Louvain School of Management

# When does onshore power supply reduce $CO_2$? A threshold model for the European cargo fleet under IMO 2050 targets

Author: Charlie Degembe
Supervisors : Per Joakim Agrell, Manuel Herrera Rodriguez
Academic year : 2025 – 2026
Dissertation for the master 120 in business engineering with a specialization in supply chain.
Daytime schedule

## Declaration on the responsible use of generative AI

This declaration is made in accordance with the LSM guidelines on the responsible use of artificial intelligence in a master's thesis.

| **Statements on responsible use of generative AI** | **Your answer** |
|---|---|
| The content of my master's thesis is my original output, even if I used generative AI to help prepare it. | ☒ **Yes** ☐ No |
| I master the theories, tools and methods that I use for the thesis. | ☒ **Yes** ☐ No |
| I verified that the output presented in the thesis is accurate and valid, and that it contains no hallucination, no false reference, etc. | ☒ **Yes** ☐ No |
| I acknowledge that I master the final output and that I am able to explain it and answer questions about it. | ☒ **Yes** ☐ No |
| I made sure not to provide generative AI tools with access to confidential data or information. | ☒ **Yes** ☐ No |

### Signature

| | **Author last name, first name(s)** | **Date** | **Signature** |
|---|---|---|---|
| 1 | Degembe, Charlie | 10 August 2026 | |

*A fuller account of the tool used, the stages at which it was used, the nature of the prompts given and the verification applied to the output is provided in Appendix B.*

**Acknowledgements**

I would like to express my sincere gratitude to my supervisor, Professor Per Agrell, for his guidance and support throughout the development of this thesis, and for feedback that helped sharpen the research question and improve the overall rigor of this work.

I am also grateful to Professor Manuel Herrera, whose insight first pointed me toward the topic of cold ironing, a suggestion that opened up a subject both academically rich and practically relevant, and whose input helped shape the direction this thesis ultimately took.

Finally, I thank both professors for the time they dedicated to this work and for the balance of freedom and guidance they offered me throughout.

## Table of contents

1. Introduction ........ 1
2. Literature Review ........ 3
  2.1 Shore power and the regulatory mandate ........ 3
  2.2 Why deployment remains slow: cost, incentives and governance ........ 4
  2.3 The emission balance and the electricity that replaces the engines ........ 6
  2.4 Time at berth, connection overhead and congestion ........ 7
  2.5 Research gap ........ 9
  2.6 Positioning of this study ........ 10
3. Methodology ........ 11
  3.1 Research design ........ 11
  3.2 Data source and sample construction ........ 12
  3.3 The net balance and its notation ........ 14
  3.4 Deriving the three thresholds ........ 17
  3.5 The queueing extension ........ 18
  3.6 Limitations ........ 19
  3.7 Reproducibility ........ 21
4. Results and analysis ........ 21
  4.1 The emission pool and its concentration ........ 21
  4.2 The three thresholds ........ 22
    4.2.1 Threshold 1: break-even grid carbon intensity ........ 22
    4.2.2 Threshold 2: minimum viable berthing time ........ 23
    4.2.3 Threshold 3: the critical overhead and the two-hour boundary ........ 25
  4.3 The queueing effect ........ 26
  4.4 Sensitivity to the open axes ........ 29
5. Conclusion ........ 31
  5.1 Contributions to the literature ........ 31
  5.2 Policy contributions ........ 32
  5.3 Critique ........ 33

5.4 Further work ........ 34
References ........ 36
Appendix A. Figures ........ 40
Appendix B. Statement on the use of generative AI ........ 46
Tool and period of use ........ 46
Stages at which it was used ........ 46
Nature of the prompts ........ 46
Verification of the output, and errors identified ........ 47
Statement ........ 47
Appendix C. The queueing computation, step by step ........ 48

# 1. Introduction

International shipping accounts for approximately 2.9% of global $CO_2$ emissions from human activity (IMO, 2020). A large share of the air pollution that goes with it is concentrated in ports, where ships at berth keep their auxiliary engines running to power onboard systems (Le, 2024). Onshore power supply, also called cold ironing, lets a berthed ship switch those engines off and take electricity from the quay instead (Le, 2024). Across the European Economic Area and the United Kingdom, Stolz et al. (2021) estimate that supplying auxiliary power from national grids would avoid 3 Mt of $CO_2$ a year, out of 8.11 Mt emitted at berth, and that the figure would reach 5 Mt only if the electricity were carbon neutral. The gap between those two numbers is the subject of this thesis. The revised IMO strategy aims for net-zero greenhouse gas emissions from international shipping by or around 2050, with checkpoints for 2030 and 2040 (Comer & Carvalho, 2023). These are ambitions, not binding rules. The instrument that would make them binding, the IMO Net-Zero Framework, was approved in April 2025 but its adoption was adjourned for twelve months by an extraordinary session of the Marine Environment Protection Committee held from 14 to 17 October 2025 (IMO, 2025). What obliges a ship to connect today is therefore European law, not global law: FuelEU Maritime requires containers and passenger ships to connect at major EU ports from 2030 (European Commission, 2023). That obligation is narrower than the IMO ambition, which covers international shipping of every type, and this study is narrower still, as Section 3.2 explains.

