling upgrades C^{upg} (the host flank interface onboard, the BV power build), and the host stern interface I_j^{iface} — share the 20 yr capital-recovery factor γ; the battery pack, with its shorter 12 yr life, carries its own γ_b (§2.3). The battery-capital term charges each installed MWh the annualised pack capital $\gamma_b K_b$; the port-coupler-build term charges only capacity built *above* the free-headroom baseline; the two power-handling-upgrade terms charge the peak power a vessel draws above the 20 MW coupler (onboard $\mathbf{u}_j$, BV $\mathbf{u}_i^{\mathrm{BV}}$); and the interface is summed over the *hosts* j a BV fits — a fixed fitting on the host ship's stern, shared by whichever BV meets it, scaled by its peak delivery power (App. C). The BV platform term is a single annual per-vessel flow combining the annualised hull γK^{BV} and the fixed operating cost F^{OM} (insurance, berth dues, unmanned maintenance). The onboard-premium term is the passenger-hull safety and space premium c^{prem} an unmanned BV avoids; the onboard weight penalty enters through the SOC balance (9). With every asset charged at its real capital cost, the objective V is the fleet's annual profit.

The last two terms handle out-of-corridor operation. *Credit:* each onboard pack also charges at the n_j mainland Cat. 1 EU/EEA ports ship j calls outside the corridor, delivering $\mathbf{o}_j\geq 0$ at each (calls are modelled identically), valued at μ — the onboard fleet's realised corridor margin, avoided value net of charging cost per delivered MWh. That per-call energy is bounded *both* by the port's charging headroom over the berth dwell and by the pack (exact caps in App. B), so a large pack cannot be filled in a short call and its credit is headroom-limited. *Weight-fuel:* the pack is dead weight on *every* leg the ship sails outside the corridor — the credited Cat. 1 calls, the lighter Cat. 2 calls, and legs wholly outside the EU — burning host fuel in proportion to that leg's propulsion. We charge this over the ship's full published itinerary: θ_j^{Σ} sums the anchor-normalised propulsion of all of ship j's out-of-corridor legs and carries the fuel cost $\hat{p}_{\mathrm{f}}$, while θ_j^{c} is the same sum weighted by

the EU maritime carbon scope (full intra-EU/EEA, half for a single EU/EEA endpoint, none fully outside) and carries the ETS-plus-FuelEU cost $\hat{p}_{\mathrm{c}}$ — the two additive components of the avoided cost (14) (App. B). Because μ depends on the solution it is held *fixed* within each solve — so the program stays linear — and driven to consistency by an outer iteration.

The base case does not allow power export to the grid — standard OPS couplers (IEC/IEEE 80005) are unidirectional, shore-to-ship — so grid arbitrage is modelled only as an extension (§2.6).

Shore-power contention and the coupler count. Let $\mathcal{O}(r,h)$ be the *dispatched* fleet members (BVs and onboard hosts, whose charging the LP controls) co-docked at port r in period h, and $o_{r,h}$ the number of major-cruise ships berthed there — *every* berthed major, the dispatched onboard hosts together with nine further ships whose packs and routes we do not model but whose berth-days we read from the same calendars (App. B), since a berthed major holds a coupler regardless. Each member's charge is capped individually by its own free headroom plus a share of any capacity built at the port,

$$\mathbf{c}_{m,h} \le P_m + \mathbf{x}_r, \qquad \text{(individual charge cap)} \tag{4}$$

where the per-member free headroom P_m takes one of two values: $P_m = P_{\mathrm{ob}} = R - H_c$ for an onboard ship (its coupler committed to the hotel load) and $P_m = P_{\mathrm{BV}} = R = 20\,\mathrm{MW}$ for a BV on an idle coupler. The built capacity $\mathbf{x}_r$ is a single *shared* port asset, so the co-docked members' *summed* draw satisfies:

$$\sum_{m\in\mathcal{O}(r,h)} \mathbf{c}_{m,h} \;\le\; \sum_{m\in\mathcal{O}(r,h)} P_m + \mathbf{x}_r, \qquad \text{(shared-build port cap)} \tag{5}$$

The individual cap (4) bounds any one vessel; (5) bounds their sum. Independently, each vessel's charge is limited by the rating of the grid tie it draws through, $\mathbf{c}_{m,h} \le P^{\mathrm{grid}} - H_m$ with $P^{\mathrm{grid}} = 50\,\mathrm{MW}$ (Table 6), capping opportunistic charging in the cheapest hours.

Free-coupler charging availability. An onboard host always holds its own coupler for the hotel load; a BV has no reserved berth and may charge only through couplers the hosts leave *free*. With $f_{r,h} = [\,K_r - o_{r,h}\,]^+$ couplers left free by the $o_{r,h}$ berthed hosts, the BVs co-docked at r share those free couplers — their *total* draw is bounded by the free-coupler power plus any capacity built at the hub, and when no coupler is free they cannot charge at all:

$$\sum_{i\in\mathcal{I}\cap\mathcal{O}(r,h)} \mathbf{c}_{i,h} \;\le\; f_{r,h}\,P_{\mathrm{BV}} + \mathbf{x}_r, \qquad \mathbf{c}_{i,h} = 0 \;\text{ when }\; f_{r,h} = 0. \tag{6}$$

Host flank interface and BV power build. A vessel that draws above its sunk 20 MW coupler needs extra drawing-side receiving equipment (switchgear and cabling above the AFIR baseline): the *host flank interface* where an onboard ship takes charge from the port coupler at berth, and the *BV power build* — the BV's own power electronics. For an onboard ship the LP sizes the flank interface from the peak charge above its free headroom P_{ob}; a BV pays the *same* 50 €/kW on the peak power it handles above 20 MW, taken as the larger of its peak charge and its peak at-sea discharge (delivered energy over the leg's underway hours):

$$\mathbf{u}_j \ge \mathbf{c}_{j,h} - P_{\mathrm{ob}} \qquad \forall\, h \in \mathcal{H}_j,\;\; j \in \mathcal{J}, \tag{7}$$

$$\mathbf{u}^{\mathrm{BV}}_i \ge \mathbf{c}_{i,h} - P_{\mathrm{BV}}, \qquad \mathbf{u}^{\mathrm{BV}}_i \ge \frac{\mathbf{d}_{i,l}}{h^{\mathrm{sea}}_l} - P_{\mathrm{BV}} \qquad \forall\, h \in \mathcal{H}_i,\; l \in \mathcal{L}_i,\;\; i \in \mathcal{I}. \tag{8}$$

Both peaks are charged the same power-handling-upgrade capital $C^{\mathrm{upg}} = 50\,€/\mathrm{kW}$ in the objective (3) (Green Cruise Port hotel-connection anchor, annualised at γ). This drawing-side equipment is distinct from and additive to the host stern interface I_j^{iface} a BV requires on each host (App. C); the BV's own coupler and converter are covered to 20 MW within its platform cost, so only power above that anchor is upgraded — symmetric to the onboard ship's flank interface.

Battery dynamics and endogenous pack. For each member m, state of charge $\mathbf{s}_{m,h} \geq 0$ evolves as

$$\mathbf{s}_{m,h} = \mathbf{s}_{m,h-1} + \sqrt{\eta_b}\,\mathbf{c}_{m,h}\,\delta - \sum_{l \in \mathcal{L}_m:\, h_l^{\mathrm{dep}} = h} \Big(\frac{\mathbf{d}_{m,l}}{\sqrt{\eta_b}} + \kappa_m\,\theta_l\,\mathbf{E}_m \Big), \quad \text{SOC balance} \tag{9}$$

$$0 \leq \mathbf{s}_{m,h} \leq \mathbf{E}_m, \quad \text{SOC within pack} \tag{10}$$

$$\mathbf{s}_{m,0} = \mathbf{s}_{m,H}, \quad \text{cyclic (periodic) boundary} \tag{11}$$

with $\kappa_m = \kappa_{\mathrm{BV}}$ for BVs (self-propulsion to follow the host) and $\kappa_m = \kappa_{\mathrm{onb}}$ for onboard ships (the small weight-of-pack propulsion penalty). The per-leg factor θ_l is the leg's host propulsion energy relative to a reference (anchor) crossing, so a long, fast haul costs more self-propulsion than a short, slow hop, with the small/large class distinction already embedded through that host-class propulsion. The draw is proportional to installed capacity $\mathbf{E}_m$ because self-propulsion is a cost of *sailing* the leg with a given hull mass — which its battery dominates — incurred regardless of how much it delivers. Its value is derived from BV mass by displacement$^{0.7}$ scaling (App. C): $\kappa_{\mathrm{BV}} \approx 0.065$ of usable capacity at the profit-max pack, about 9.9 % of the host's leg propulsion, held constant as a size-robust linearisation, with a doubled $\kappa = 0.13$ (a fuller, speed-mismatched barge hull) carried as a sensitivity (§4.2).

At-sea delivery draws $\mathbf{d}_{m,l}/\sqrt{\eta_b}$ from the pack (symmetric round-trip loss); charging adds $\sqrt{\eta_b}\,\mathbf{c}_{m,h}\,\delta$. Because $\mathbf{E}_m$ is a continuous LP variable, both the SOC cap (10) and the self-propulsion draw in (9) are linear. Each pack's state of charge is periodic over the year (11): the charge it carries into the first hour equals the charge it holds at the last, the standard cyclic boundary for annual energy-system models. The non-negativity in (10) ensures a pack never delivers more than it has charged. At the optimum the endogenous $\mathbf{E}_m$ equals the pack's peak state of charge (the rent term prices no more capacity than the schedule uses). We add no operational reserve floor above this non-negativity: $\mathbf{E}_m$ is already *usable* capacity, levelised over a depth-of-discharge of 0.80 so the modelled band sits inside the cell's physical range, and neither vessel's delivery is firm — a pack that runs low simply stops exporting and the host falls back on shore power or its own engines — so there is no firm load for a reserve to protect.

Corridor delivery target. For some simulations, a fleet-wide energy delivery target D is imposed:

$$\sum_{m \in \mathcal{M}} \sum_{l \in \mathcal{L}_m} \mathbf{d}_{m,l} = D. \tag{12}$$

Mutual exclusion. When the fleet contains multiple BVs, the same physical host-leg cannot be doubly served:

$$\sum_{i \in \mathcal{I}:\, l \in \mathcal{L}_i} \mathbf{d}_{i,l} \leq D_l \qquad \forall\, l. \tag{13}$$

2.5 At-sea revenue: avoided fuel and the EU regulatory stack

The avoided cost $\hat{p}$ (€per delivered MWh) is what the host avoids by taking battery power instead of running its own generators. It is a single corridor scalar, not leg-specific: the displaced fuel and its price are annual constants (§3), the EU ETS allowance and FuelEU compliance prices are corridor-wide market values, and battery energy is credited at zero grid carbon under FuelEU's shore-power zero-rating (below), so every delivered MWh carries the same avoided cost. We model the corridor fleet as liquefied-natural-gas (LNG) dual-fuel — the architecture of recent Western-Mediterranean cruise newbuilds — so the displaced fuel is LNG, with the existing marine-gas-oil (MGO) and LNG fleet reported as a sensitivity. LNG's bunker price and combustion CO_2 sit below MGO's, while once Otto-cycle methane slip is counted (Pavlenko et al., 2020) its well-to-wake intensity is barely better, leaving the dominant FuelEU term almost unchanged. The avoided cost has four additive components: a fuel cost, a cap-and-trade charge on *tank-to-wake* CO_2e (combustion plus methane slip), an *intensity*-standard (FuelEU) compliance cost, and the variable genset maintenance a displaced MWh of generator running saves.

$$\hat{p} = \underbrace{\sigma\, p_{\text{fuel}}}_{\text{fuel}} + \underbrace{\sigma\, \varepsilon\, \phi\, p_{\text{EUA}}}_{\text{EU ETS}} + \underbrace{m_{\text{FuelEU}}}_{\text{FuelEU Maritime}} + \underbrace{m_{\text{o\&m}}}_{\text{avoided O\&M}} \quad . \tag{14}$$

Here σ is the host's specific fuel consumption (t fuel per MWh electric), p_{fuel} the Mediterranean bunker price of that fuel, ε its tank-to-wake CO_2-equivalent emission factor (tCO_2e per t fuel, combustion plus methane slip),[1] ϕ the EU ETS phase-in fraction (full scope for the intra-EU corridor from 2026), p_{EUA} the EU allowance price, m_{FuelEU} the marginal FuelEU compliance value, and $m_{\text{o\&m}} \approx$ €9.2/MWh the avoided variable genset maintenance (lube oil, wear parts, overhaul-interval consumption — about a tenth of the fuel line). The FuelEU term is built explicitly as

$$m_{\text{FuelEU}} = 10^{-6}\, p^{\text{pool}} \underbrace{\sigma\, \iota_{\text{fuel}}\, L_{\text{fuel}}}_{\text{host WtW, gCO}_2\text{e/MWh}} \quad , \tag{15}$$

where p^{pool} (€/tCO_2e) is the price at which FuelEU compliance surplus clears in the pooling market and $\sigma\, \iota_{\text{fuel}}\, L_{\text{fuel}}$ is the host genset's well-to-wake intensity per delivered MWh — the emissions a delivered MWh displaces. The marginal value is the point: in the pooling market a delivered zero-carbon MWh earns a surplus credit equal to the *full* intensity it displaces, valued at p^{pool}. The mechanism is symmetric across the host's compliance position: a host in surplus sells the freed intensity into the pool at p^{pool}, while a host in deficit avoids buying an equal quantity of compliance at the same price — so the marginal MWh is worth p^{pool} times its full displaced intensity either way, provided the host (or its pool) transacts at that clearing price. The credit does *not* scale with the fleet's aggregate deficit fraction f^{def} — that fraction determines how tight the pool is, and hence the price level p^{pool}, but drops out of the per-MWh credit at the margin. Because the corridor is a small share of the EU fleet it is a price-taker in this pool, and its own electrification does not move p^{pool}; the pool level itself is the robustness axis we sweep (§4.2). We take $p^{\text{pool}} \approx$ €200/tCO_2e, an indicative marginal-compliance level on the OceanScore pooling index, which ran ∼€195–220/tCO_2e over the 2025 compliance year (OceanScore, 2026); this sits well below the statutory penalty ceiling (∼€642/tCO_2e) because compliance surplus — much of it generated at near-zero cost by LNG

[1] From 2026 the EU ETS extends beyond CO_2 to ship CH_4 and N_2O on a tank-to-wake basis. Under the full-scope regime we adopt we therefore price the LNG host's Otto-cycle methane slip alongside combustion CO_2: at the FuelEU Annex I default slip fraction (3.1 % of fuel mass for a medium-speed dual-fuel Otto engine; Pavlenko et al., 2020) and $GWP_{100} = 28$, ε rises from 2.75 (combustion) to 3.62 tCO_2e/t, adding ∼€10/MWh to the ETS charge. This is the same slip already carried in the FuelEU well-to-wake term, so nothing is double-counted; we price only tank-to-wake emissions, excluding the upstream well-to-tank leakage the ETS does not cover.

and LPG carriers — is currently in oversupply. FuelEU scores a ship on the GHG intensity of the energy it uses on board, defined to include energy used for propulsion at sea and from substitute sources such as electricity (European Union, 2023b): a battery-delivered MWh stays in that base and displaces fuel, so the numerator falls by the full host intensity, and the regulation expressly requires alternative propulsion supplied by "direct transfer as propulsion" to be recognised in the compliance balance. We take the delivered electricity's well-to-tank carbon as zero, extending to at-sea supply the treatment FuelEU already gives shore power at berth (European Union, 2023b; European Commission, 2024); the charging grid's *physical* carbon is priced separately, through the wholesale power price the battery pays.

Valuation at present-day prices. We deliberately value the stack at recent, observable carbon prices — realised 2025 markets and current spot and pool prints — not at a projected future stack. The ETS enters at the spot EUA price (∼€74/t); FuelEU enters at the current pooling price rather than the penalty; and, as p^{pool} is itself a present-day market print, pairing it with the spot ETS is a consistent same-vintage snapshot. The headline is therefore a *present-day* marginal-abatement statement that does not rest on the regulatory stack tightening: the FuelEU limit-reduction schedule shapes how the pool price will evolve, not the per-MWh credit, so the case is near-flat across the 2035–2045 horizon and turns on the pool and ETS prices rather than on the year. The widely projected rise in the ETS price — from ∼€74/t today to ∼€130/t by 2040 (Enerdata, 2025) — and a tighter FuelEU pool as the intensity target bites are carried as *upside* sensitivities (§4.1); a wide FuelEU pool-price range (€120–300/tCO_2e) is the FuelEU robustness axis.