A ship plugged into the grid has not stopped emitting; its emissions have moved ashore. How much that move is worth depends first on what the shore electricity is made from: on a grid fed mainly by fossil fuels, the net gain can be small or nil (Hu, 2022). It also depends on time. Connecting and disconnecting take time, the auxiliary engines keep running while it happens, and the extra time at berth cuts into the berth's availability for the next vessel (Abu Bakar et al., 2022). So the answer changes from one call to the next, with the grid, with the length of the stay, and with how busy the berth is. The literature and the regulation frame the question differently, by vessel segment: which ship types to equip first (Abu Bakar et al., 2022; Pruyn & Willeijns, 2022). That framing hides the operational conditions that actually decide the outcome.

This study asks a deliberately conditional question: under what circumstances does onshore power supply reduce net $CO_2$ emissions for the European cargo fleet, rather than merely displace or, under some conditions, increase them? It answers by building a deterministic model anchored on measured at-berth emissions from the EU Monitoring, Reporting and Verification (EU-MRV) dataset, in which every parameter is either measured, drawn from a citable source, or swept across a range where no single public value exists. Some of those ranges are themselves sourced: the connection overhead is swept, but between bounds a published study provides. The model derives three thresholds that a port call must clear for shore power to reduce carbon in net terms: a break-even grid carbon intensity, a minimum viable berthing time, and a critical connection overhead beyond which the regulatory mandate itself becomes counterproductive. A queueing extension then relaxes the assumption that calls can be treated in isolation and quantifies the congestion cost that the connection overhead imposes on the vessels waiting behind.

Two findings follow. First, conditional on the same call duration, fuel, engine, grid and connection overhead, the retained share does not vary with vessel type: the segment commonly used to prioritize adoption acts, in that conditional sense, as a proxy for call duration. This narrows rather than overturns the segment-specific conclusions of earlier work such as Pruyn and Willeijns (2022), whose results reflect the parameters typical of a segment rather than an independent segment effect. Second, the connection overhead carries a second-order cost that static assessments omit: by reducing berth throughput it lengthens the queue of vessels waiting at anchorage, so that under congestion shore power may raise net emissions rather than lower them. Building on the fleet-level net-balance approach of Stolz et al. (2021), the contribution of this work is to integrate the grid, duration and congestion effects into a single threshold framework that is computable from public data and generalizable across the fleet rather than tuned to a single port.

The remainder of this thesis is organised as follows. The literature review surveys the technical, economic, environmental and operational dimensions of cold ironing, and positions the present contribution against the closest existing work. The methodology sets out the data, the model and its parameter classification. The results present the three thresholds and their sensitivity, and the discussion draws out the implications for port investment, for the design of the two-hour exemption boundary, and for the literature. All figures are collected in Appendix A and are referred to in the text by number.

# 2. Literature Review

The literature on shore power is large, but it is not evenly spread. It says a great deal about why the technology is hard to finance and to govern, and much less about the physical conditions under which connecting a ship cuts emissions. This review is arranged to move from the first question to the second. Section 2.1 sets out what shore power is and what European law now requires. Section 2.2 covers the question the literature has concentrated on, why deployment is slow, and says why this thesis sets it aside. Sections 2.3 and 2.4 then take the three quantities the emission balance turns on: how clean the shore electricity is, how long the ship stays and how long connecting takes, and what congestion does to both. Each section ends by saying what its strand settles and what it leaves open. Section 2.5 gathers those open points into the gap this thesis addresses, and Section 2.6 places the study against the three works closest to it.

## 2.1 Shore power and the regulatory mandate

Cold ironing, also known as shore-side electricity or onshore power supply (OPS), is a technology that allows ships at berth to shut down their auxiliary engines and connect to the port's electrical grid, thereby eliminating emissions during port stays (Le, 2024). From a technical standpoint, this works by replacing onboard diesel generators that normally power lighting, refrigeration, air conditioning and cargo equipment with shore-supplied electricity through specialized cables (Abu Bakar et al., 2023). As ports increasingly adopt zero-emission models, cold ironing has emerged as one of the most promising technical solutions. In Europe, all ports were required to have cold ironing provision by 2025 (Zis, 2019). A survey of port experts and stakeholders found that only 40% of Croatian state-owned ports have a connection capacity above 1 kV, the practical minimum for implementation (Glavinović et al., 2023), highlighting the gap between regulatory ambition and operational reality.

From a technical perspective, cold ironing requires synchronization of frequency and voltage between the ship and the shore, as most ships operate at 60 Hz while European ports operate at 50 Hz, necessitating costly frequency converters (Abu Bakar et al., 2023). Furthermore, the coordination between multiple stakeholders, port authorities, shipping lines, policymakers and terminal operators adds another layer of complexity to its implementation.