Reading the stack as an implied carbon price. Because the fleet is single-fuel and the charging grid is credited at zero carbon, the avoided cost splits linearly into the displaced LNG fuel cost (∼€100/MWh) and the policy-plus-maintenance stack (∼€200/MWh), and each policy component is linear in the fuel a delivered MWh displaces. The ETS-plus-FuelEU charge is therefore *exactly* equivalent to a carbon price: dividing it by the delivered MWh's well-to-wake carbon content (0.73 tCO_2e/MWh) gives ∼€260/tCO_2e at today's prices.[2] More generally, netting the displaced fuel cost off *any* per-MWh storage cost and dividing by the same carbon content converts it to the carbon price that would rationalise that delivery — the device we use to read the profit-optimal marginal-abatement curve of §4.1 as an implied carbon price.

Table 2: Regulatory and market inputs to the avoided-cost stack. Base = prices observable today; the ETS path and the FuelEU pool range are carried as sensitivities (§4.1). The implied carbon price expresses the combined ETS+FuelEU burden per tonne of avoided well-to-wake CO_2e (LNG host).

	Base (today)	**Sensitivity range**
EU ETS allowance price (€/t)	∼**74**	74–130 (2040, Enerdata)
FuelEU pool price (€/tCO_2e)	∼**200**	120–300
Implied all-policy carbon price (€/tCO_2e, WtW)	∼**260**	—

[2]Expressed per tonne well-to-wake CO_2e: the ETS enters on tank-to-wake CO_2e (the €74/t EUA on the host's 0.60 tCO_2e/MWh tank-to-wake content) but is divided by the larger WtW content, so on this basis it appears as ∼€61/t rather than €74 — the metric is WtW throughout, by construction.

2.6 Extensions

Three optional extensions enter the results as upside cases, each a modification to the base model collected here: grid arbitrage (exporting stored energy in idle dock hours), coupler substitution (the BV standing in for a short shore-power coupler), and the adjusted schedule (direction switching).

2.6.1 Grid arbitrage

When grid export is enabled, each member $m \in \mathcal{M}$ may discharge to the grid at rate $\mathbf{g}_{m,h} \geq 0$ (MW) during dock periods, through a dedicated bidirectional shore connection rather than the unidirectional OPS coupler. Three modifications apply to the base LP. First, the objective (3) gains an export revenue term:

$$V \;\leftarrow\; V \;+\; \sum_{m \in \mathcal{M}} \sum_{h \in \mathcal{H}_m} \mathbf{g}_{m,h}\, \pi_{z_m(h),h}\, \delta. \tag{16}$$

Second, the SOC balance (9) gains an export drain with symmetric round-trip loss:

$$\mathbf{s}_{m,h} \;\leftarrow\; \mathbf{s}_{m,h} \;-\; \frac{\mathbf{g}_{m,h}\, \delta}{\sqrt{\eta_b}}. \tag{17}$$

Third, the combined charge and export power is bounded by the connection capacity:

$$\mathbf{c}_{m,h} + \mathbf{g}_{m,h} \;\leq\; P_m + \mathbf{x}_{z_m(h)}. \tag{18}$$

The symmetric $\sqrt{\eta_b}$ loss on both charge and discharge means that charging then immediately exporting retains only $\eta_b < 1$ of the energy, so the LP arbitrages only when the price spread exceeds the round-trip loss. Both onboard packs and the BV may export in this sensitivity: a standard OPS inlet is unidirectional (IEC/IEEE 80005), so export presumes the added bidirectional connection, its power bounded by each member's free headroom P_m plus any build and priced through the same connection capital as charging.

2.6.2 Coupler substitution: the BV as peaking shore power

By standing in for a marginally-used coupler, a BV lets a port build *fewer* couplers, avoiding shore-power capital.

At-berth supply with one fewer coupler. At a selected port we remove one coupler, leaving $K_r - 1$. On any hour when the remaining couplers are all busy ($o_{r,h} > K_r - 1$) an arriving ship would run its auxiliary generators; a BV standing in supplies that one unserved host's hotel load through a mobile cable, and it must do so for the host's *entire* berth — coverage is all-or-nothing, never partial (a bumped host is served on every one of its all-couplers-busy hours or on none):

$$\mathbf{G}_{r,h} = H_m\, \beta_b, \qquad h \in \mathcal{H}_b, \;\; \beta_b \in \{0, 1\}, \tag{19}$$

where $\mathcal{H}_b = \{h : o_{r,h} > K_r - 1\}$ collects the all-couplers-busy hours of a bumped host's berth b, $\beta_b = 1$ marks the BV standing in for it (feasible only if the BV is docked through all of $\mathcal{H}_b$ with charge to spare), and H_m is the host's class hotel load — a single berth, since one BV serves one host at a time. This energy draws on the pack; the fleet covers a berth only where doing so pays, with no forced obligation. The natural target is the port whose all-couplers-busy hours are small and most often coincide with the hours a BV is docked there; we screen the corridor and re-solve each candidate hub with its marginal coupler removed (App. E), where **Livorno** and **Naples** qualify.

Value of the removed coupler. Removing the coupler avoids its annualised capital and, against that, forgoes the electrification of any bumped host the BV cannot cover:

$$V \leftarrow V + \underbrace{\gamma\,(c_{\text{chg}} + c_{\text{grid}})\,R}_{\text{avoided coupler capital}} - \underbrace{\sum_{b}(1-\beta_b)\sum_{h\in\mathcal{H}_b}\hat{p}\,H_m\,\delta}_{\text{bumped-host fallback to generators}} \quad . \tag{20}$$

The second term nets the electrification the removed coupler would have delivered on the all-couplers-busy hours the BV cannot cover ($\beta_b = 0$) — those host calls fall back on generators, valued at the same avoided margin $\hat{p}$ (definitions and the per-port sum in App. E). Only ports whose served berth-days stay above the 90 % coverage floor with the coupler removed are retained, so the network still meets the same coverage standard.

2.6.3 Adjusted schedule (direction switching)

As an upside case we let *any* marked BV host run *any* of its circuits in reverse. A circuit is a full home-port-to-home-port loop; reversing the *whole* loop relocates its corridor hops to the other end of the week and flips their direction, opening return-direction legs a BV can slot into what would otherwise be idle time. The directions are chosen to **maximise the joint fleet's delivered host energy** by a hill-climb from the all-forward schedule; because the reversal that best feeds several vessels differs from the one best for a single chain, the chosen schedule depends on the battery-vessel count. The construction, the origin-admission rule and the resulting per-host switches are set out in Appendix B (Table 9); the operational challenges of reversing a published itinerary are taken up in the discussion (§5).

3 Data

The quantities in this section — electricity, fuel and allowance prices, the ship schedules and port calls, and the sea states — are *measured*: realised market and reanalysis data on which nothing is assumed. The *market* inputs are electricity, marine bunker fuel (LNG in the base case, MGO as a sensitivity) and EU allowance (EUA) prices. We run the model on the **2025** calendar year, and report 2023–2025 in Table 4. We choose 2025 because it is the *latest* year and so best reflects the ongoing clean-energy transition (solar capacity is expanding fast), making it the most representative of the near-future market that matters here. The qualitative case is robust to the year: what the optimisation exploits is the *relative* temporal and spatial structure of prices — the depth of the daily and cross-zone troughs a roaming pack charges into — and that structure is similar across years, so the dispatch and the case for a well-utilised pack do not depend on the absolute price level.

Corridor sea states hardly differ across the years: the operability gate (§2.1) blocks only about 0.4 % of corridor legs at the base 4 m wave limit and 5 % at the stricter 2.5 m sensitivity in 2025, with broadly similar shares in 2024 and 2023 (Table 4). The cruise schedule is built from the fleet's real, forward-looking port-call calendars — the coming year's published itineraries (PortMinder lists only forward departures; §2.1, Appendix B). The regulatory inputs are valued at present-day prices — spot EU ETS and FuelEU at the pool price — with that valuation, its one-sided ETS conservatism, and the ETS-price upside set out in §2.5. Table 4 collects the headline annual figures and Table 3 the inputs, their sources and treatment; the charging-price and zero-intensity-credit discussions are in Appendix F.

Table 3: Data inputs, sources and treatment. The charging-price and zero-intensity-credit discussions are in Appendix F.

Input	**Source**	**Resolution / treatment**
Electricity price $\pi_{z,h}$	ENTSO-E Transparency Platform	day-ahead hourly, seven bidding zones, 2025; clean UTC index, rare gaps filled
Marine fuel p_{fuel} (LNG base, existing-fuel sens.)	Dutch TTF (Protergia); Ship & Bunker G20	LNG: TTF annual mean $\times$ 13.64 MWh/t + €110/t small-scale bunker premium; MGO (existing-fuel sens.): G20 index at annual EUR/USD
EU allowance p_{EUA}	ICAP; Ember Carbon Price Viewer (EU ETS front-December contract)	realised 2025 annual mean, spot basis (€74; €130 by 2040 as sensitivity)
Cruise schedule	PortMinder iCalendar	28 ships (nineteen make regular corridor sailings); real calls $\rightarrow$ legs (App. B)
Significant wave height H_s	Open-Meteo Marine reanalysis	hourly, one open-water point per sea segment
Regulatory stack (EU ETS, FuelEU)	European Union (2023b, c)	valued present-day: spot EU ETS allowance plus FuelEU pool price (not the statutory penalty); mechanism in §2.5

Table 4: Annual data summary, 2023–2025. Electricity prices are realised day-ahead zonal means with the within-year hourly standard deviation in parentheses (ENTSO-E); LNG is the Dutch TTF annual average plus a small-scale bunkering premium and EUA the realised 2025 annual mean (ICAP; Ember, spot basis). The seven bidding zones cover the thirteen corridor hubs; hubs without their own zone share one (Savona and Naples price at IT-NORD and IT-CSUD, Valencia and Palma at ES, Valletta and Messina at IT-SICI). Sea-state blocked counts are corridor legs whose transit window exceeds the wave limit anywhere (Open-Meteo; §2.1).

	2023	**2024**	**2025**
Day-ahead electricity, zonal mean (SD) (€/MWh)			
ES (Barcelona)	87 (41)	63 (46)	65 (48)
FR (Marseille)	97 (46)	58 (41)	61 (46)
IT-NORD (Genoa)	128 (36)	107 (28)	116 (29)
IT-CSUD (Civitav.)	126 (37)	110 (32)	116 (30)
IT-CNOR (Livorno)	128 (36)	109 (31)	117 (30)
IT-SICI (Palermo)	126 (39)	112 (36)	116 (33)
IT-SARD (Cagliari)	123 (42)	106 (38)	112 (37)
LNG (€/t, TTF+premium)	772	571	615
EUA (€/tCO_2)	84	65	74
Sea-state gate — corridor legs blocked (% of total)			
$H_s > 4.0$ m (base)	12 (0.7%)	8 (0.4%)	7 (0.4%)
$H_s > 2.5$ m (sens.)	122 (6.7%)	100 (5.6%)	91 (5.0%)

4 Results

We discuss our base results and then sensitivities.

4.1 Base electrification frontier

Figure 4a plots optimal profit against net annual corridor delivered energy targets for the full 19-ship corridor fleet. The total at-sea leg demand is 705 GWh/yr in 2025. We start by considering the onboard case alone; profit rises to a peak of **+€13.4M/yr** at about 36 % of that potential (254 GWh/yr). Adoption follows value order, led by the two year-round corridor ships Costa Toscana and MSC World Asia. As delivery is forced to rise, profit falls because more *seasonal* ships, whose packs idle while the host leaves the EU, are adopted.

Battery vessels lift and extend that frontier from about 18 % of potential, and *how many* to field is itself a scale decision. The optimal BV count steps up with delivery: onboard-only at the smallest scales, then one, two, and three vessels through the mid-range, reaching **four** at the overall peak; a fifth only past it. Fielding four BVs lifts the profit-maximising envelope's peak to about half the corridor's electrifiable potential (349 GWh/yr), of which the BVs carry the majority — concentrated on the seasonal ships an onboard pack serves poorly, while onboard packs stay largest on the near-year-round hosts (per-ship division in Fig. 6) — avoiding of order 255 ktCO_2e/yr and earning an annual fleet profit of +€23.5M/yr. Each vessel's pack is sized independently in the optimisation, and the four settle at broadly similar capacities, each adapted to the leg set the joint route assignment hands it.

The gain from BVs is one of utilisation. Each BV works its pack about 150 equivalent full cycles a year, in the same busy band as the two year-round corridor hosts, whose onboard packs are the hardest-worked in the fleet, and far above the *seasonal* onboard packs the fleet must otherwise recruit to reach the same delivery, which cycle only a fraction as often and sit idle for the months their host leaves the EU. A roaming vessel is thus not used more intensively than the *best* onboard pack; rather it sustains near-year-round utilisation on capacity that would otherwise be stranded seasonally, so the fleet-average onboard pack, dragged down by those seasonal ships, works markedly less. The fleet stays profitable well past the peak but the margin per delivered MWh, about €67 at the peak, falls with every further MWh as progressively worse-utilised storage is recruited. The four-vessel optimum is capital-intensive — of order €430 M across battery packs, the four vessels and shore-side build — and the headline profit is surplus above a full 5 % real capital charge (plus battery replacement) already levied on every asset in the objective, not an accounting return on that capital.

Panel (b) reads the same frontier as the *marginal cost of electrification*: the cost of the last MWh delivered along the profit-optimal path. It rises as delivery grows and the fleet must recruit ever worse-utilised storage, and crosses the flat avoided-value line (∼301 €/MWh) near the envelope profit peak — beyond that scale the last MWh costs more to store than the fuel and carbon it displaces, coinciding with the profit peak. The curve is *not* monotone: onboard capacity scales smoothly, so its marginal cost climbs steadily, but each BV is a lumpy, fixed-cost investment, so the envelope can dip back after a BV entry. The text labels in panel (b) convert this marginal electrification cost into an *implied carbon price* — the marginal *abatement* cost: net of the displaced fuel cost and divided by the delivered MWh's well-to-wake carbon content, the carbon price that would rationalise each delivery target (§2.5). Along the profit-optimal path it rises from ∼€95/tCO_2e on the cheapest deliveries — onboard packs in the most active hosts, exploiting sunk OPS capacity — to today's stack's ∼€260/tCO_2e at the profit peak.

Panel (c) levelises the profit-max envelope into the cost bands each delivered MWh must cover

— charging energy, the battery pack at its annualised capital $\gamma_b K_b$, the onboard passenger-hull premium, the port coupler build and the power-handling upgrade, and the BV's hull-and-O&M and stern interface — leaving as headroom the *profit per MWh* delivered, a scale-invariant net margin. Introducing BVs trades higher vessel and interface cost for lower battery, safety and coupler-build cost; past four BVs the fixed vessel cost dilutes with delivery while coupler-build and battery costs climb.

Panel (d) decomposes *why* each delivered MWh is more profitable under the *combined* fleet — onboard packs plus four battery vessels — than under an onboard-only fleet, on an all-in basis — full pack and safety cost, net of the out-of-corridor OPS credit — so its endpoints read directly as the headline profit per MWh. The two fleets are compared at a *common* corridor delivery, the combined fleet's profit peak ($\sim$50 % of potential), which the onboard-only fleet is forced to match; earning the same gross avoided value on each MWh, the combined fleet delivers it more cheaply: the avoided passenger-hull premium, a smaller coupler build and a smaller pack outweigh the BV's added hull, O&M and stern interface, lifting profit per MWh from about €28 to €67. Charging energy nets to a small saving that conceals a larger split — the BVs buy power cheaper, roughly half offset by their self-propulsion and round-trip losses — with a smaller net weight penalty and lighter power-handling upgrade completing the balance. The largest single saving is the avoided passenger-hull premium — an engineering allowance rather than an observed price — with the smaller coupler build close behind: the same utilisation-for-capacity substitution read in cost terms, since a BV charges at a sunk coupler's full 20 MW against the $\sim$5 MW an onboard pack draws behind its hotel load, so pressing onboard delivery at speed forces dedicated build while the BV instead lifts utilisation of capacity AFIR already funds (below $\sim$50 GWh/yr the corridor builds none, charging purely on spare headroom). The underlying split of delivery across host ships is shown in App. G (Fig. 7b).

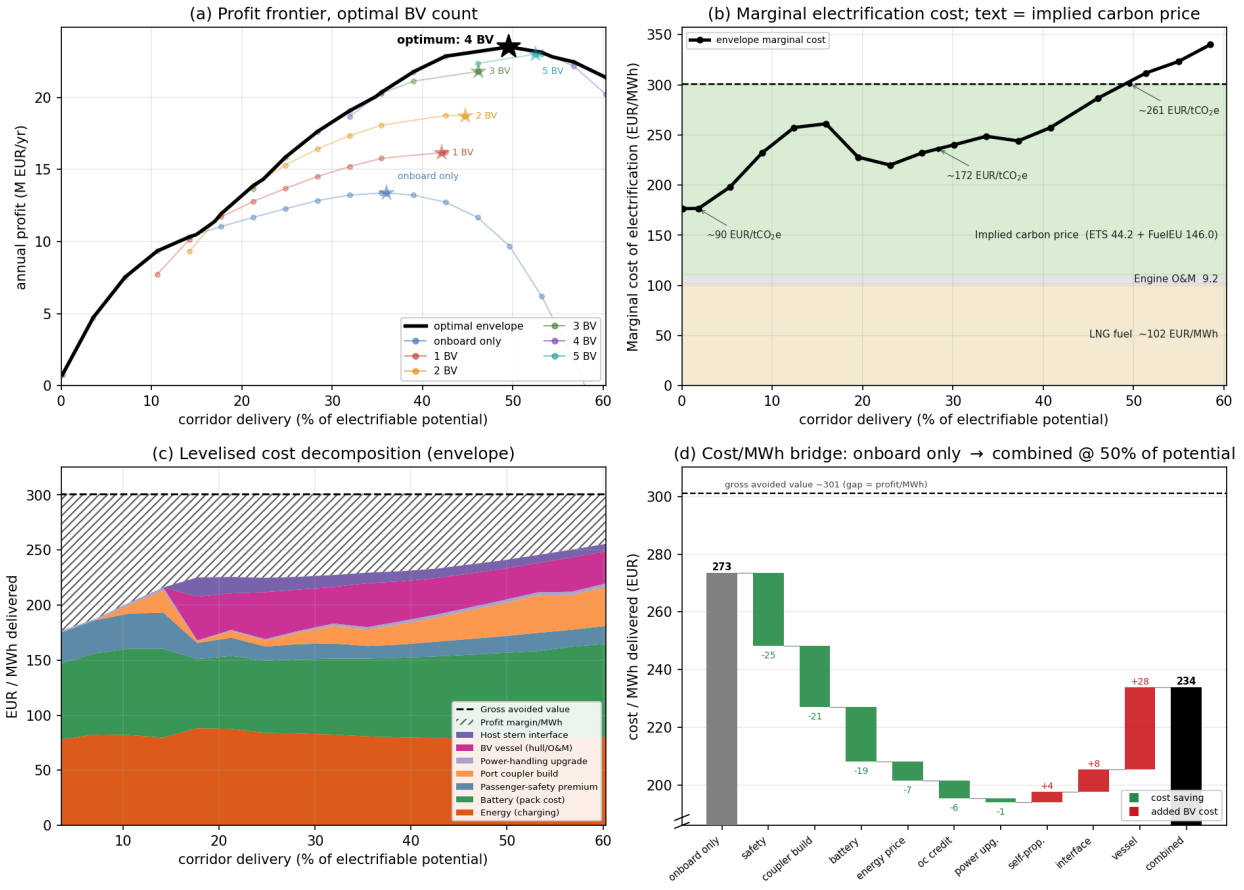


Figure 4: Profit-max frontier over battery-vessel count. At each corridor delivery the fleet fields the profit-maximising number of BVs (0–5). All three delivery axes are expressed as a share of the corridor's electrifiable potential ($\sum_l D_l \approx 705\,\text{GWh/yr}$). *(a)* Annual profit (€M/yr) against corridor delivery for onboard-only and each fixed BV count (faint arcs, each spanning its own profit peak) and their upper envelope (bold). The bold star marks the profit-max optimum (**four vessels**); faded stars mark each other count's own peak. *(b)* Marginal cost of the last delivered MWh along the profit-optimal envelope (bold). The shaded bands decompose the flat avoided value (∼€301/MWh, dashed top) into the displaced LNG fuel (∼€102/MWh), the avoided engine O&M (€9.2/MWh) and the carbon-priced remainder (ETS + FuelEU, ∼€190/MWh = the **implied carbon price** of ∼€260/tCO$_2$e); where the envelope marginal cost rises through the top band the last MWh's cost meets the avoided value. Text labels read the marginal cost at each point as an implied carbon price. *(c)* Levelised cost decomposition along the profit-max envelope (onboard-only at low delivery, then successively more BVs): each delivered MWh earns the gross avoided value (dashed); the stacked bands are the costs it must cover and the narrowing hatched headroom is the profit-per-MWh margin. *(d)* Cost-per-MWh bridge from the onboard-only fleet to the combined fleet (onboard packs plus four battery vessels), both at the combined fleet's profit peak (∼50 % of potential, the onboard-only fleet forced to match), with the energy bucket split into a cheaper-electricity (price/timing) effect and the extra self-propulsion draw. Unlike panel (c) (corridor-attributed), this bridge is on an all-in basis — full pack and safety cost less the net out-of-corridor OPS credit — so its endpoints equal the headline profit per delivered MWh. Green bars are savings from the BVs, red the costs they add (itemised in §4.1).

4.2 Sensitivities

Figure 5 and the full table (App. D) vary one modelling lever at a time from the profit-max base. Each case re-solves the fleet both to its own profit-max delivery *and* over battery-vessel count. Figure 5(a) recasts the single-lever sensitivities in intensive terms — profit per delivered MWh against electrification. Panel (b) traces how the profit-maximising fleet responds as the carbon price rises: sweeping the FuelEU pool price steps the optimum from three vessels to five and lifts electrification with it, a policy-robustness check on the base result. Where Figure 4(b) reads the marginal cost of electrification against today's fixed stack, here the stack itself is swept (implied-carbon conversion in §2.5).

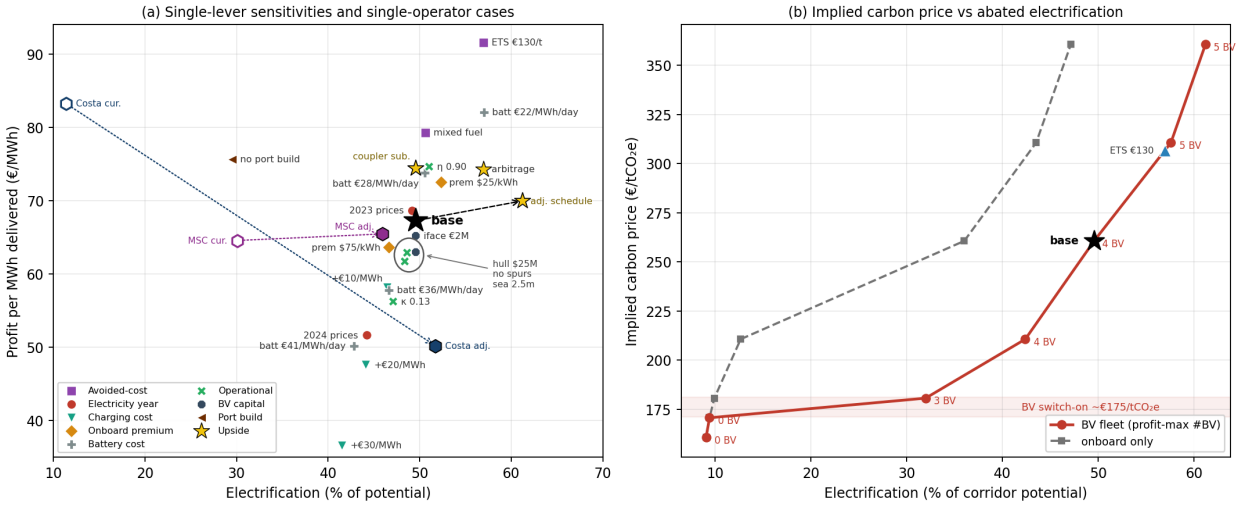


Figure 5: Profit margin and abatement against electrification. *(a)* Single-lever sensitivities and single-operator cases in intensive axes: profit per MWh delivered against electrification (share of the corridor's potential $\sum_l D_l$ delivered). Each single-lever case is its own profit-max optimum, grouped by colour; the base (four vessels) is the black star, grid arbitrage the gold star, and the Costa- and MSC-only cases (normalised by each operator's own potential) the hexagons — open on the current schedule, filled on the adjusted schedule, the dotted arrow marking the route-flexibility shift. Lever detail is in App. D. *(b)* The pooled fleet as a marginal-abatement curve: electrification against the implied carbon price of the policy stack that rationalises it (EU ETS held at spot while the FuelEU *pool price* is swept from €100 to €300/tCO$_2$e — extended below the €120–300 headline robustness range to expose the switch-on point — implied all-in carbon price ∼€160–360), for the profit-maximising fleet (re-solved over BV count) and an onboard-only fleet. Battery vessels switch on around 175 €/tCO$_2$e; because each carries a large fixed platform cost the optimum jumps from onboard-only to three vessels, then rises to four and five (the one- and two-vessel optima appear only on Fig. 4a's forced-delivery axis).

The BV case needs a threshold carbon price, which 2025 prices clear. Read as an implied carbon price (§2.5), the base stack is worth ∼€260/tCO$_2$e, comfortably above the ∼€175 at which battery vessels first pay. What moves the optimal fleet is the FuelEU pool price, the stack's largest and most variable component: sweeping it across a wide range steps the optimum from three vessels through four to five and lifts the electrified share from about a third of corridor potential to three-fifths (Table 10). Even a soft pool well below today's still keeps three vessels, while a tighter pool, or the widely projected rise in the ETS price to €130/t, pushes the fleet to five vessels.

Fuel and charging-price assumptions move the profit more than the electrified share. Valuing each host at its *actual* fuel — the mixed fleet of a few LNG dual-fuel newbuilds and mostly conventional MGO ships, rather than the all-LNG base — lifts profit by about a fifth but the electrified share barely at all. The all-LNG base is thus conservative on profit — the roaming fleet, whose seasonal hosts are mostly MGO, gains most — while realistic for the newbuilds that will dominate the corridor as it turns over. Pulling the other way, a network/levy adder on charging falls almost entirely on the roaming fleet, whose value is grid-charged corridor volume: as the adder rises the four-vessel optimum slides down and to the left in Figure 5(a), electrifying less of the corridor, but it holds four vessels across the whole €10–30/MWh range tested, the BV increment thinning from ∼€9 M to ∼€5 M rather than vanishing.

The result holds on different price years. Re-solving the whole fleet on realised 2024 and 2023 market data in place of 2025 leaves it clearly profitable in both, still electrifying much of the corridor (44 and 49 % of potential, against 50 % in 2025) and holding four vessels in each. Cheaper fuel and carbon in 2024 lowers the avoided value each MWh earns, while 2023's much dearer day-ahead electricity raises charging cost. The zonal price *structure* the pack exploits is similar across years, so the dispatch and fleet architecture are year-robust; what moves the level is the year's balance of charging cost against fuel-and-carbon value.

Robustness to the battery price. The pack is perhaps the most uncertain cost. Across the discount-rate band around the $87/kWh estimate (€28–36/MWh/day) electrification moves only a few points either side of the base, clearly profitable throughout and with all four vessels retained at both edges. Only beyond that band — the wider €22–41/MWh/day bracket in Table 10 (App. C) — does the optimal count shift: a fifth vessel enters at the cheap extreme, while four still hold at the dear.

The onboard premium sets the split between onboard and BV capacity. The passenger-hull premium a BV avoids ($c^{\text{prem}} = \$50$/kWh base) is the parameter the BV-versus-onboard comparison turns on. Cutting it to $25/kWh cheapens onboard storage and lifts the peak; raising it to $75/kWh does the reverse, with the BV then carrying about 63 % of delivery, since dear onboard storage throws more of the corridor onto the BVs. The BV keeps supplementing the onboard core throughout — the premium moves the *level* and the split, not the finding that a roaming vessel electrifies the seasonal hosts more cheaply than a year-round onboard pack.

Operational levers: sea state small, self-propulsion material. Tightening the sea-state operability gate from 4 m to 2.5 m significant wave height barely moves the electrified share — the BV re-routes onto the legs that survive — while trimming profit by about a tenth. Dropping the two delivery-only spurs (Tarragona and La Spezia, whose off-host repositioning energy the model does not separately price) is likewise minor, four vessels still optimal. Self-propulsion matters more: doubling the per-leg propulsion draw (from the calibrated $\kappa = 0.065$ to 0.13) lowers the peak appreciably, the four vessels retained but each delivering less. Pulling the other way, raising round-trip storage efficiency to $\eta_b = 0.90$ cuts charging losses.

Barring new port-coupler build widens the BV's edge. If the fleet may build no charging capacity beyond the sunk AFIR couplers, the four-vessel peak falls to €15.7M and the electrified share to 30 % of potential, but four vessels stay optimal. The restriction bites the onboard fleet hardest: an onboard-only fleet electrifies only 11 % of potential — against 36 % with build. Battery vessels, which take a whole spare coupler at its full rating, nearly triple that reach without any new capacity.

Grid arbitrage, coupler substitution and host routing flexibility present large upsides for BVs. Grid arbitrage is a *large* upside for the roaming fleet: enabling fleet-wide export lifts the profit-max peak by roughly a quarter (to €29.8 M) and draws two further vessels into the money — six rather than four — without cannibalising corridor delivery. Because dispatch assumes perfect foresight of hourly prices, this uplift is best read as a ceiling on the arbitrage available. The same lever barely

moves an onboard-only fleet — a roaming vessel banks far more idle dock-hours in which to cycle than a ship that berths briefly on its own itinerary.

A BV standing in for a coupler on the rare all-berths-busy hours lets the port build *one fewer* (§2.6.2). Screening the corridor and re-solving each candidate with its coupler removed (App. E), two hubs qualify — **Livorno** and **Naples** — where the routed fleet backfills enough of the marginal coupler's duty to hold coverage above the 90 % AFIR floor and the tighter charging costs less than the avoided build. Together they shed two couplers for a net **€2.0M/yr** credited to fleet profit — the avoided shore-power capital less the tighter charging and the few host-calls left un-electrified when no BV is present to stand in (decomposition in App. E, Eq. 20). As an avoided-*capital* value accruing to the port, realising it requires a contract transferring the saving to the BV operator, so we report it as an upside outside the headline.

Reorganising itineraries so that marked hosts run each typically-weekly circuit in whichever direction best feeds the roaming fleet — the *adjusted schedule* of §2.6.3, which opens return-direction legs a BV can slot into otherwise-idle time — lifts the current-schedule outcome of §4.1, raising the corridor share the fleet electrifies from 50 to 61 % of potential (profit €24 M → €30 M) without adding a hull. The reversal set is chosen to maximise the *joint* fleet's delivered energy: at the four-vessel optimum it reverses 150 of 530 typically-weekly circuits, spread across seventeen of the nineteen hosts (Table 9, App. B). We keep it out of the headline because it presumes operator cooperation in re-timing published itineraries rather than a pure dispatch choice — but it is the single most valuable "non-carbon" lever in the study.

Traffic density helps — but route flexibility can substitute for much of it. Restricting the fleet to a single owner (Table 10, fleet-density rows) shows how much the profit-max case rests on host density: MSC's thirteen ships support two BVs and Costa's five support none — against four for the pooled corridor — because a sparser network cannot keep a shared pack busy enough to cover its hull, and each operator electrifies far less of its own potential (11 % and 30 % for Costa and MSC, against 50 % pooled). Here alone we deliberately combine two levers — the single-operator restriction and the adjusted schedule of §2.6.3, applied to that operator's own circuits — because the point is precisely that a small fleet can trade one for the other. Allowing the operator to direction-switch its own itineraries closes much of the gap, making Costa's vessel viable and lifting Costa and MSC to 52 and 46 %: route flexibility can substitute for the traffic density a small fleet lacks.

5 Discussion

5.1 Why the BV lifts the frontier, and what deployment requires

A BV *substitutes utilisation for capacity*: one roaming pack kept busy across many hosts meets a delivery target that seasonal onboard packs — idle while their host is away — could match only with far more installed storage and *dedicated* port-coupler build (Fig. 7c,d).

Several of our simplifications cut against the BV and others flatter it; all are probed in the sensitivities and weighed in the Limitations (§5.2). The harder questions lie not in the model but in deployment and policy, which we group under three heads.