What European law now requires is narrower than the technology allows, and the boundaries matter for what follows. Under the FuelEU Maritime Regulation, containerships and passenger ships of 5 000 gross tonnage and above must, from 1 January 2030, connect to shore power and cover their whole electrical demand from it while moored in the main EU ports of the trans-European transport network, with other equipped ports following in 2035 (European Commission, 2023; Osipova & Carraro, 2023). The obligation carries no condition on how long the call lasts. A duration appears only in the derogations, which release stays shorter than two hours, the recital explaining that this is the minimum time required for connection. Vessels using an alternative zero-emission technology are exempt on the same terms.

That two-hour figure is treated in the literature as a given rather than as a question. Osipova and Carraro (2023) use it as a filter, counting the roughly 15,700 ships that spent more than two hours at berth in 489 EU ports in 2019 and estimating from them a demand of nearly 5.9 TWh; they then assess what the regulations achieve, concluding that the current level of ambition would cut only 24% of the 4.37 Mt of at-berth $CO_2$ emitted annually in the Union. Their critique is one of coverage: too few ship types, too few ports, boilers left out. What no study asks is whether two hours is the right length, that is whether a call of that duration is long enough for connecting to reduce emissions at all. The threshold is used to select vessels and never tested as a condition, and Section 2.5 returns to that omission.

This strand settles three things: what shore power is, how a ship connects to it, and that European ports must now provide it. It leaves two questions open, and both decide the answer to the question asked here. The first is how much $CO_2$ connecting avoids, which depends on what the shore electricity is made from; Section 2.3 takes that up. The second is that this strand treats connecting as instantaneous. A ship is described as either burning fuel or drawing from the quay, with nothing in between, and no time is charged to the act of connecting; Section 2.4 returns to that. Before either, Section 2.2 covers the question the literature has spent most effort on, which is why a technology this well understood is installed in fewer than half of European ports.

## 2.2 Why deployment remains slow: cost, incentives and governance

Port sustainability is usually described through the Triple Bottom Line, which sets commercial growth against environmental and social costs. The framework is Elkington's, and it is used here as Alamoush et al. (2021) present it, since that is the reading this literature has adopted; they sort port sustainability into sixteen linked categories tied to the United Nations Sustainable Development Goals. Notteboom et al. (2020) then map what ports do, in five groups of green actions observed across the Rhine-Scheldt Delta. Shore power falls into two of those groups at once, because it needs investment on the quay from the port authority and a retrofit on board from the operator. That double requirement is where the coordination problem starts, since neither side will commit first. Their survey of Belgian and Dutch port actors finds greening on the agenda since 2010 while implementation lags, which they put down to discontinuous public policy.

The capital requirement is substantial on both sides. Shoreside installation runs from $300,000 to $4 million per berth and vessel retrofitting from $300,000 to $1 or $2 million, with operation and maintenance adding up to 12% of capital cost (Abu Bakar et al., 2023); for a medium-sized United Kingdom port the total reaches £6.6 million with a payback of seven years, or 3.5 years with EU subsidy (Innes & Monios, 2018). Segment-level appraisals are harsher. Pruyn and Willeijns (2022) find shore power economically unviable for tankers in every scenario tested, driven by the fixed cost of supply, and Le (2024) identifies lack of capital and absence of standardization as the two structural barriers to adoption. A socio-economic reading reverses the conclusion: Canepa et al. (2023) shows that once the external health costs of the pollutants covered by the CAFE method, among them NH3, NOx, SO2, volatile organic compounds and PM2.5, are internalized, shore power at the Port of Genoa becomes justifiable for every stakeholder. Which conclusion holds depends on where the accounting boundary is drawn, and funding remains one of the barriers to those studies list, alongside the coordination of port authorities, shipowners and energy companies (Peddi et al., 2024).

The barriers are not only financial. Williamsson et al. (2022) group them into technology and operations, institutions, economics and stakeholders, and identify a structural asymmetry: the investment costs land on ports and shipowners while the benefits spread across communities and public health systems, which discourages anyone from moving first. Bignucolo et al. (2025) arrive at the same point from the other side. In their model of the Italian port system, a cold ironing plant only pays under a narrow set of conditions holding at once: a wholesale electricity price at the bottom of its range, one particular management model, and a marine fuel price more than 25% above its base level. Across this strand the conclusion keeps returning in the same form: left to itself, shore power does not spread.

This is the largest body of work reviewed here, and this thesis sets it aside on purpose. What it establishes is clear: shore power is hard to finance, the costs fall on ports and operators while the benefits spread across communities and health systems, and that mismatch discourages unilateral investment. Those findings explain why deployment is slow, and they are why the European legislator stopped encouraging and started requiring. But every one of them assumes the point this thesis questions. A cost-benefit study asks whether an emission cut is worth paying for; it takes the cut itself for granted. Here the deployment decision is treated as settled by law, and the earlier question is asked instead: under what conditions does connecting a vessel cut emissions at all? No cost enters the model that follows, and none is needed. A call that increases emissions cannot be saved by a good tariff.