Feasibility and the policy stack. Unlike an onboard pack, which rides technology already in service, the BV is a speculative asset, and its most consequential operational assumption is at-sea power transfer. The *initial* coupling is not the difficulty — it can be made close to port in sheltered water — but maintaining a high-power connection and transferring propulsion-scale power to a host *under way* in an open seaway is not yet an established maritime operation; the closest precedent is

naval underway replenishment (United States Navy, 2007; Smith, 1992). We gate operability on sea state but otherwise abstract from the mechanics of the connection and station-keeping. The host-side stern interface is costed as a propulsion-scale connection for which no catalogue price yet exists, and is carried as a flagged engineering estimate. The engineering feasibility of the underway connection, including seakeeping analysis, is the subject of ongoing companion work. Its economics also rest on the EU regulatory stack, as does any decarbonisation pathway: the baseline uses present-day carbon and fuel prices and is near-flat across the horizon — a tightening ETS cap and FuelEU pool are carried as upside, not relied on — so the stability and predictability of that pricing is a determinant of viability here as elsewhere. More particular to the BV is at-sea supply: reading FuelEU's zero-intensity treatment of shore power onto energy transferred to a host *under way* — which the regulation recognises in the compliance balance but does not spell out for delivery by a separate vessel — is a modest clarification of a mechanism the stack already contains, not a new policy (§2.5, App. F).

Joint electrification and cost recovery. A useful public action at corridor hubs is to support grid connections sized beyond the hotel-scale supply AFIR mandates. Such connections serve wider port electrification and shoreside storage, and their scale economics argue for co-planning BV charging with the cruise-OPS rollout rather than treating the two as separate assets. Greater coupler capacity is valuable but not decisive — even with no build permitted the four-vessel fleet stays the profit-max choice (§4.2). The couplers a BV borrows were built to cut *at-berth* emissions; significant *at-sea* electrification instead turns on cheaper power, so how port infrastructure is funded — electricity tariffs, berth fees, or general charges — is decisive, and BV economics are especially sensitive to the electricity price. The efficient arrangement rewards the BV for filling coupler-hours that would otherwise sit idle — a port selling that spare capacity to a price-sensitive bulk buyer would rationally price below its premium hotel-service tariff — and the network/levy-adder sensitivities (§4.2) show the case tolerates a material charging markup in any event. The BV's other coupler-side value — raising utilisation so a hub can defer a marginal coupler — is an avoided *capital* cost accruing to the port, so realising it likewise requires attributing the saving across port, host and operator by contract. Siting is easier on this side too: an unmanned vessel can tie in at a cheaper industrial or utility point away from passenger terminals, and is smaller and more manoeuvrable in port than the ships it serves.

Scale and density. Both the onboard and BV architectures depend on scale, but differently. An onboard pack shares the port-coupler build across a host's repeated calls, lowering the connection cost per delivered MWh. A BV carries more fixed cost — a stern interface on each host and the flat platform cost — so its viability turns on traffic density and on scheduling that keeps one roaming pack busy across many hosts. Value is therefore set by network density, not fleet size alone: a hub serving twenty hosts rather than five, or several operators rather than one, offers more coupling opportunities per idle hour — the same mechanism that drives our profit peak — and the economics resemble airline-hub and shared rolling-stock utilisation, in which an asset shared across many users outperforms captive ones tied to a single one. Our owner-subset solves (Table 10) put MSC alone at the scale to field vessels of its own; reaching the corridor-wide optimum, or making a smaller operator's vessel viable, rests on electrification across operators — eased by the corridor's concentrated ownership, MSC and Costa between them running eighteen of the nineteen ships — on a shared or third-party vessel, or on the route flexibility that partly substitutes for density (§4.2). The flexibility of BV routing means initial deployment could be confined to a single dense pair such as Barcelona–Marseille.

5.2 Limitations

Several caveats bound these conclusions. We group them by the likely direction each pushes the value of electrification — conservative (understating it), optimistic (overstating it), or broadly neutral — noting where an over- or understatement bears mainly on one storage form. Table 5 summarises them; the text elaborates the reasoning.

Table 5: The main modelling simplifications not separately quantified, and the direction each pushes the value of electrification: conservative (understates), optimistic (overstates), or neutral. Parenthetical tags mark where an over- or understatement bears mainly on the onboard pack, the BV, or both. Levers we *do* quantify — fuel choice, network/levy adder, FuelEU pool price, ETS price, battery cost, self-propulsion and sea state — are in the sensitivity table (Table 10).

Simplification	**Likely impact**
Corridor sized to today's tracked fleet (2040 growth and untracked ships omitted)	Understates: growth adds hosts and coupling opportunities, favouring a roaming BV in particular.
Realised 2023–25 electricity prices (a future, solar-heavier grid not modelled)	Understates: a more volatile grid richer in near-zero solar troughs would cheapen charging for both forms.
Finer re-timing unpriced (call times, en-route speed)	Understates: catching cheap power would raise the value of electrifying either form.
Onboard retrofit unpriced (space, weight, stability and class approval)	Overstates (onboard): its retrofit burden goes uncounted; a BV needs only a stern interface.
Dry-dock idle not modelled	Overstates (onboard): an onboard pack earns nothing in dock while a fleet BV keeps working elsewhere.
Coupler outages and congestion not modelled	Overstates (onboard): an onboard pack waits at its berth, whereas a BV reroutes to a free coupler.
Whole-circuit reversal treated as costless	Overstates (BV): real frictions (Mistral, berth slots) make the direction-switching uplift an upper bound.
Price-taker (no price formation)	Overstates (both): at scale the fleet's own demand would lift the trough prices it exploits, symmetrically.
Annual-average bunker price	Neutral: operators buy forward, so a smoothed price is the relevant one.
Host demand from cube-law scaling (no weather/occupancy variation)	Small: the headline is insensitive where demand meets or exceeds the deliverable.
Sea state enters energy and a binary operability gate only	Uncertain: schedule-level weather effects (delayed calls, missed charging windows) are not modelled.

Direction of the unquantified simplifications. The *conservative* omissions in Table 5 — today's tracked fleet, realised rather than future solar-heavy prices, and only whole-circuit reversal — understate either storage form, their omitted density and coupling opportunities favouring a roaming asset in particular. The *optimistic* ones mostly fall on the *onboard* pack, whose retrofit, dry-dockings and coupler-outage idle we do not price, whereas a fleet BV needs only a stern interface and reroutes to a free coupler. Two overstatements are not onboard-specific: taking prices as given flatters both forms at scale, and treating whole-circuit reversal as costless overstates the *BV's* direction-switching uplift, which real frictions — the Mistral's one-way advantage on the Barcelona–Marseille crossing,

berth-slot patterns favouring a fixed direction — make an upper bound. That uplift needs only the minority of circuits (about a third) whose reversal fills an otherwise-idle stretch, so a few well-chosen reversals likely capture most of it.

6 Conclusion

Batteries can decarbonise at-sea propulsion on a busy cruise corridor in two ways — carried *together* with the host as fixed onboard packs, or sent *alone* as dedicated battery vessels that power hosts under way and recharge from sunk port couplers. Assessing both against the Western Mediterranean's real schedules, sea states and market prices, we find each earns according to its comparative advantage, and that combining them electrifies far more of the corridor than onboard batteries alone. Onboard batteries are profitable first and unaided, but only on the handful of ships that stay on the corridor near-year-round; forcing further onboard delivery recruits seasonal packs that sit idle for the months their host leaves the EU, and profit falls. Battery vessels break this ceiling: at the profit peak a four-vessel roaming fleet supplementing the onboard core electrifies about 50 % of the corridor's at-sea energy, against 36 % without them. Their advantage is redeployability. Cycled no more often than the best onboard pack, a roaming vessel stays busy where a fixed pack cannot, simply moving to the hosts that remain when a seasonal host leaves the corridor. It also draws on a coupler the host fleet has left idle at its full rating, rather than the remainder an onboard pack gets behind its own hotel load, raising the utilisation of the publicly funded couplers AFIR is installing — and can even let a hub defer a marginal one.

The findings are less corridor-specific than they may seem. Our results rely on a set of sunk OPS chargers free for a BV to borrow, and the EU policy stack that prices the avoided emissions. The drivers behind the result, though, are common to cruising more generally: fuel-electric hosts, calm and densely-spaced adjacent hops, and the overnight-sail, day-in-port rhythm that enables regular battery cycling and leaves couplers idle by day for a roaming vessel to use. With port charging infrastructure and an effective carbon price, the qualitative case should therefore carry to other cruise regions sharing these features, the quantitative economics turning on local charging prices, traffic density and sea state rather than on anything unique to this corridor. Finally, we price only two storage architectures — a permanently installed onboard pack and a dedicated BV — and not the intermediate options, chief among them containerised, swappable packs charged ashore and exchanged dockside, which share the BV's separation of storage from host and its shore-charging economics but trade the BV's at-sea coupling for dockside crane handling, the port dwell — or lost utilisation of an overnight wait — of moving a pack to a spare charger, and a host retrofit far heavier than a stern interface, the pack still riding the host's space and weight at sea. We leave a fuller comparison to future work.

Battery vessels, though, require sufficient traffic density. The corridor is concentrated in just two operators, and the case weakens sharply if only one of them electrifies; schedule flexibility recovers much of the shortfall, but reaching the corridor-wide optimum turns on the operators pooling their ships and schedules behind a shared or third-party vessel, while grid arbitrage and the vessels' ability to defer coupler build may each add materially on top. Viability rests above all on an effective carbon price, supplied here by the EU regulatory stack, and on charging at port. If the at-sea power transfer can be made to work, the BV would carry direct electrification past the short, port-anchored duty cycles to which stored energy has so far been confined, onto the longer and denser routes that today's fuels still own.

References

ABB. (2024). *ABB and Seafjord Energy introduce pioneering e-bunker vessel concept* [Press release]. ABB.

Anwar, S., Zia, M. Y. I., Rashid, M., Zarazua de Rubens, G., & Enevoldsen, P. (2020). Towards ferry electrification in the maritime sector. *Energies*, 13(24), 6506.

BloombergNEF. (2024). *Lithium-ion battery pack prices see largest drop since 2017, falling to $115 per kilowatt-hour*. BloombergNEF.

Bullock, S., Higgins, E., Crossan, J., & Larkin, A. (2023). Improving shore power project economics at the Port of Aberdeen. *Marine Policy*, 152, 105625.

CE Delft. (2022). *The role of shore power in the future maritime fuel mix* (Publication 22.210314.164; J. Faber, R. van den Berg, & J. de Vries). CE Delft.

Cruise Lines International Association. (2024). *State of the cruise industry report*. CLIA.

CruiseMapper. (2025). *Cruise ship itineraries and schedules* [Data set]. `https://www.cruisemapper.com`

DNV. (2024). *2024 Battery scorecard*. `https://www.dnv.com`

Ember. (2026). *Carbon Price Viewer: EU ETS allowance price* [Data set; front-December contract]. Ember. `https://ember-climate.org/data/data-tools/carbon-price-viewer/`

Enerdata. (2025). *Carbon price forecast under the EU ETS* [Executive briefing]. Enerdata. `https://www.enerdata.net/publications/executive-briefing/carbon-price-forecast-under-eu-ets.html`

ENTSO-E. (n.d.). *Transparency Platform: Day-ahead prices* [Data set]. European Network of Transmission System Operators for Electricity. `https://transparency.entsoe.eu`

European Commission. (2024). *Questions and answers on Regulation (EU) 2023/1805 on the use of renewable and low-carbon fuels in maritime transport (FuelEU Maritime)* [Q&A 2.5, on-board battery-electric technologies]. Directorate-General for Mobility and Transport. `https://transport.ec.europa.eu/transport-modes/maritime/decarbonising-maritime-transport-fueleu-maritime_en`

European Union. (2023a). Regulation (EU) 2023/1804 of the European Parliament and of the Council of 13 September 2023 on the deployment of alternative fuels infrastructure (AFIR). *Official Journal of the European Union*, L 234, 1–47.

European Union. (2023b). Regulation (EU) 2023/1805 of the European Parliament and of the Council of 13 September 2023 on the use of renewable and low-carbon fuels in maritime transport (FuelEU Maritime). *Official Journal of the European Union*, L 234, 48–100.

European Union. (2023c). Directive (EU) 2023/959 of the European Parliament and of the Council of 10 May 2023 amending Directive 2003/87/EC establishing a system for greenhouse gas emission allowance trading within the Union [extending the EU ETS to maritime transport]. *Official Journal of the European Union*, L 130, 134–202.

Frković, L., Ćosić, B., Falkoni, A., Pukšec, T., & Vladimir, N. (2024). Shore-to-ship: Enabling the electrification sustainability of maritime transport with the nexus between berthed cruise ships and renewables in the isolated energy systems. *Ocean Engineering*, 302, 117537.

Green Cruise Port. (2018). *Onshore power supply for cruise vessels: Assessment of opportunities and limitations for connecting cruise vessels to shore power* [INTERREG V B project report]. Green Cruise Port.

He, W., Valøen, L. O., Olsen, K. V., Kjeka, K. M., Fredriksen, B. M., Petiteau, M., Touat, A., Sætendal, H., Howie, A., Howey, D., Kandepu, R., & Hammershøj, C. F. (2024). Lessons learned from the commercial exploitation of marine battery energy storage systems. *Journal of Energy Storage*, 87, 111440.

International Carbon Action Partnership. (2026). *EU Emissions Trading System (EU ETS): Allowance price data* [ETS detail page; 2025 auction and secondary-market annual averages]. International Carbon Action Partnership. `https://icapcarbonaction.com/en/ets/eu-emissions-trading-system-eu-ets`

International Maritime Organization. (2020). *Fourth IMO GHG study 2020*. International Maritime Organization.

International Maritime Organization. (2023). *2023 IMO strategy on reduction of GHG emissions from ships* (Resolution MEPC.377(80)). International Maritime Organization.

International Transport Workers' Federation. (2018). *Autonomous ship "Yara Birkeland" — how far has the industry reached?* International Transport Workers' Federation.

Kanchiralla, F. M., Grunditz, E., Nordelöf, A., Brynolf, S., & Wikner, E. (2025). Environmental and economic assessment of electric ferries with different lithium-ion battery technologies. *Applied Energy*, 396, 126274.

Kersey, J., Popovich, N. D., & Phadke, A. A. (2022). Rapid battery cost declines accelerate the prospects of all-electric interregional container shipping. *Nature Energy*, 7(7), 664–674.

Kistner, L., Bensmann, A., & Hanke-Rauschenbach, R. (2024). Potentials and limitations of battery-electric container ship propulsion systems. *Energy Conversion and Management: X*, 21, 100507.

Li, R., Li, H., Huang, W., Tao, H., Xu, W., Tai, N., & Li, C. (2025). Accelerating green shipping with spatially optimized offshore charging stations. *Nature Energy*, 10(2), 243–254.

Mackenbach, K., Singhal, H., Zhao, T., Pradhan, S., Wolf, R. M., & Nevenzeel, K. (2026). Techno-economic feasibility of electrification for short sea shipping. *npj Clean Energy*, 2, 14.

MAN Energy Solutions. (2019). *Batteries on board ocean-going vessels*. MAN Energy Solutions.

Moon, H. S., Park, W. Y., Hendrickson, T., Phadke, A., & Popovich, N. D. (2025). Exploring the cost and emissions impacts, feasibility and scalability of battery electric ships. *Nature Energy*, 10(1), 41–54.

Nazir, C. P. (2021). Offshore electric ship charging station: A techno-economic analysis. *International Journal of Marine Engineering Innovation and Research*, 6(4), 210–225.

OceanScore. (2026). *OceanScore Pool-Price Index (OPX): FuelEU Maritime compliance surplus pricing* [Market index]. OceanScore. `https://oceanscore.com/solutions/fueleu-pooling-marketplace/`

Open-Meteo. (n.d.). *Marine weather API: significant wave height reanalysis* [Data set]. `https://open-meteo.com`

Osipova, L., & Carraro, C. (2023). *Shore power needs and CO_2 emissions reductions of ships in European Union ports: Meeting the ambitions of the FuelEU Maritime and AFIR* (Working Paper 2023-24). International Council on Clean Transportation.

Pavlenko, N., Comer, B., Zhou, Y., Clark, N., & Rutherford, D. (2020). *The climate implications of using LNG as a marine fuel* (Working Paper 2020-02). International Council on Clean Transportation.

Perčić, M., Vladimir, N., & Fan, A. (2020). Life-cycle cost assessment of alternative marine fuels to reduce the carbon footprint in short-sea shipping: A case study of Croatia. *Applied Energy*, 279, 115848.

PortMinder. (2025). *Cruise ship schedules and port-call calendars* [Data set]. `https://www.portminder.com`

Protergia. (2026). *TTF natural gas prices per month* [Data set; monthly Dutch TTF settlement prices]. Protergia. `https://www.protergia.gr/en/home/natural-gas/ttf-prices-per-month/`

Ren, B., Li, Y., Wen, X., & Ma, H.-L. (2025). The offshore mobile charging vessel location and electric vessels charging scheduling problem. *Ocean Engineering*, 339, 122181.

Risposte Turismo. (2026). *Speciale Crociere 2026: Il traffico crocieristico in Italia nel 2025 e le previsioni per il 2026* [Special report]. Risposte Turismo.

Ship & Bunker. (2026). *Global 20 Ports Average bunker prices (G20): MGO* [Data set; annual-average marine-gas-oil bunker prices]. Ship & Bunker. `https://shipandbunker.com/prices/av/global/av-g20-global-20-ports-average`

Smith, T. C. (1992). *An examination of underway replenishment seakeeping criteria* (Technical Report). Canadian National Defence, Defence Research and Development Branch.

Sruthy, V., & Preetha, P. K. (2024). Implementation and operational feasibility of an offshore floating charging station for sustainable marine transportation. *Environment, Development and Sustainability*, 26,

20931–20962.

Stopford, M. (2009). *Maritime economics* (3rd ed.). Routledge.

Transport & Environment. (2023). *The return of the cruise: How luxury cruises are polluting Europe's cities.* Transport & Environment (T&E).

United States Navy. (2007). *Underway replenishment* (Naval Warfare Publication NWP 4-01.4). Department of the Navy.

Wiskich, A. (2025). How can long-distance battery-powered container ships stack up? A speculative Asia–Europe illustration. *Research in Transportation Economics*, 113, 101636.

Wiskich, A. (2026). *Coupling maritime electrification with renewable power systems: The potential of battery vessels* [Working paper].

Zhang, Y., Sun, L., Fan, T., Ma, F., & Xiong, Y. (2023). Speed and energy optimization method for the inland all-electric ship in battery-swapping mode. *Ocean Engineering*, 284, 115234.

A Nomenclature

Table 6: Nomenclature. Decision variables (bold) are continuous; non-bold are parameters or data. A leg l encodes its own origin, destination and departure/arrival periods; the thirteen charging hubs and their seven price zones are enumerated in Table 7.

Sets / indices	**Symbol**	**Elements**
Port charging hub	r, ρ	13 hubs (enumerated in Table 7); a corridor leg runs $r \to \rho$ between distinct, adjacent hubs
Day-ahead price zone	z	7 zones; $z(r)$ is hub r's zone and $z_m(h)$ the zone member m occupies at h (in-zone hubs share $\pi_{z,h}$)
Period ($\delta = 1$ h)	h	$1 \dots T$ (one year)
Regulatory stack (avoided-cost)	—	present-day: ETS spot, FuelEU at pool price (120–300 €/t sens.)
Leg (directed sailing)	l	origin $r(l)$, dest. $\rho(l)$, periods $h_l^{\text{dep}}, h_l^{\text{arr}}$
Fleet member (ship or BV)	m	indexed over $\mathcal{M} = \mathcal{J} \cup \mathcal{I}$
Leg set; docked-period set of member m	$\mathcal{L}_m, \mathcal{H}_m$	$\mathcal{L}_m$: m's legs (chain-DP route for a BV, own corridor sailings for a ship); $\mathcal{H}_m$: its docked periods
Decision variables		
Usable battery capacity of member m (MWh)	$\mathbf{E}_m$	endogenous; ≥ 0
Additional charger build at port r (MW)	$\mathbf{x}_r$	endogenous; ≥ 0
Energy delivered by m on leg l (MWh)	$\mathbf{d}_{m,l}$	$0 \leq \mathbf{d}_{m,l} \leq D_l$
State of charge of member m (MWh)	$\mathbf{s}_{m,h}$	continuous; ≥ 0
Charge rate of member m at h (MW)	$\mathbf{c}_{m,h}$	continuous; docked periods only
Out-of-corridor OPS energy per call, ship j (MWh)	$\mathbf{o}_j$	≥ 0; $\leq \sqrt{\eta_b}\mathbf{E}_j$, $\leq \sqrt{\eta_b} P_j^{\text{oc}} \tau_j$ (Eq. (21)); earned at n_j calls
Parameters / data		
Onboard ships; BVs in fleet	J; I	$J = 19$ (all corridor ships offered; 4-ship core adopts); I exogenous, varied (4 at the profit-max peak)
Fleet sets	$\mathcal{J}, \mathcal{I}, \mathcal{M}$	$\mathcal{M} = \mathcal{J} \cup \mathcal{I}$
Port coupler capital (€/MW)	C_r	$c_{\text{chg}} + c_{\text{grid}} \approx$ €850k; $\gamma C_r \approx$ €68k/MW/yr
Free charge headroom: onboard; BV (MW)	$P_{\text{ob}}, P_{\text{BV}}$	$R - H_c = 5$; $R = 20$

continued on next page

Table 6, continued

Description	Symbol	Elements / value
Grid-connection charge cap (MW)	P^{grid}	50 (BV); $50 - H_c = 35$ (onboard)
Hosts a BV couples to	H_i	distinct hosts BV i delivers on
Power-handling upgrade capital: host flank / BV (€/MW)	C^{upg}	50 €/kW (GCP); $\gamma C^{\text{upg}} \approx$ €4010/MW/yr
BV hull capital (€)	K^{BV}	$\sim$ €18.4M ($20M); $\gamma K^{\text{BV}} \approx$ €1.48M/yr
BV fixed O&M (€/yr)	F^{OM}	$\sim$ €1.02M ($3000/day; §2.3)
Host stern interface capex (€)	I_j^{iface}	€1.5M + €50/kW $(P_j - 20\,\text{MW})^+$ (P_j peak power delivered to host j) per host j a BV fits; shared across BVs (App. C)
Host flank interface: peak onboard charge above free headroom (MW)	$\mathbf{u}_j$	decision var.; ≥ 0, $\forall j \in \mathcal{J}$
BV power build: peak BV power (charge/discharge) above 20 MW	$\mathbf{u}_i^{\text{BV}}$	decision var.; ≥ 0, $\forall i \in \mathcal{I}$
Self-propulsion / pack-weight coefficient of member m	κ_m	BV self-propulsion reserve $\kappa_{\text{BV}} \approx 0.065$ (derived from BV mass; 0.13 sens.; written κ where unambiguous); onboard weight penalty κ_{onb} ($\approx$1.1% of leg propulsion per 100 MWh usable)
Per-leg propulsion scale	θ_l	host leg propulsion energy / Barcelona–Marseille anchor leg
Out-of-corridor propulsion sums, ship j	$\theta_j^{\Sigma}, \theta_j^{c}$	$\sum_l \theta_l$ (all oc legs); $\sum_l s_l \theta_l$ (EU-carbon-weighted); App. B
EU maritime carbon scope of oc leg l	s_l	$\in \{1, \frac{1}{2}, 0\}$ by count of EU/EEA endpoints
Realised day-ahead price (€/MWh)	$\pi_{z,h}$	ENTSO-E hourly, 2025 (per price zone z)
Avoided at-sea cost (€/MWh)	$\hat{p}$	Eq. (14) (corridor scalar)
Fuel; carbon components of $\hat{p}$ (€/MWh)	$\hat{p}_{\text{f}}, \hat{p}_{\text{c}}$	$\sigma p_{\text{fuel}} \approx 102$; ETS+FuelEU ≈ 190 (Eq. (14))
Realised corridor net margin (€/MWh)	μ	avoided value net of charge cost per delivered MWh; fixed-point (Eq. (22))
Out-of-corridor Cat. 1 calls; headroom, mean dwell	n_j; P_j^{oc}, τ_j	per ship j (5/10 MW; $\tau_j \leq 14$ h; App. B)
Specific fuel consumption (t/MWh elec.)	σ	0.165 LNG (0.19 MGO)
Tank-to-wake CO_2e factor (tCO_2e/t fuel)	ε	3.62 LNG (combustion 2.75 + slip 0.87); 3.206 MGO
Bunker fuel price (€/t)	p_{fuel}	Table 4 (LNG base; MGO sens.)
EU allowance (EUA) price (€/tCO_2)	p_{EUA}	Table 4
EU ETS phase-in fraction	ϕ	1 (full scope, from 2026)
FuelEU marginal compliance value (€/MWh)	m_{FuelEU}	$p^{\text{pool}}\times$ full WtW intensity; $\approx$€146/MWh at base (Eq. (15))
Round-trip efficiency	η_b	0.86, applied $\sqrt{\eta_b}$ each way
Installed propulsion; max speed	$P_{\max}^{\text{prop}}, V_{\max}$	44 MW; 22.7 kn (World Europa/Asia; cube law)
Coupled at-sea window on leg l (h)	h_l^{sea}	gap − buffer, $\in [4, 20]$
Host energy demand on leg l (MWh)	D_l	Eq. (1) ($\sim$390 avg)
Per-leg deliverable cap (routing)	$\bar{E}$	chain-DP leg weight $\min(D_l, \bar{E})$ (24); one BV's single-pass deliverable, iterated to the solved packs; $\approx$227 MWh
Port couplers at hub r (integer)	K_r	90% major-cruise stop coverage
Port coupler rating; at-berth hotel draw (MW)	R, H_c	20; 15 (large), 10 (small)
Cruise occupancy at r, h	$o_{r,h}$	real ICS schedule (all berthed majors, dispatched or not)
Significant wave height; gate threshold (m)	H_s, $H_s^{\star}$	reanalysis; 4 (base, 2.5 sens.)
Battery pack capital (€/MWh usable)	K_b	$\sim$ €100k; $\gamma_b K_b \approx$ €11.3k/MWh/yr, band $\sim$10.0–13.3k (App. C)

continued on next page

Table 6, continued

Description	Symbol	Elements / value
Onboard passenger-hull premium (€/MWh usable)	c^{prem}	certification/space capital a BV avoids; \$50/kWh nameplate base $\rightarrow \gamma c^{\text{prem}} \approx$ €4.6k/MWh/yr (\$25/\$75 sens.; App. C)
Real discount rate; infra, battery life	i, L, L_b	5%; 20 yr, 12 yr
Capital-recovery factor: infra; battery	γ, γ_b	0.080 (20 yr); 0.113 (12 yr)
Period length (h)	δ	1

B Cruise host schedule

The host schedule (§2.1) is built from real port-call calendars: year-long itineraries for the Western- and Central-Mediterranean cruise fleet, downloaded from PortMinder. A parser reads each ship's chronological calls and, taking each corridor port as a charging hub (Table 7), emits a directed overnight leg for each pair of *consecutive* calls between two *distinct, adjacent* hubs whose origin can charge; each leg keeps its real departure time (week-shifted, preserving weekday and time-of-day) and its energy is set by the speed-based model (Eq. (1)). The *nineteen* ships that make regular corridor sailings form the study fleet (Table 8): onboard packs are offered on all nineteen and adopted in value order by the joint LP, and all ships are also eligible BV hosts. Projected onto 2025 the nineteen-ship pool yields 1814 corridor legs. Barcelona is the natural western anchor: it was the busiest cruise port in the entire Mediterranean in 2025 (≈4.0 M passengers), and all five of the busiest Mediterranean cruise ports — Barcelona, Civitavecchia, Marseille, Palma and Naples — are corridor hubs (Risposte Turismo, 2026).

By calendar month the 2025 corridor legs run 75, 77, 105, 118, 184, 184, 203, 196, 181, 221, 174, 96 (January–December), plateauing across May–October. The Central-Mediterranean hubs (Naples, Palermo, Valletta, Cagliari and the Iberian pair Valencia/Palma) extend the corridor's reach into the shoulder seasons in particular.

Table 7: Corridor charging hubs. The corridor is thirteen charging hubs across seven day-ahead price zones — each hub a corridor port with an OPS berth. A leg runs between distinct, adjacent hubs and may originate only at a charging hub. Two further ports enter as *delivery-only* spurs (Tarragona and La Spezia; mechanism below); Salerno is grouped with Naples (both directions).

Hub	Price zone	Delivery-only spur
Barcelona	ES	Tarragona (delivery-only, from FR)
Valencia	ES	—
Palma	ES	—
Marseille	FR	—
Genoa	IT-NORD	—
Savona	IT-NORD	—
Livorno	IT-CNOR	La Spezia (delivery-only, from IT-CSUD)
Civitavecchia	IT-CSUD	—
Naples	IT-CSUD	Salerno (grouped, bidirectional)
Cagliari	IT-SARD	—
Palermo	IT-SICI	—
Messina	IT-SICI	—
Valletta	IT-SICI[†]	—

[†]Valletta/Malta proxied to the Sicilian (IT-SICI) price series.

Delivery-only spurs. A host sailing Marseille → Tarragona, for instance, is recorded as an FR → ES leg served in the from-FR direction only: the BV couples near Marseille, powers the host underway across the Gulf of Lion, disconnects at sea as the host nears Tarragona, and then recharges at Barcelona — the ES-zone hub it passes en route — rather than at Tarragona itself, which has no assumed high-power tie. In the reverse (from-ES) direction the recharge hub no longer lies between origin and destination, so the spur is not served. La Spezia works the same way on the Tyrrhenian side, served only from Civitavecchia with the BV recharging at Livorno as it passes. Ports with weak island grids or tender-only access, or that would sit as a distorting through-waypoint — Ibiza, Toulon, Cannes, Olbia, Alicante, Tunis and others — are excluded from BV service entirely. A spur's recharge run — the BV's leg from the mid-sea disconnect to the hub it passes — is not separately charged self-propulsion; as a check, re-solving with both spurs removed altogether lowers the electrified share by roughly one percentage point, with four vessels still optimal (Table 10, Fig. 5a).

Leg-window and direct-return tests. The coupled at-sea window h_l^{sea} that sizes each leg's energy (Eq. (1)) is the scheduled arrival-to-departure gap less a ∼1.5 h port-manoeuvring buffer, floored at 4 h — so a very short gap cannot imply an unphysical speed — and capped at 20 h for an adjacent overnight hop (34 h for a direct circuit-closing return), since a longer gap means the two calls are days apart (the ship has left the corridor between them) rather than one continuous crossing. A *direct* return (§2.1) is admitted as a single BV leg only where the schedule shows one genuine crossing: departure-to-arrival gap ≤36 h and inferred cruise speed ≤25 kn, which excludes both seasonal repositioning legs and schedule artefacts.

Beyond this dispatched fleet, nine further major-cruise ships — Costa Diadema and Costa Fortuna, Harmony and Oasis of the Seas, MSC Bellissima and MSC Fantasia, Silver Ray, and Viking Sky and Viking Vela — make regular corridor calls. We do not size onboard packs or route BVs for these ships, but we do read their berth-days from the same PortMinder calendars, so they enter the coupler-occupancy count $o_{r,h}$ that governs shore-power contention (§2.4.2): a berth these ships fill is unavailable to a BV, whether or not the LP dispatches its host. Only the small yacht and expedition traffic is dropped entirely.

Schedule completeness and the data gap. PortMinder publishes only forward-looking departures, so a ship's calendar can thin out or break where later itineraries had not yet been loaded. To confirm these are genuine redeployments rather than missing data, we cross-checked the five ships with visible gaps (Costa Smeralda, Costa Pacifica, MSC Magnifica, MSC Musica and MSC Orchestra) month-by-month against an independent cruise-itinerary source (CruiseMapper, 2025). For four of the five the gaps are real seasonal absences from the corridor: Costa Smeralda winters in the Canary Islands (December–March), MSC Magnifica spends the summer in the Norwegian fjords and Baltic, MSC Orchestra moves to the Aegean (eastern Mediterranean, out of Piraeus and Istanbul) from late autumn, and MSC Musica repositions to South America over the winter — so no corridor legs are lost to the gap. The exception is Costa Pacifica, whose PortMinder feed omits a genuine winter Western-Mediterranean season (weekly loops from Barcelona, Marseille and Savona, December–March, before it repositions to the Aegean in spring); CruiseMapper independently confirms this presence, which we reconstruct from the operator's own published departure listing (Table 8; embark hubs only, mid-cruise days unmodelled). No other ship shows corridor presence hidden by a data gap.