## 2.3 The emission balance and the electricity that replaces the engines

Beyond emission reduction, cold ironing can be improved by drawing the shore electricity from renewable sources rather than from the national grid (Kelmalis et al., 2024). Their case study at Mytilene, on Lesvos, shows both the scale and the reason. The island grid runs on a diesel power plant, so the electricity a connected ship would draw is itself carbon-intensive; a simulated 20 MW wind installation would produce about 72,080 MWh a year against a port demand of 6,118 MWh a month, enough to cover it with a surplus, but demand fluctuations would still leave the port drawing some 924 MWh a year from the main grid, which is why the authors call for storage and demand response alongside it. Yuksel et al. (2025) quantify what the choice of supply is worth. Grid electricity as the cold ironing source cuts $CO_2$ equivalent emissions by 58.92%, a photovoltaic and battery system by 80.73 to 85.82% depending on irradiation, and their most ambitious configuration, a photovoltaic plant feeding proton exchange membrane fuel cells on hydrogen cracked from green ammonia, by 99.92 to 100%, but at a payback of 23.69 years against 2.06 for grid electricity. The general point the strand makes is that the same connection can help a great deal or hardly at all depending on the country it is made in.

One study takes this dependence past the qualitative statement. Hu (2022) builds a carbon emission correlation model between marine fuel and shore power for vessels at berth and applies it across the Chinese coastal provinces, whose generation mixes differ sharply. Shore power is, in that paper's terms, by no means certain to lower carbon emissions: only a minority of coastal provinces are suitable for it, the condition depends on the berthing time of the vessel, and it is governed not by ship size but by the share of each generation technology in the provincial mix. Hu also charges a duration to the connection itself, fixed at one hour, and derives from it a basic berthing time below which connection does not pay (Hu, 2022, p. 6, equations 13 to 15). That is the closest published precedent to the first two thresholds of this study, and Section 2.6 sets out what separates the two.

The dependence is therefore established, and established qualitatively. Every reduction percentage reported here is conditional on a supply mix rather than being a property of shore power: the 58.92% on grid electricity and 80.73 to 85.82% on a photovoltaic and battery system (Yuksel et al., 2025) would both be different figures on a different grid. What no study in this strand provides is the value at which the benefit disappears. Hu (2022) states the mechanism plainly, that shore power is by no means certain to lower emissions where the generation mix is carbon-intensive, and derives a threshold for it; the rest of this strand asserts the dependence without quantifying it, so a port cannot locate itself within it. That threshold is the first of the three derived in Section 3.4.

## 2.4 Time at berth, connection overhead and congestion

The operational integration of cold ironing raises scheduling challenges that the literature often understates. Abu Bakar et al. (2022) make berth duration a variable to be forecast, since longer stays raise shoreside consumption while reducing berth availability for incoming vessels, and they identify ship type, size, arrival hour, mode of operation and cargo capacity as the drivers, interacting non-linearly. They single out cargo vessels and tankers as the segments with the greatest potential, precisely because their regular year-round calls make berthing patterns predictable. The implication is that connection cannot be assessed in isolation from the berth allocation system it sits inside.

Two recent studies close part of that distance from the congestion side, though neither models the channel examined here. Rhodes et al. (2025) study the queuing system the ports of Los Angeles and Long Beach introduced in 2021, under which a vessel's place in the queue is set when it leaves its previous port instead of when it arrives. Across ten thousand voyages by 1,157 container ships over six and a half years, $CO_2$ per voyage fell by 16 to 24% after the change. The authors put this down to ships slowing across the Pacific instead of racing to wait at anchor, and they are careful with the causal claim, since control ports without such a system also improved over the period. Even so, the result is the closest empirical demonstration that the way a port runs its queue moves as much $CO_2$ as the technologies usually discussed as decarbonisation measures. What it does not do, and what this study does, is trace any part of that queue back to a connection procedure.

Akakura (2023) measures the other side of the mechanism. From AIS data he identifies ships waiting offshore at sixteen large container terminals, works out how long they wait, and estimates the extra $CO_2$ that waiting causes, comparing each ship with one that sails straight in. His figure is a total over a month at a terminal rather than a rate per ship-hour, so it is not the parameter this model needs, but it fixes the order of magnitude: waiting emissions are not a rounding error beside berthing emissions. What he does not do, since shore power is not his subject, is trace any of that waiting back to a connection procedure.

Time at berth is therefore known to matter, and its two components are known separately. Abu Bakar et al. (2022) make berthing duration a variable to be forecast and identify the connection and disconnection procedures as a risk to turnaround; Zis (2019) reports the duration of those procedures directly. What neither does, and what no study in this review does, is convert that time into emissions. It is counted against berth availability and vessel turnaround, never as fuel burned during a period in which the vessel draws on neither source. The same holds for the congestion that follows: it is described as a risk in the scheduling literature and quantified only inside port-specific optimizations that require proprietary data. Table 1 sets out how each of the five strands reviewed here treats the three quantities that govern the balance.

*Table 1. How each strand of the reviewed literature treats the three quantities that govern the net emission balance. No strand carries all three. Hu (2022) carries the grid and the call duration, including the time the connection consumes, but no congestion channel; Akakura (2023) and Rhodes et al. (2025) measure congestion without relating it to shore power.*

| Strand | Studies | Grid carbon intensity | Time at berth and connection overhead | Congestion |
|---|---|---|---|---|
| **Port sustainability and governance** | Alamoush et al. (2021); Notteboom et al. (2020); Williamsson et al. (2022); Le (2024) | Not addressed | Not addressed | Not addressed |
| **Technology and energy supply** | Abu Bakar et al. (2023); Kelmalis et al. (2024); Yuksel et al. (2025); Glavinović et al. (2023) | Central, but reported as case results for a given supply mix rather than as a function | Connection treated as instantaneous | Not addressed |
| **Investment and cost appraisal** | Innes and Monios (2018); Pruyn and Willeijns (2022); Canepa et al. (2023); Peddi et al. (2024); Bignucolo et al. (2025); Zis (2019) | Enters through the electricity tariff, not through carbon content | Connection duration reported (Zis, 2019, p. 85); berth time enters as a cost driver | Noted as a risk to turnaround, not quantified |
| **Berth allocation and scheduling** | Abu Bakar et al. (2022); Innes and Monios (2018) | Not addressed | Central: berth duration is the forecast variable | Modelled, but inside port-specific optimizations requiring proprietary data |
| **Fleet-level net balance** | Stolz et al. (2021); Martínez-López et al. (2021); Hu (2022) | Fixed by Stolz et al. (2021); made the variable of interest by Hu (2022) | Connection and disconnection carried by Martinez-Lopez et al. (2021); a connection duration of one hour and a berthing-time condition derived by Hu (2022) | Not addressed |

| | | | | |
|---|---|---|---|---|
| **Congestion and the emissions of waiting vessels** | Akakura (2023); Rhodes et al. (2025) | Not addressed | Waiting time measured from AIS data, but not attributed to any connection procedure | Measured empirically, and waiting emissions expressed as a share of at-berth emissions |

Read across the rows, the table shows that no strand carries all three quantities through to a condition a port could apply. Read down the columns, it shows why. Grid carbon intensity is treated seriously only where the generation mix is the object of study, as in Hu (2022), and is elsewhere fixed to the case at hand. Connection overhead is measured only where the concern is the schedule and is then expressed in minutes of turnaround rather than in tons of $CO_2$. Congestion is measured empirically by Akakura (2023) and Rhodes et al. (2025), but neither relates it to a connection procedure, while the scheduling literature that does model that relation requires proprietary terminal data. The three quantities are separated not by oversight but by the disciplinary boundaries of the strands that study them.

## 2.5 Research gap

The gap this thesis addresses is the last column of Table 1 read as a whole. Literature provides a comprehensive understanding of shore power across technical feasibility, economic viability, environmental impact and regulatory frameworks, but it treats costs, emissions and operations as separate variables rather than as elements of a single balance. Environmental assessments take the connection itself as costless in emission terms. The scheduling and berth-allocation literature does model the operational cost of connection but does so inside port-specific optimizations that require proprietary terminal data and cannot be generalized. Neither treats the carbon intensity of the substituting grid as a variable whose value determines the sign of the result.

Two consequences follow. The first is that no study identified in this review brings the grid, the duration and the congestion conditions together in one framework computable from public data. Hu (2022) comes closest: he joins the grid to a berthing-time condition that already charges a duration to the connection, but he treats each call on its own and carries no congestion channel, while the congestion studies measure the waiting that connection would worsen without linking the two. The question of when shore power cuts $CO_2$ has therefore been answered for one terminal, or for one national generation mix, or under one fixed set of assumptions, but never as a set of conditions that any port call either meets or fails. The second is more specific: the mandate itself has never been tested against the conditions it presupposes. FuelEU Maritime requires containerships and passenger ships above 5 000 gross tonnage to connect at major EU ports and exempts stays shorter than two hours, yet no study establishes whether that two-hour boundary falls in the right place, in either direction, for connection to pay in net terms. This study addresses both gaps: it derives, for the European cargo fleet, the thresholds that decide whether shore power cuts net emissions, and it evaluates the regulatory boundary against them.

## 2.6 Positioning of this study

The present study sits at the intersection of three strands of shore-power literature. The first is the fleet-level net-balance assessment, of which Stolz et al. (2021) is the reference case: using EU-MRV and AIS data, they quantify the $CO_2$ reduction potential of shore-side electricity across European ports by netting the grid emissions against the auxiliary-engine emissions displaced. The net-balance logic used here follows theirs directly, and the study is indebted to it for the treatment of conversion efficiencies. Where it departs from them is in what happens to the balance once the time cost of connecting is considered.