Out-of-corridor OPS. The onboard packs are credited for charging at the thirteen corridor hubs and at EU and EEA ports outside the corridor where AFIR — or, for the Norwegian fjords, a national zero-emission mandate from 2026 — will oblige the same ∼20 MW OPS. We take a deliberately narrow view of where that headroom is real: only mainland, strong-grid ports with confirmed or mandated high-voltage OPS are counted (Cat. 1); island, weak-grid and tender-only ports — Ibiza, Olbia, the Balearics, the Canaries and most Greek islands — are assumed to receive at most a ∼10 MW hotel-only connection with no charging headroom, and are excluded. Each of ship j's n_j credited calls is modelled identically, at the ship's mean out-of-corridor call dwell τ_j (capped at 14 h) and in-port charging headroom P_j^{oc}, so each delivers the same energy $\mathbf{o}_j$ — limited both by the pack (at most $\sqrt{\eta_b}\,\mathbf{E}_j$, a single charge and discharge) and by the coupler over the dwell:

$$\mathbf{o}_j \;\le\; \sqrt{\eta_b}\,\mathbf{E}_j, \qquad \mathbf{o}_j \;\le\; \sqrt{\eta_b}\,P_j^{\mathrm{oc}}\,\tau_j. \tag{21}$$

The credit is this per-call delivery summed over the ship's n_j calls (Eq. (3)). Which cap binds is what makes it informative: a single call charges only 50 MWh to 70 MWh (headroom × dwell), well under a corridor-sized pack, so the year-round core ships — which in any case make no out-of-corridor Cat. 1 calls — would be headroom-limited, whereas a seasonal ship that would otherwise carry *no* pack hits the pack cap, so its credit rises with $\mathbf{E}_j$ and the solver builds it a small pack purely to harvest this OPS. The 14 h cap on τ_j merely stops a multi-day idle layover from being mis-scored as one continuous session.

The pack's weight-fuel is charged separately, over every out-of-corridor leg. While the credit accrues only at the n_j Cat. 1 calls, the pack is dead weight on *all* of the ship's non-corridor sailing — Cat. 1 and Cat. 2 calls and legs wholly outside the EU alike — and that mass burns host fuel whether or not the call earns OPS. We read each dispatched ship's full published itinerary and, for every consecutive pair of calls not both in the corridor, take the at-sea time (the schedule gap less a port-manoeuvring buffer, clipped at ten days to guard data gaps) at a nominal 18 kn service speed and size the leg's propulsion by the same anchor-normalised cube law as the corridor legs (§3), giving a dimensionless θ_l. Summing over the ship's out-of-corridor legs gives $\theta_j^{\Sigma} = \sum_l \theta_l$ (all legs, carrying the fuel cost $\hat{p}_{\mathrm{f}}$) and $\theta_j^{c} = \sum_l s_l\,\theta_l$ (carrying the ETS-plus-FuelEU cost $\hat{p}_{\mathrm{c}}$), where the EU maritime carbon scope $s_l \in \{1, \frac{1}{2}, 0\}$ is set by how many of the leg's two endpoints lie in the EU/EEA. Fuel is thus charged on every out-of-corridor leg and the policy stack only on its EU-scoped share (EU ETS island and outermost-region derogations are not applied — a small, conservative simplification). This matters most for ships that leave EU waters for a season: Musica's and Divina's South-American winters haul the pack for months at fuel cost but little carbon cost, and Musica earns *no* Cat. 1 credit at all.

The credit rate μ is the onboard fleet's realised corridor margin — avoided value net of charging cost per delivered MWh:

$$\mu = \frac{\sum_{j\in\mathcal{J}} \Big(\sum_{l\in\mathcal{L}_j} \hat{p}\,\mathbf{d}_{j,l} - \sum_{h\in\mathcal{H}_j} \mathbf{c}_{j,h}\,\pi_{z_j(h),h}\,\delta \Big)}{\sum_{j\in\mathcal{J}} \sum_{l\in\mathcal{L}_j} \mathbf{d}_{j,l}}. \tag{22}$$

As an output of the solve, μ is held fixed within each solve so the program stays linear; an outer loop re-solves with μ reset to its realised value until the two agree (within 3 €/MWh), converging in about one pass because μ tracks fuel, carbon and charging prices and is nearly independent of out-of-corridor volume. Table 8 lists, per ship, these out-of-corridor Cat. 1 calls (columns *+EU/+EEA*) alongside its corridor legs. The credit concentrates on the seasonal hosts (Magnifica, Euribia, Sinfonia, Orchestra), whose out-of-corridor Cat. 1 presence can rival or exceed their corridor presence.

Table 8: Corridor fleet pool (PortMinder schedules projected onto 2025), with per-ship out-of-corridor OPS. *Stops*: corridor berth calls. *Legs*: directed overnight corridor hops (one fewer than stops per continuous voyage-chain). *+EU/+EEA*: calls at out-of-corridor Cat. 1 ports (mainland strong-grid EU/EEA ports with confirmed or AFIR/fjord-mandated ∼20 MW OPS). *Hotel*: in-port load (15 MW large, 10 MW small). "—" marks no material out-of-corridor Cat. 1 presence.

Ship	**Hotel**	**Legs**	**Stops**	**+EU**	**+EEA**	**Typical corridor loop**
MSC World Asia	15	261	312	—	—	BAR→MAR→GEN→NAP→MES→MLT
Costa Toscana	15	230	283	—	—	MAR→BAR→TUN→PAL→CIV→SAV
MSC Sinfonia	10	136	148	62	0	BAR→MAR→GEN
Costa Smeralda	15	133	165	—	—	MAR→BAR→IBZ→PAL→CIV→SAV
MSC World Europa	15	129	158	—	—	MAR→GEN→NAP→MES→MLT→BAR
MSC Seaview	15	126	157	—	—	GEN→CIV→PAL→IBZ→BAR→MAR
MSC Splendida	15	111	133	22	0	BAR→MAR→LIV→CAG→PAL→MLT
MSC Euribia	15	100	126	59	23	BAR→MAR→GEN→NAP→PAL→TUN
MSC Musica	10	87	110	—	—	CIV→GEN→MAR→VAL→IBZ→CAG
Costa Pacifica	15	87	137	6	0	CIV→SAV→TOU→PLM→VAL→OLB
MSC Orchestra	10	78	98	48	0	CIV→GEN→MAR→TAR→VAL→LIV
Norwegian Epic	15	76	103	17	0	MAR→BAR
Costa Fascinosa	15	62	85	50	0	SAV→BAR→MAR
MSC Virtuosa	15	49	62	53	5	MAR→BAR→TUN→PAL→NAP→LIV
MSC Magnifica	10	47	61	94	22	MAR→BAR→GEN
Costa Favolosa	15	37	50	26	23	BAR→MAR→SAV
MSC Divina	15	29	57	—	—	CIV→GEN→ $\cdots$ (repositioning)
MSC Grandiosa	15	20	40	—	—	BAR→CAN→GEN→SPE→CIV→PLM
MSC Opera	10	16	28	48	0	MAR→MAL→CAD→LIS→ALC→MAH→ $\cdots$
Total		**1814**	**2313**			

Loop = the ship's modal weekly corridor circuit, closing back to the first port. Port codes: BAR Barcelona, MAR Marseille, GEN Genoa, SAV Savona, CIV Civitavecchia, NAP Naples, PAL Palermo, CAG Cagliari, MES Messina, MLT Valletta, TUN Tunis (La Goulette), IBZ Ibiza, VAL Valencia, TAR Tarragona, LIV Livorno, TOU Toulon, PLM Palma, OLB Olbia, CAN Cannes, SPE La Spezia, MAL Málaga, CAD Cádiz, LIS Lisbon, ALC Alicante, MAH Mahón.

Direction switching in the adjusted schedule. Reversal eligibility is open to every marked BV host (§2.6.3); the route optimiser reverses a typically-weekly circuit only where doing so raises the *joint* fleet's delivered energy, so the chosen set depends on the number of vessels fielded. Table 9 reports the per-host switches at the profit-optimal **four-vessel** fleet: 150 of 530 typically-weekly circuits are reversed, spread across seventeen of the nineteen hosts. Two of the three one-directional hosts (†) are reversed well above their share, but reversal helps broadly — wherever a return-direction circuit fills a gap a vessel would otherwise sit through — and at this optimum two hosts (World Europa and Orchestra) are left as-sailed.

Per-ship division of labour. Figure 6 reads the two profit-max solves ship by ship, ordered by each ship's corridor at-sea demand — the year-round proxy. Onboard packs are taken up *widely* (sixteen of the nineteen ships carry one at the onboard-only optimum), but their delivered energy concentrates on the near-year-round ships: the two year-round hosts (World Asia and Costa Toscana) supply just over half of all onboard delivery, the top four about three-quarters. Battery vessels invert this. Their share of a host's delivered energy rises monotonically as the host becomes more seasonal — from ∼28% on World Asia and 50% on Costa Toscana, through ∼60% on the mid-season ships, to 80–100% on the seasonal tail (Euribia, Virtuosa, Favolosa, Musica, Divina, Opera). The two architectures thus partition the fleet by schedule rather than compete for the same hosts.

Table 9: Circuit direction in the adjusted schedule: typically-weekly circuits reversed per host, of the total the host sails, under the joint-fleet energy-maximising rule of §2.6.3 at the profit-optimal four-vessel fleet. † sails the corridor in one direction only in its real schedule. The *of which CW→CCW* column splits the reversals by resulting compass direction: reversing a clockwise circuit opens counter-clockwise (return-direction) legs and vice versa. The per-circuit compass sense is not stored, so we approximate at host level — the three one-directional (†) hosts sail counter-clockwise, so their reversals are CCW→CW; the remaining sixteen are taken as CW→CCW (132 of 150).

Host	**Circuits**	**Reversed**	**of which CW→CCW**
MSC World Asia	52	26	26
Costa Toscana	53	17	17
Costa Fascinosa	27	14	14
MSC Sinfonia	31	13	13
MSC Euribia†	26	13	0
Costa Pacifica	50	11	11
MSC Musica	23	10	10
MSC Grandiosa	20	10	10
Costa Smeralda	32	8	8
MSC Seaview	31	6	6
Costa Favolosa†	17	5	0
MSC Divina	28	4	4
Norwegian Epic	28	4	4
MSC Magnifica	14	4	4
MSC Opera	13	2	2
MSC Virtuosa	13	2	2
MSC Splendida	22	1	1
MSC World Europa†	30	0	0
MSC Orchestra	20	0	0
Total	**530**	**150**	**132**

C Cost parameter derivations

This appendix derives each cost parameter of Table 1 and gives its sources. Three power-side costs sit at different points of the link and should not be confused: the *port coupler* is the shore charger at the berth; the *host stern interface* is the aft power connection built on the host to receive a BV under way; and the *power-handling upgrade* is the drawing-side receiving switchgear — the *host flank interface* on an onboard ship taking charge from the port coupler, and the *BV power build* on the BV itself.

Port coupler capital (C_r). A port coupler is $R = 20\,\text{MW}$ of port-side capacity: a shore-side charger (rectifier, converter, switchgear) at $c_{\text{chg}} \approx$ €650/kW (Mackenbach et al.) plus a grid connection (substation and feed) at $c_{\text{grid}} \approx$ €200/kW (Green Cruise Port, 2018), *both* on the full nameplate — the sunk coupler carries its whole 20 MW in the model, a berthed host drawing its hotel load *plus* its charge headroom through it with no build. Together $C_r = c_{\text{chg}} + c_{\text{grid}} \approx$ €850k/MW, so $\gamma C_r R \approx$ €1.36M/yr annualised over $L = 20\,\text{yr}$ at $i = 5\%$ ($\gamma = 0.080$/yr), roughly three-quarters charger and one-quarter grid connection. The *coupler-substitution* analysis (§2.6.2) prices a whole such coupler a port could *avoid building*; the endogenous *port coupler build* instead *adds power* to the sunk couplers, at the same C_r per added MW.

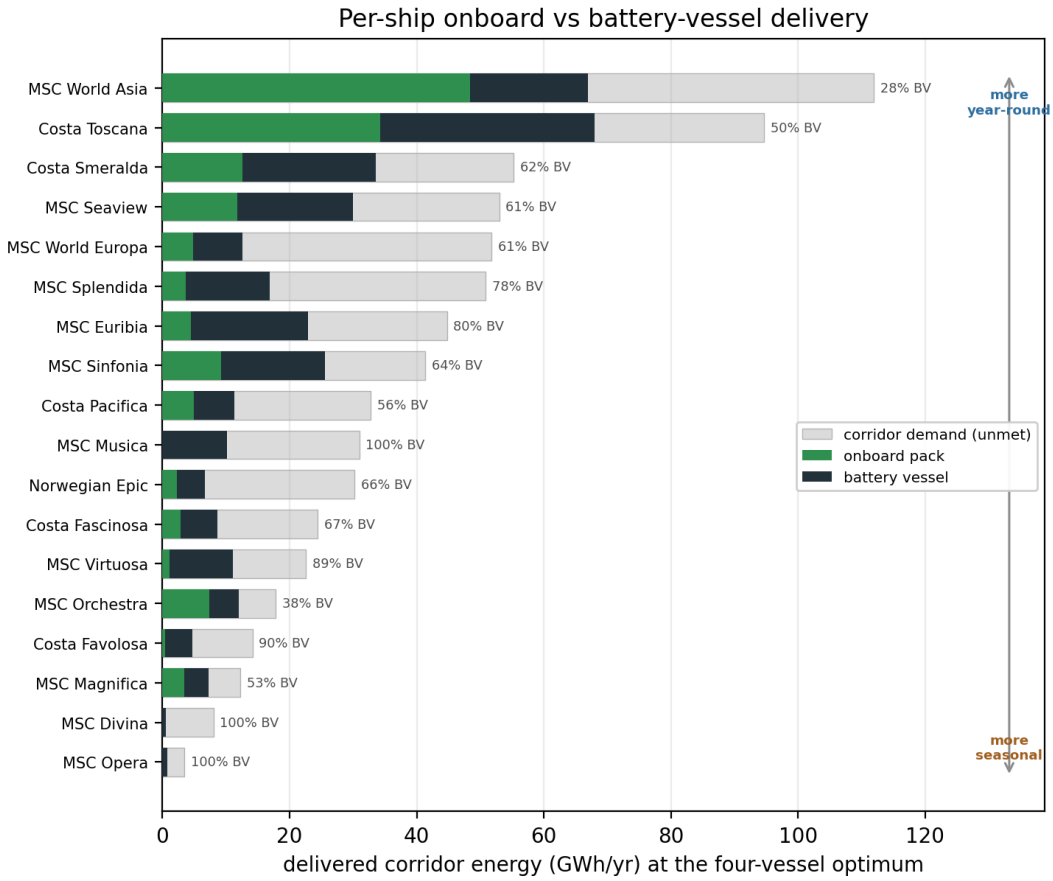


Figure 6: Per-ship onboard versus battery-vessel delivery at the profit-max four-vessel optimum, one row per host, ordered by the ship's corridor at-sea demand (the year-round proxy; most year-round at top). Green is energy delivered by the ship's own onboard pack, dark by battery vessels; the faint bar is the ship's full corridor demand, so the electrified share and the unmet remainder read off directly. Labels give each host's battery-vessel share of delivered energy. Onboard delivery concentrates on the near-year-round ships (top); battery vessels carry a rising majority down the seasonal tail.

Host stern interface (I^{iface}). The host-side interface is a *stern* (aft) power connection — aft inlet, switchgear and electrical uprating to accept the coupling — rated to the *peak power delivered to that host*, $P_j = \max_{l \to j} \mathbf{d}_l / h_l^{\text{sea}}$ — the underway hotel-plus-propulsion power (delivered energy over the coupled underway hours h_l^{sea}), floored at the 20 MW anchor the connection is assumed to cover. The propulsion component applies only while under way, so the peak is set by the highest-power leg served, not a pack-size proxy. We price it as a flat anchor for that covered power plus a *linear* uprating above it, $I_{\text{ship}}(P) = €1.5\text{M} + €50/\text{kW} \cdot (P - 20\,\text{MW})^+$. The €1.5 M anchor is Green Cruise Port's €0.5 M for a 10 MW hotel connection scaled up and increased for the novel stern coupling and propulsion-bus uprating; the €50/kW rate above it is the *same* per-MW figure as the power-handling upgrade below. Each host j carries its own stern receptacle $I_j^{\text{iface}} = I_{\text{ship}}(P_j)$ sized to the power delivered *to it*; assuming every host in the fleet $\mathcal{J}$ has one, the fleet interface cost is $\sum_{j \in \mathcal{J}} I_j^{\text{iface}}$, as in the objective (3). Higher delivery power scales the interface *up* from the 20 MW anchor, with no cap. The cost is annualised at $\gamma = 0.080$/yr. Because no catalogue price yet exists, a €2 M anchor is carried as a sensitivity.