Hu (2022) belongs to the same strand and reaches the threshold logic independently, on a different fleet and a different regulation. His equation (15) gives a basic berthing time that is the ratio of the connection duration to one minus the share of the on-board emission rate claimed by the grid, the same structure as Threshold 2 here, and his equation (14) gives a grid cleanliness criterion equivalent to Threshold 1. The first two thresholds of this study are therefore a European counterpart to a published result rather than a new mechanism, and their agreement is evidence for both. Three things separate the two. Hu models the displaced emissions from a tonnage regression, whereas this study takes them from verified EU-MRV declarations. He fixes the connection duration at one hour, whereas it is carried here as an axis, which is what makes the third threshold possible. And he treats each call in isolation, with no congestion channel, which is where the contribution of this study lies.

One study comes closer to the mechanism modelled here than the rest. Martínez-López et al. (2021), assessing cold ironing for short sea shipping in Spanish ports, carry a connection and disconnection lag time as an explicit term and observe that its influence on the outcome is inversely proportional to berthing time. That observation is the qualitative form of the second threshold derived below. Three things distinguish the present treatment. The lag time is not fixed at an assumed value but carried across the published range; the relationship is stated in closed form as a minimum viable call duration rather than reported as a sensitivity; and it is applied to the whole EU cargo fleet rather than to a set of short sea routes. The congestion channel, which no study in this strand quantifies, is added on top.

Where this study departs from Stolz et al. (2021) is in the treatment of connection overhead. They acknowledge, as one possible direction of effect, that shore-power connection may delay port operations and prolong the stay at berth, with vessels recovering the lost time by sailing faster afterwards (Stolz et al., 2021, Figure 7, arrow II) but note it qualitatively without quantifying it. The mechanism modelled here is distinct: rather than a speed effect on the connecting vessel, the connection overhead is treated as a reduction in berth throughput that lengthens the anchorage queue for the vessels behind it, and it is quantified through an M/M/c queueing model.

The second strand is the berth-allocation and scheduling literature, which does capture the trade-off between shore power and waiting time. These studies decide which ships go to which berths, and which of them connect, for one terminal at a time, and they find consistently that giving priority to shore power lengthens the queue at anchorage. They therefore establish that the trade-off is real. What they cannot do is generalize it: each result is tuned to one terminal and rests on berth-level data that ports do not publish.

The contribution is therefore neither the net balance, which follows Stolz et al. (2021), nor the existence of a shore-power and congestion trade-off, which the scheduling literature has established. It is the integration of the two into one framework: a break-even grid intensity, a minimum viable berthing time and a congestion limit, computable from public data, generalizable across the EU cargo fleet rather than tuned to one port, and treating connection overhead as a driver of anchorage congestion. To our knowledge this congestion channel has not previously been quantified in the shore-power literature.

# 3. Methodology

## 3.1 Research design

This study is a deterministic feasibility model rather than a statistical inference. It quantifies the net $CO_2$ outcome of replacing at-berth auxiliary power with shore power across the EU cargo fleet and derives the conditions under which that outcome is positive.

The design follows directly from the data environment. At-berth emissions are publicly measured and verified at vessel level. The operational parameters that determine whether shore power pays off are call duration, connection overhead, berth count, berth utilisation, and the emission rate of a vessel waiting at anchorage which are not available in any free public source at the level of detail required.

The model therefore treats every parameter as belonging to exactly one of three categories, and the three results it produces do not share the same status. The break-even grid intensity is an analytical threshold, a property of the technologies compared and independent of any operational assumption. The minimum viable duration and the critical overhead are conditional operational thresholds: they hold for a given fuel, engine and grid, and move when those move. The queueing results are scenario outcomes, computed on parameters that no port publishes. Where a value can be neither measured nor sourced it is swept across its plausible range rather than fixed, but sweeping a range is not the same as having no assumptions, and the reader should treat the three sets of results accordingly. The categories are:

*Table 2. The three-category treatment of model parameters. A parameter is measured when it is reported vessel by vessel in EU-MRV, sourced when a citable authority gives it, and carried as an axis when neither holds. No fourth category is used, so that no value in the model rests on an assumption that is not visible as a range.*

| **Category** | **Treatment** | **Parameters** |
|---|---|---|
| Measured | Taken directly from data | At-berth $CO_2$ by segment |
| Sourced | Taken from a citable authority | Fuel emission factor, grid carbon intensity, regulatory floor |
| Axis | Never assigned a value; results presented across its plausible range | *SFOC*, call duration, connection overhead, berth count, berth utilisation, anchorage emission rate |

## 3.2 Data source and sample construction

The primary source is the EU Monitoring, Reporting and Verification dataset (EU-MRV), which EMSA publishes through the THETIS-MRV portal. The file for the 2024 reporting year holds verified annual emissions and fuel consumption for every vessel above 5,000 GT that called at an EU port.

The variable used here is the $CO_2$ a vessel emitted at berth in ports under Member State jurisdiction, reported per vessel in tonnes per year. It is measured, it is verified, and it is exactly the emission stream shore power would remove. Building the model on it avoids reconstructing at-berth emissions from assumed hotel loads, which much of the literature does and which is a recurring source of error.

The 2024 file lists 15,176 vessels. Of these, 1,027 appear only in the separate partial-reporting sheet and are dropped throughout, because part-year coverage distorts annual figures and nothing can correct for it; 14,149 vessels with full reports remain. Another 2,959 of those report exactly zero at-berth $CO_2$. They are not ships that berthed without emitting; they are ships that did not berth in an EU port that year, or did not report it, so they add nothing to the pool under study. That leaves 11,190, from which 589 non-cargo vessels are removed, giving the final sample of 10,601.

Passenger, cruise, offshore and unclassified segments are excluded. This is a scope decision rather than a data limitation: those segments do not constitute the traffic of high-throughput cargo ports, which is the object of this study. No claim is made here about whether shore power serves them well or badly, since none of the quantities this model turns on has been examined for them.

Ro-pax vessels are retained although they carry passengers, and the criterion is worth stating because it is the traffic rather than the manifest. EU-MRV classifies ro-pax ships separately from passenger ships and from cruise passenger ships, and the sample keeps every EU-MRV category that carries freight. Ro-pax carry rolling freight on scheduled services: roll-on roll-off units accounted for 275 million tons, or 16.7% of EU short sea shipping of goods in 2024, and 28.9% of it in the Baltic, and the units in question are road goods vehicles, unaccompanied trailers and rail wagons (Eurostat, 2026). Excluding ro-pax would remove a sixth of European short-sea freight from a study of European freight shipping. The cost of retaining them is that their hotel load is higher than a pure cargo vessel's, which shows in the anchorage emission rate discussed in Section 3.5; they are the segment for which the label cargo fleet sits least comfortably. Cruise and pure passenger vessels are excluded on the same criterion, since they carry no freight, but that exclusion is a scope decision rather than a finding, and Section 5.4 returns to it.

EU-MRV covers EU-scope emissions only: voyages to, from and between EU ports, plus what is emitted at berth in EU ports. It reports the emissions, not how long the calls that produced them lasted. A vessel trading worldwide reports only its EU-related activity, so EU exposure differs enormously between ships of the same segment: among container ships, the ninetieth percentile of at-berth $CO_2$ is roughly twenty-three times the tenth.

An average per vessel would therefore mix a ship calling twenty times a year with one calling once, and would describe no ship that exists. It would also carry a bias whose direction cannot even be signed without knowing how EU exposure is distributed, which the dataset does not say.

## 3.3 The net balance and its notation

Before the model is stated it is useful to fix the notation. Table 3 lists every symbol used in equations (1) to (12), with its unit, its status in the three-category architecture of Section 3.1, and the source or the reason it is carried as an axis.

*Table 3. Notation, parameter status and sources. Every symbol used in equations (1) to (12) appears once. "Status" follows the three-category architecture of Section 3.1 measured from EU-MRV, sourced from a citable authority, or carried as an open axis; "derived" quantities are computed from the others and introduce no new assumption. The final column gives, for each input, the source and the quantified effect on the result.*

| Symbol | Meaning | Unit / value | Status | Source, and effect on the result |
|---|---|---|---|---|
| **INPUTS** | | | | |
| *E* | At-berth $CO_2$ of a segment (the gross saving) | 6.15 Mt/yr in total; t $CO_2$/yr | Measured | EU-MRV 2024, own analysis. Sets the gross saving: the absolute size of every result scales with it. |
| *SFOC* | Specific fuel oil consumption of the auxiliary engines | 156–225 g/kWh by fuel and era | Sourced axis | IMO (2020), Table 19, p. 71, auxiliary engines. Set jointly with the emission factor by the fuel the engine burns; the eight documented combinations are listed in Table 4. |
| *EF* | Fuel emission factor: $CO_2$ released per tonne of fuel burned | 2.750–3.206 t $CO_2$/t by fuel | Sourced axis | IMO (2020), Table 21, p. 74; FuelEU Maritime, Annex II (European Commission, 2023). HFO 3.114, MDO 3.206, LNG 2.750. With *SFOC* it drives the break-even range: 429 g$CO_2$/kWh to 701. |
| *Igrid* | Carbon intensity of the shore electricity grid | 30 to 566 g$CO_2$/kWh by country | Sourced | Ember (2025): EU average 213 g$CO_2$/kWh in 2024. Country values from the European Environment Agency (2025) indicator on the emission intensity of electricity generation. Sets the grid share: 33% at the EU average, 89% for Poland. |
| $t_o$ | Connection overhead: connection plus disconnection, engines running | 10 to 90 min | Axis | Zis (2019), p. 85: under 10 min at an optimal setting (Gothenburg terminal operators), 1 to 1.5 h described as realistic. Martínez-López et al. (2021) assume 1 h. Sets the minimum viable call duration: 0.25 h at 10 min, 1.50 h at 60 min, 2.25 h at 90 min. Above 80 min the viability threshold exceeds the 2 h exemption boundary. |
| *d* | Duration of the call at berth | 2–36 h | Axis | Not published per segment. Sets how much overhead is absorbed. Below 2.25 h (at 90 min overhead) the balance is negative. |
| *dfloor* | Regulatory duration floor below which no connection is mandated | 2 h | Sourced | FuelEU Maritime / AFIR; Osipova and Carraro (2023). Reference line for Threshold 3, not an input to the balance itself. |
| *c* | Number of berths operating in parallel | 1–15 | Axis | Not published per port. At a 12 h call and $\rho$ = 75%, the retained share rises from -100% at c = 1 to +26% at c = 5 and +42% at c = 10. |
| $\rho$ | Berth utilisation | 0–100% | Axis | Not published per port. At a 12 h call, the retained share falls from 52% at $\rho$ = 50% to 26% at $\rho$ = 75%, turns negative just below 80%, and saturates at 88.9%. |
| $\alpha$ | Emission rate of a waiting vessel, as a share of its at-berth rate | 10–150% of at-berth | Sourced axis | IMO (2020), Table 17, pp. 68-69: power at anchor over power at berth, boiler plus auxiliary, converted at the specific consumptions of Table 19, p. 70. Some cargo types exceed 100%. Not published per port, so swept. |
| n.a. | Conversion efficiencies (alternator, shore-to-ship losses) | omitted | n.a. | Omitted here, which is a simplification of this study and not of Stolz et al. (2021): they include a generator |