Power-handling upgrade: host flank interface and BV power build (C^{upg}). These two drawing-side upgrades share one marginal rate with the stern interface — €50/kW for delivery power above 20 MW — at *both* ends of the transfer; they differ only in the base. The host stern interface is a whole fitting built from nothing, so it carries the €1.5 M anchor for its first 20 MW. The drawing side does not: its 20 MW already exists — the AFIR coupler for an onboard ship's flank interface, the BV's own power electronics (part of its flat vessel cost) for the BV power build — so only the increment above 20 MW is new. An onboard ship pays $C^{\text{upg}} = €50$/kW for charge power above its 20 MW coupler, and, symmetrically, the BV pays the same €50/kW on its peak power — max of peak charge and peak at-sea discharge — above 20 MW. Keeping the increment linear also serves the solver: the upgrade increment is a decision variable, so a nonlinear rate on it would forfeit the LP, whereas the interface is sized off the *solved* delivery post-hoc.

Battery benchmark (K_b). The round-trip efficiency is $\eta_b = 0.86$. The battery pack capital K_b (€per usable MWh) is the pack's nameplate purchase price spread over its usable depth-of-discharge, and it enters the objective annualised by the pack's own capital-recovery factor γ_b:

$$K_b = \frac{p_{\text{pack}}\,(€/\$)}{\text{DoD}}, \qquad \gamma_b = \text{CRF}(i, L_b) = \frac{i\,(1+i)^{L_b}}{(1+i)^{L_b} - 1}, \tag{23}$$

the capital-recovery (annuity) factor γ_b carrying both the time value of money at real discount rate i and, through the finite life L_b, the pack's replacement cost; no residual or second-life value is credited at retirement, consistent with Mackenbach et al. (2026) and conservative for a pack that leaves service at 80% state of health. We take the pack price from Mackenbach's maritime cell-cost forecast at the **2040** deployment horizon — **$87/kWh**[3] — at €0.92/$, a usable depth-of-discharge DoD = 0.80, and a real discount rate $i = 5\%$. The usable window and the life both reflect the LFP/sodium-ion chemistries an unmanned pack can adopt, and both are modest departures from Mackenbach's generic-pack inputs: we take DoD = 0.80 where Mackenbach adopts 0.70 (a 20–90% swing), since LFP and sodium-ion tolerate deeper cycling. Service life is degradation-set rather than assumed: with a per-cycle fade of 3.75×10^{-5} (DNV, 2024), the ∼150–200 cycles/yr both fleets run reach the 20% end-of-life fade only after ∼5,300 cycles, so the pack is *calendar*-limited, not cycle-limited. We adopt a central calendar life $L_b = 12$ yr for LFP/sodium-ion: longer than the ∼8–11 yr Mackenbach's generic (NMC-derived) fade implies at this cycling, reflecting the slower calendar ageing of these chemistries. Equation (23) then gives a central pack capital $K_b \approx €100$k/MWh usable, annualised to $\gamma_b K_b \approx$ **€11,300/MWh/yr** (equivalently €30.9/MWh/day), with a **€10,000–13,300/MWh/yr** band spanned by a 3–8% discount rate — the residual capital-market uncertainty once price, chemistry and life are fixed. Table 10 samples this band with $\gamma_b K_b = 10.2$ and 13.1 k rows inside it, and brackets it more widely still with $\gamma_b K_b = 8.0$ and 15.0 k rows reaching beyond the discount-rate range, alongside the round-trip-efficiency and self-propulsion levers that move the same pack economics.

Because the pack is calendar-limited, cycling adds no marginal cost to the objective, and since the delivery target (12) pins annual throughput the cycle count is not a free decision; the linear profit objective (3) is therefore unaffected by cycling (save the grid-arbitrage variant, which adds endogenous cycles) and no state-of-health variable is carried in the LP. Cycling would have to rise roughly 2.4× before cycle life displaces the 12 yr calendar limit, and the throughput-based fade coefficient's neglect of SOC-window, C-rate and temperature effects is second-order in this regime.

[3] The 2040 point of the maritime forecast in the supplementary calculations accompanying Mackenbach et al. (2026), sheet `calc.3.Battery`, matched to the 2035–2045 horizon over which the present-day stack of §2.5 is argued near-flat; the 2035 value is $112.

BV self-propulsion (κ_{BV}), derived from pack mass. A BV must propel *itself* between hosts. The power to follow a host at transit speed scales with displacement as $(W_{\mathrm{BV}}/W_{\mathrm{host}})^{0.7}$ (Admiralty/Stopford; IMO, 2020), so on any leg the BV's self-propulsion is that *same fraction of the host's propulsion energy on the leg* — the very quantity the BV is there to displace. Hotel load does not enter: a larger pack changes the vessel's weight, hence its propulsion, but not a host's hotel draw. We build the BV mass bottom-up at a 2040 basis, $\rho_{\mathrm{BV}} \approx 9\,\mathrm{t/MWh}$ at system level: a $\sim$3 t/MWh cell pack at the 2040 energy density (consistent with the 2040 pack *price* we adopt) plus $\sim$3 t/MWh of balance-of-system, scaled up by $\approx$1.5 for a small purpose-built hull; this sits below the 11 t/MWh present-day system low end (MAN Energy Solutions, 2019, Table 2) and well below the 19 t/MWh of passenger-certified packs (He et al., 2024), whose fire/thermal structure an unmanned hull avoids. A BV of usable capacity $\mathbf{E}$ carries a nameplate pack of mass $\rho_{\mathrm{BV}}\mathbf{E}/\mathrm{DoD}$; anchored on the large host ($W_{\mathrm{host}} \approx 90\,\mathrm{kt}$), at a representative profit-max pack ($\approx$300 MWh usable, 375 nameplate — a single calibration size; the four vessels of the base optimum size differently around it, §4, but κ_{BV} is held fixed across them, below) the mass ratio gives $(W_{\mathrm{BV}}/W_{\mathrm{host}})^{0.7} \approx 0.10$ — **self-propulsion $\approx$9.9% of the host's leg propulsion**. On the *anchor leg* — the large host's representative corridor crossing, $\approx$196 MWh of host propulsion, against which each leg's self-propulsion is normalised via θ_l — that is $\approx$19.5 MWh, i.e. $\approx$6.5% of the usable pack. We enter it linearly in the SOC balance (Eq. (9)) as $\kappa_{\mathrm{BV}}\,\theta_l\,\mathbf{E}_m$, with θ_l the leg's host propulsion relative to the anchor and $\kappa_{\mathrm{BV}} \approx 0.065$ (of usable capacity) fixed at the profit-max size. The mass fraction varies only as $\mathbf{E}^{0.7}$ — from $\sim$8% at 250 MWh to $\sim$12% at 500 — so holding κ_{BV} constant is a mild, size-robust linearisation (no fixed-point iteration), conservative in over-charging self-propulsion as the pack grows above the calibration size, with 0.13 carried as an upper sensitivity (Table 10). On the reduced-energy small-class legs (§2.1) — similar distance, less delivered — the same reserve is a somewhat larger share of that leg's smaller delivery. At-berth supply (§2.6.2) carries no such reserve, the BV being stationary.

Onboard weight penalty. An onboard pack of installed mass $m = \rho_{\mathrm{inst}}\mathbf{E}_j/\mathrm{DoD}$ is hauled at sea by its host, so the host burns more under way. Unlike the BV reserve this is a *marginal* uplift on a $\sim$90 kt hull, not a whole small vessel: by the same displacement$^{0.7}$ scaling the extra propulsion is $0.7\,m/W_{\mathrm{host}}$ per leg, so at a representative onboard pack (the size the LP builds at the profit peak) it is only $\approx$3.1 % of the host's leg propulsion — less than a third of the BV's $\approx$9.9 % (§2.3), the price of putting the battery on its own small hull rather than piggybacking a cruise ship. We build the mass bottom-up: a passenger-certified marine system is $\sim$19 t/MWh today (He et al., 2024; $\sim$8 t/MWh cell at $\sim$125 Wh/kg, plus $\sim$11 t/MWh balance-of-system), and by 2040 the cell fraction lightens with energy density (Mackenbach, $\sim$285 Wh/kg$\rightarrow\sim$3 t/MWh, consistent with the 2040 pack *price*) while the balance-of-system shrinks modestly to $\sim$8, giving $\rho_{\mathrm{inst}} \approx 11\,\mathrm{t/MWh}$. It enters as an extra per-leg pack draw $\kappa_{\mathrm{onb}}\theta_l\mathbf{E}_j$ ($\approx$1.1 % of propulsion per 100 MWh of *usable* pack) flowing through the energy balance as extra *charging* on corridor legs, self-limiting on pack size. The pack is also hauled on the host's *out-of-corridor* legs, burning host fuel on *every* such leg (not only the credited Cat. 1 calls); this weight-fuel is charged over the ship's full itinerary as constructed in §2.4.2 and App. B.

Onboard passenger-hull premium (c^{prem}). A bundled per-nameplate-kWh capital cost dominated by the passenger fire/thermal certification premium (A-60 barriers, redundant suppression and venting) an unmanned BV — which can seal-and-flood — escapes, plus space opportunity and other passenger-hull costs. MAN Energy Solutions (2019) put the marine non-cell system cost at \$250–300/kWh, the majority safety and thermal; against that anchor we adopt **\$50/kWh**

nameplate as a *representative engineering allowance* for the slice a passenger hull carries and a BV avoids — a judgement figure rather than a directly observable cost — annualised at γ and €0.92/$, and expressed per usable MWh (dividing by DoD) as a flat **€4610/MWh/yr** rent on onboard pack capacity. Because it is not directly measurable and is chemistry-dependent, we do not rest the conclusions on the point value but sweep it in the sensitivities; the space component is held at zero but can be considered a part of the premium.

Vessel cost (K^{BV}, F^{OM}). Unlike an onboard pack, which rides an already-crewed ship, a BV is a *dedicated* platform whose whole cost the electrification must carry; we charge a flat per-vessel cost, independent of pack size. The all-in capital — excluding the pack and the coupling interface — is taken as $20 M: a specialised hull scaled from the Stopford (2009) newbuild curve (as in Wiskich, 2025), plus BV-side power electronics and autonomy, cross-checked against Nazir's (2021) moored charging platform. This is a deliberately *top-down* figure for an *unmanned*, single-purpose vessel that carries none of the accommodation, lifesaving or comfort systems of a crewed ship, read as an at-scale (series-build) cost for the 2040 horizon. It sits at the *generous* end of the benchmarks: the autonomous, electric *Yara Birkeland* was built for a reported ∼$25M *including* its battery — of order $15–20 M ex-pack for a smaller vessel that additionally carries cargo-handling gear a BV does not (International Transport Workers' Federation, 2018, citing the *Wall Street Journal*); battery-bunker-vessel concepts point the same way (ABB, 2024). The $20 M hull (∼€18.4M at €0.92/$) is the raw capital K^{BV}, annualised at γ to **∼€1.48M/yr**; unmanned operations (insurance, berth dues and maintenance at ∼$3000/day, below the ∼$5000/day labour-and-insurance allowance of the RETREC calibration (Wiskich, 2025)) are a separate annual flow F^{OM} **≈€1.02M/yr**. The two together are **∼€2.5M/yr** per vessel — the hull capital K^{BV} and fixed O&M F^{OM} of the objective (3). The flat, pack-independent form is what makes a BV economic only once its pack is large enough to dilute it, so BV adoption begins at scale while the onboard fleet, carrying no platform cost, electrifies incrementally. A $25 M variant (∼€1.85M/yr capital, ∼€0.4M/yr per vessel more) trims the four-vessel peak by of order €1.5M without moving the optimal count — small beside the regulatory-year lever, and carried as a row in Table 10; it also stands in for a higher risk premium on a first-of-kind, unmanned platform than the common 5 % rate — applied uniformly here, including to sunk shore infrastructure — would imply.

D Single-lever sensitivity detail

This appendix tabulates the single-lever sensitivities summarised in Figure 5

Table 10: Single-lever sensitivities on the base. Each row re-solves the fleet to its own profit-max delivery and battery-vessel count (0–6), reporting the profit-maximising configuration — its profit, delivery, optimal number of BVs, and the share of that delivery the BVs carry — beside an onboard-only fleet re-optimised at its own optimum. The *Fleet density* rows restrict electrification to one owner, delivery measured against that owner's own potential (Costa 221, MSC 453 GWh).

Group	**Lever**	**Profit-max BV fleet**				**Onboard-only**	
		Profit (€M)	Deliv. (GWh)	#BV opt	BV (%)	Profit (€M)	Deliv. (GWh)
Base	Current, today's stack, LNG	+23.5	349	4	57	+13.4	254
FuelEU pool (base €200)	Pool €120/t	+6.1	225	3	61	+5.2	70
	Pool €150/t	+12.1	299	4	63	+6.9	89
	Pool €250/t	+36.8	406	5	57	+23.3	307
	Pool €300/t	+52.1	431	5	53	+34.6	332
Avoided-cost stack	Existing fuel use (mixed)	+28.3	357	4	56	+16.4	281
	ETS €130/t (Enerdata)	+36.8	402	5	57	+23.5	302
Electricity year	2024 prices	+16.1	312	4	62	+8.2	197
	2023 prices	+23.8	346	4	56	+14.2	252
Charging cost	Adder €10/MWh	+19.1	328	4	60	+10.2	222
	Adder €20/MWh	+14.8	311	4	62	+7.5	151
	Adder €30/MWh	+10.7	293	4	64	+5.8	85
Onboard premium	$c^{\text{prem}} = \$25$/kWh	+26.8	369	4	51	+18.6	294
	$c^{\text{prem}} = \$75$/kWh	+20.9	329	4	63	+9.2	213
Battery cost	$\gamma_b K_b =$ €22/MWh/day	+33.0	402	5	57	+21.1	311
	$\gamma_b K_b =$ €28/MWh/day	+26.3	356	4	55	+15.6	277
	$\gamma_b K_b =$ €36/MWh/day	+19.0	329	4	60	+9.9	217
	$\gamma_b K_b =$ €41/MWh/day	+15.1	302	4	63	+7.1	147
Operational	Self-propulsion $\kappa = 0.13$	+18.7	332	4	48	+13.4	254
	Round-trip eff. $\eta_b = 0.90$	+26.8	360	4	55	+16.3	278
	No delivery-only spurs	+21.6	343	4	56	+12.8	253
	Sea-state gate $H_s \leq 2.5$m	+21.0	341	4	53	+13.4	254
BV capital	Hull \$25 M	+22.0	349	4	57	+13.4	254
	Host interface €2 M	+22.8	349	4	57	+13.4	254
Fleet density	Costa only, current (5 hosts)	+2.1	25	0	0	+2.1	25
	Costa only, adjusted	+5.7	114	1	50	+2.1	25
	MSC only, current (13 hosts)	+8.8	136	2	49	+7.3	71
	MSC only, adjusted	+13.6	208	2	56	+7.3	71
Port build	No endogenous coupler build	+15.7	208	4	64	+9.0	77
Upside	Grid arbitrage	+29.8	402	6	62	+13.6	253
	Coupler substitution	+25.5	349	4	57	+13.4	254
	Adjusted schedule	+30.2	432	4	57	+13.4	254

E Coupler substitution: marginal-coupler utilisation

This appendix reports the corridor-wide coupler-substitution screen (§2.6.2). Each hub's built coupler count K (the K_r of the base model) is the smallest that holds served berth-days at the AFIR ∼90 % floor, on the daily peak simultaneous major-cruise occupancy. The marginal (K-th) coupler lifts coverage from just below the floor to just above it, so it is used only on the rare peak-occupancy days — the *last-coupler days* (occupancy reaching K) — yet costs as much as any other. The screen asks what happens if the port builds one fewer ($K-1$) and leans on the routed 4-BV fleet to backfill that rare duty. A first pass keeps only hubs where the fleet dwells on enough last-coupler days to restore the floor; the thin and two-coupler hubs (Valletta, Palma, Cagliari, Savona, Valencia, Messina and Palermo) fail here — either the route never brings a BV to their peak days, or, at a two-coupler hub, removing the only spare leaves no free coupler for the BV to charge on. Each surviving candidate is then *re-solved* with its coupler removed, so the LP reprices the tighter charging. The **marginal net** in Table 11 is the re-solved *operating-profit change* — the value of any berth charging the fleet itself gains or loses when the coupler goes — plus the €1.36M/yr avoided coupler capital (Eq. 20), credited only where it is positive *and* coverage clears 90 %. To this we apply one further, post-hoc deduction the fleet LP does not carry, scored hour-by-hour rather than by whole days: on the all-couplers-busy hours when the routed BV is *absent* — about a quarter of them at Livorno and over two-fifths at Naples — the host the removed coupler would have served runs its hotel load off gensets rather than shore power, a lost-electrification cost we value at the corridor avoided-cost margin (§2.5), €0.15M/yr at Livorno and €0.39M/yr at Naples. Each is sized on the hub's class-mix of bumped hosts. Removals are taken greedily, one hub at a time, so each is judged against those already removed; a final berth-energy check confirms the covering BV holds enough charge, read from its solved state-of-charge, to power a 15 MW host call on the days it stands in. The all-or-nothing coverage of Eq. (19) is not imposed but *verified*: on every sub-need day the metric counts covered, the BV is present for all the day's all-couplers-busy hours (Livorno 16/16, Naples 20/20; its berth is a continuous block spanning the congested window), so no day is credited on partial presence.

Table 11: Coupler-substitution candidates: the hubs where the congestion-routed 4-BV fleet dwells on enough last-coupler days to test removal, each re-solved with its coupler removed. K is the port's built coupler count; "Cov. K" is served-berth-day coverage as built (smallest integer K clearing the 90 % floor); "Cov. $K-1$" removes one coupler with no BV; "Cov. +BV" adds the energy-gated BV backfill. "Marginal net" is the port's contribution to fleet net — the re-solved operating-profit change plus €1.36M/yr avoided capital, with Naples valued after Livorno's removal (greedy order). A coupler comes out only where this is positive *and* coverage clears 90 %; the Outcome column records the verdict, discussed below.

Hub	K	Last-coupler days	Cov. K	Cov. $K-1$	Cov. +BV	Marginal net (€M)	Outcome
Livorno	**3**	**20**	**0.989**	**0.885**	**0.969**	**+1.37**	**remove**
Naples	**3**	**35**	**0.981**	**0.843**	**0.921**	**+1.15**	**remove**
Barcelona	4	88	0.938	0.823	0.896	+0.45	keep: coverage
Marseille	4	48	0.953	0.873	0.950	−0.76	keep: charging
Genoa	2	37	0.992	0.857	0.952	−1.81	keep: charging

Removing the marginal coupler at both qualifying hubs together costs only €0.2M/yr of fleet operating profit — the tighter charging, from the joint re-solve — against €2.72M/yr of avoided

coupler capital (two 20 MW couplers at C_r, annualised; Eq. 20). Netting the €0.5M/yr of host calls left un-electrified on days no BV backfills leaves a **net €2.0M/yr** (Livorno €1.21M, Naples €0.76M). At both hubs the covering BV's solved state-of-charge (pack ∼300 MWh) clears the ∼180 MWh a 15 MW half-day call needs on all but one backfilled day per hub. Marseille and Genoa carry too much of the fleet's charging, so removal costs €0.8–1.8M/yr more than the avoided capital, and Barcelona slips to 89.6 % once both qualifying hubs draw on the same four vessels.

F Data sources and treatment

This appendix expands the inputs of Table 3: how the price, fuel and sea-state series are sourced and treated, and the charging-price and zero-intensity-credit discussions.

Electricity prices. The price series $\pi_{z,h}$ are realised day-ahead hourly prices for the seven bidding zones, taken from the ENTSO-E Transparency Platform and resampled to a clean hourly UTC index over the calendar year (rare gaps forward/back-filled).

The charging price the BV faces. Price structure reaches the BV only through its *charging* cost — the at-sea displacement value it earns is set corridor-wide by fuel, carbon and the regulatory stack, essentially independent of zone (§2.5). On the charging side the zonal figures (Table 4) show Spain and France consistently cheaper and more volatile than the five Italian-priced zones. The western within-year standard deviation is €40–48/MWh, half or more of the mean, against €30–40 for the Italian zones, and Spanish and French prices increasingly fall at or below zero in solar-rich hours. The three extended-hub zones price near the Italian mainland (€110–125), adding corridor reach but no deeper trough. A roaming BV charges preferentially in the western troughs, but charging is a small fraction of the avoided value, so this is secondary to the throughput the schedule allows. These are moreover *wholesale* day-ahead prices, excluding network and transmission tariffs, taxes and levies: a real charger pays those adders, so our charging cost is a favourable lower bound and the breakevens correspondingly optimistic, with flat network adders carried as sensitivities.

Marine fuel and carbon. The avoided-cost inputs are realised annual Mediterranean bunker prices (LNG base, MGO sensitivity) and EU allowance (EUA) prices, the EUA series a traded annual average. We use an annual-average bunker price rather than a daily spot series by design: cruise operators typically buy fuel on forward contracts and refuel only every few weeks, so the cost a host actually faces is a smoothed price, not the daily spot, and finer granularity would add noise without changing that smoothed cost. The avoided-cost stack is in any case dominated by its regulatory components (§2.5), so the headline conclusions do not hinge on the fuel-price input.

The zero-intensity credit. The avoided-cost parameters of Eqs. (14)–(15) are set in §2.5 and Table 4. One choice deserves emphasis: we credit battery-delivered energy at *zero* well-to-wake carbon in the host's FuelEU compliance balance, so it earns the full host intensity. This follows the regulation, not a grid-intensity assumption: FuelEU already applies a zero Well-to-Tank factor to electricity supplied to a ship *at berth* and treats onboard battery storage as a zero-emission technology (Annex III; European Commission, 2024). Extending the same zero-intensity treatment to energy transferred *at sea* is a modest extension of a mechanism the regulation already contains for at-berth electrification, not a wholly uncontemplated case. The credit is therefore independent of the charging grid's *physical* carbon intensity, which is priced separately through the wholesale power price the battery pays; and the asset's own operation makes that physical intensity low in

any case — it charges in cheap, renewable-rich hours (concentrated in the solar-rich midday window; Fig. 7a), can be contracted to a renewable power-purchase agreement, and draws power whose marginal intensity falls as the European grid decarbonises.

Significant wave height. Hourly H_s for the operability gate (Eq. (2)) is taken from the Open-Meteo Marine reanalysis, cached per year, at one representative open-water point per sea segment — seven in all: Gulf of Lion (42.3°N, 3.8°E), Ligurian (43.4°N, 7.8°E) and Tyrrhenian (42.9°N, 10.2°E) on the western spine, and the Balearic (39.2°N, 1.5°E), southern-Tyrrhenian (39.8°N, 13.5°E), Sicily-Channel (36.5°N, 14.5°E) and Sardinian (39.6°N, 10.5°E) basins for the central-Med and Balearic legs; each leg is gated on its own basin. The corridor is calm — 95th-percentile H_s is 2.5 m to 3 m on the exposed Gulf of Lion and below 2.5 m in every other basin — so the 4 m base gate blocks only a handful of legs per year (7 in 2025), mostly on the Mistral-exposed Gulf of Lion but with a few on the Ligurian and southern-Tyrrhenian crossings, whose winter tails also reach the gate. The 4 m threshold (upper sea state 5) is a conventional operability limit for offshore transfer and station-keeping work — the level up to which naval underway replenishment operates (§2.1). That precedent's transfer to a small, low-freeboard battery hull coupling to a cruise ship is uncertain, so we treat the 4 m gate as a working estimate rather than a demonstrated limit, probed by the stricter 2.5 m sensitivity of §2.1.

G Dispatch detail: charging profile and port build

Three features of the profit-max dispatch (the four-vessel base on the current schedule) support claims made in the main text but are not needed there in full.

Charging coincides with solar. The dispatch concentrates charging in the middle of the day. Because cruise calls are daytime events and a battery draws on a coupler only while its ship (or a co-docked neighbour, for the BV) is berthed, mean charging power across the corridor peaks around midday, an order of magnitude higher than overnight (Fig. 7a). This diurnal shape coincides with peak solar generation in the Western-Mediterranean price zones, so the fleet charges when zonal prices — and grid carbon intensity — are at their lowest, absorbing midday solar rather than adding to the evening peak; it underpins the zero grid-intensity credit of §2.5.

Delivery mix by host. Figure 7(b) shows *where* the delivered energy comes from at four delivery levels, contrasting the onboard-only optimum (left column of each pair) with the best-BV envelope optimum (right). At the envelope peak the BVs supply over half of corridor delivery, and the onboard contribution is pared back not by dropping a few seasonal hosts but broadly across the fleet — the roaming packs absorb a block of delivery that would otherwise be met by many lightly-cycled onboard packs, so nearly every host carries a smaller pack than it would in the onboard-only fleet. The utilisation-for-capacity substitution of §5.1 is thus visible directly in the delivery mix.

Port-capacity utilisation. The same profit-max frontier solves show how much shore-power infrastructure the corridor actually needs, and expose a distinct benefit of the BV. Figure 7(c,d) plots the fleet's *dedicated* port-coupler build (above the free port-coupler headroom) against delivery, for an onboard-only fleet (panel c) and the profit-max envelope (panel d). Up to ∼50 GWh/yr both build *nothing* — the fleet charges entirely on spare couplers, raising the utilisation of connections AFIR already funds, infrastructure that would otherwise sit idle between calls. Beyond that, the onboard-only fleet must build early and steeply, since each onboard pack has only its ∼5 MW

of coupler headroom above the hotel load (§2.4.2): total corridor build reaches ∼352 MW at 400 GWh/yr. The envelope with BVs meets the same delivery with about half that (∼186 MW), a BV occupying a whole idle coupler at its full 20 MW rating and so substituting spare-coupler utilisation for dedicated capex. Marseille (FR) builds first and hardest in both (the busiest western hub and the pivot between the cheap western zones and the high-value Italian ones), the central-Med hubs recruited last.

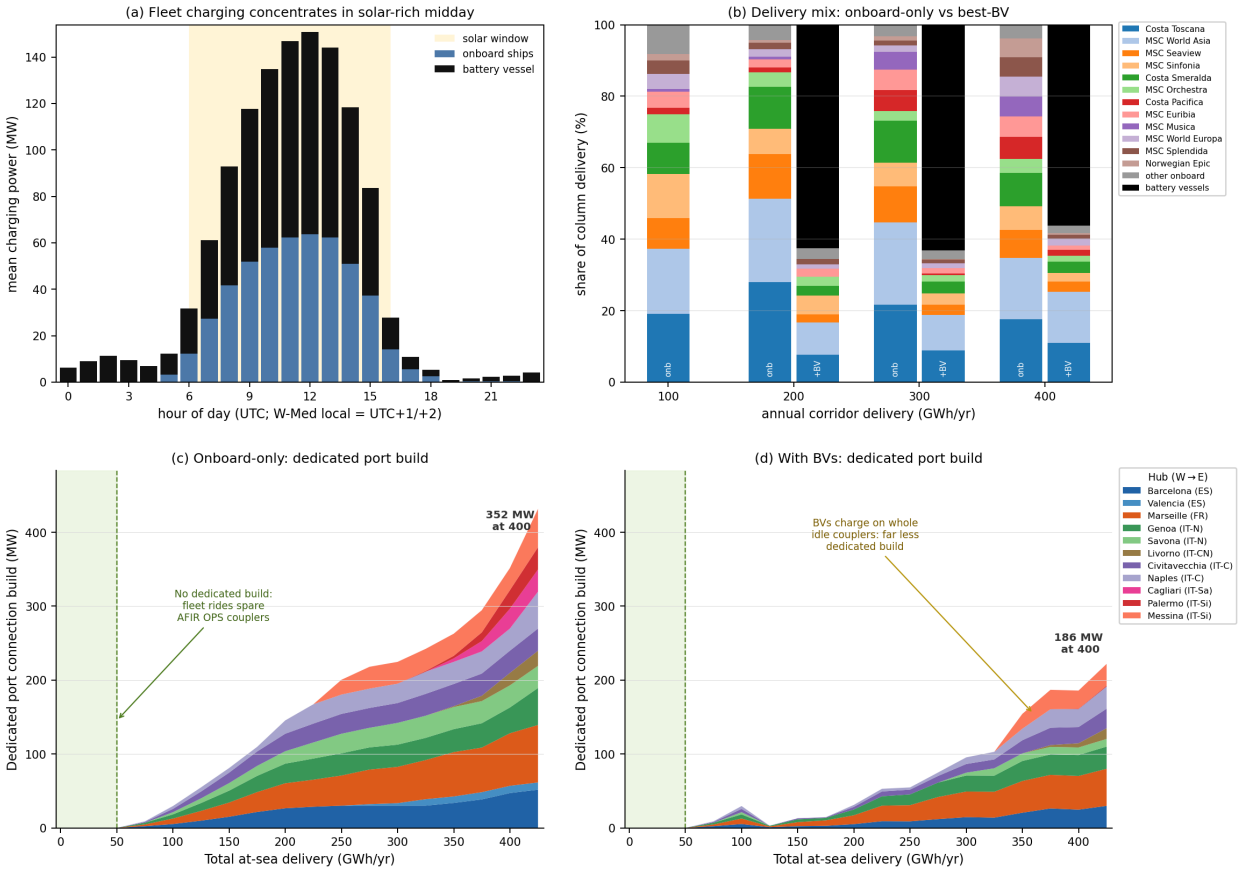


Figure 7: Dispatch detail for the profit-max four-vessel base (current schedule), read from the saved solves. *(a)* Fleet charging by hour of day: mean charging power (MW) across all corridor chargers, averaged over the year and stacked by asset — onboard ships (blue) and BVs (black); hours are UTC, so the midday peak falls near 13–14 h local, coincident with the solar window (shaded). *(b)* Delivery mix at four delivery levels, each column normalised to its own delivery: *onb* the onboard-only optimum, $+BV$ the best-BV (envelope) optimum; onboard delivery split by host (coloured, minor ships lumped grey), BVs black. *(c,d)* Dedicated port-coupler build (MW, above the free port-coupler headroom), stacked by hub, against total at-sea delivery, for *(c)* an onboard-only fleet and *(d)* the profit-max envelope; both share the shaded zero-build region where the fleet rides spare AFIR couplers.

H Coupler congestion and battery-vessel fleet size

Two model elements govern how coupler scarcity bounds the number of BVs the corridor can profitably field: the congestion-aware routing weight (24) in the first step and the free-coupler

charging availability (6) in the second. Together they make the profit-max count finite — four vessels at the base stack (Table 10), peaking where free coupler-hours run out at the two-coupler hubs, not at an imposed cap. This appendix records how each is computed.

Free couplers. At each hub r the simultaneous major-cruise host count $o_{r,h}$ is read directly from the timetable — the same hourly berth-occupancy series that underlies the contention accounting (§2.4.2). It counts every berthed major-cruise ship, the nine non-dispatched majors included, and excludes only the small yacht and expedition vessels, which are displaced first. The free couplers are $f_{r,h} = [\, K_r - o_{r,h} \,]^+$. A berthed host leaves only its charge headroom above the hotel load (§2.4.2), but an *idle* coupler carries the full 20 MW for a BV. The binding scarcity is at the two-coupler Ligurian hubs, Genoa and Savona ($K_r = 2$), where BV routing concentrates; the other two-coupler hubs (Valencia, Palermo, Messina, Cagliari) see little routing, and the busier western hubs have more couplers.

Routing weight. This per-leg weight is the edge weight the first-step chain DP maximises along each vessel's chain (the maximum-weight path for a single BV, the min-cost flow for several; §2.4.1) — a routing signal only, never entering the second-step LP. Writing $[t_l^0,\, h_l^{\mathrm{dep}})$ for the ship's dwell at the origin of leg l (its previous same-ship arrival to this departure, clamped to a berth-length maximum), the congestion-aware routing weight is

$$w_l^{\mathrm{cong}} = \min(D_l, \bar{E}) \;\cdot\; \frac{\left|\{h \in [t_l^0,\, h_l^{\mathrm{dep}}) : f_{r(l),h} \geq 1\}\right|}{h_l^{\mathrm{dep}} - t_l^0}, \tag{24}$$

the deliverable host energy $\min(D_l, \bar{E})$ scaled by the fraction of the origin berth window on which a coupler is free. A call whose couplers are never full keeps its full weight regardless of dwell length (the BV carries charge between docks); a partly-full berth is discounted in proportion, a fully congested one to zero.

Using a single $\bar{E}$ fleet-wide is a conservative simplification: it forgoes letting small-pack vessels specialise on short, dense runs and large-pack ones on long hauls, a refinement that could only raise the joint value the LP achieves.