| Symbol | Meaning | Unit / value | Status | Source, and effect on the result |
|---|---|---|---|---|
| | | | | efficiency of 92% and grid transmission losses taken from IEA data. The two act in opposite directions, the generator raising the ship-side figure and transmission losses lowering the delivered grid figure. |
| **DERIVED QUANTITIES AND RESULTS** | | | | |
| *eship* | $CO_2$ the auxiliary engine emits per MWh generated on board | t $CO_2$/MWh | Derived | Equation (1) |
| *egrid* | $CO_2$ the grid emits per MWh delivered to the vessel | t $CO_2$/MWh | Derived | Equation (2) |
| *s* | Share of the gross saving consumed by grid emissions | dimensionless | Derived | Equation (3) |
| *NET* | Net annual $CO_2$ saved by a segment | t $CO_2$/yr | Result | Equation (4) |
| *R* | Share of the gross saving retained after grid and overhead | dimensionless | Result | Equation (5) |
| *Igrid** | Break-even grid carbon intensity (Threshold 1) | g$CO_2$/kWh | Result | Equation (6) |
| *dmin* | Minimum viable call duration (Threshold 2) | h | Result | Equation (7) |
| *to,crit* | Critical connection overhead (Threshold 3) | h | Result | Equation (8) |
| $\lambda$ | Vessel arrival rate at the berth group | calls/h | Derived | Equation (9) |
| *S* | Berth service time per call | h | Derived | Equation (9) |
| *a* | Offered load, $\lambda$S | erlangs | Derived | Equation (9) |
| *C(c, a)* | Erlang C probability that an arriving vessel has to wait | dimensionless | Derived | Equation (10) |
| *Wq* | Mean waiting time at anchorage | h | Derived | Equation (10) |
| *ΔW* | Additional waiting per vessel caused by the connection overhead | h | Derived | Equation (11) |
| *Rqueue* | Share of the gross saving retained once congestion is counted | dimensionless | Result | Equation (12) |

For each segment, the net annual saving is the measured at-berth $CO_2$ (the gross saving), less the $CO_2$ the grid emits to supply the same energy, less the $CO_2$ emitted while the auxiliary engines run during connection and disconnection. Writing E for the measured at-berth $CO_2$ of a segment, the auxiliary engine emits per unit of energy: The rate *e_ship* is that of the auxiliary engine alone: the boiler share of at-berth fuel is not represented in it, which is the limitation set out in Section 3.6.

$$eship = SFOC \times EF / 1000 \qquad [\text{t } CO_2 / \text{MWh}] \qquad (1)$$

Dividing the gross saving by this rate recovers the energy the vessel would otherwise generate on board. That energy, drawn instead from the grid, carries the grid emission rate, and the share of the gross saving lost to the grid, *s*, is:

$$egrid = Igrid / 1000 \qquad [\text{t } CO_2 / \text{MWh}] \qquad (2)$$

$$s = egrid / eship \qquad [–] \qquad (3)$$

Two durations must be distinguished before the balance is written, since the distinction governs everything that follows. Throughout this study d denotes the berthing duration the vessel would have had without shore power. EU-MRV does not report that duration, which is why d is carried as an axis in Table 3 rather than measured. The connection overhead, written t with subscript o, is the connection and disconnection time that shore power adds to d. A connecting vessel therefore occupies its berth for $d + t_o$, not for d. The convention is applied identically in three places: in the static balance below, where to removes a fraction to / d of the saving; in the queueing extension of Section 3.5, where the service time rises from d to $d + t_o$; and in the comparison with the regulatory floor, which is a floor on d rather than on $d + t_o$. Reading to as already contained in the reported duration would remove the second threshold entirely, since a connection that consumes no additional time cannot impose a minimum viable duration. The net saving and the share of the gross saving retained are:

$$NET = E\,(1 - s - to\,/\,d) \qquad [\text{t } CO_2 \text{ / yr}] \tag{4}$$

$$R = 1 - s - to\,/\,d \qquad [-] \tag{5}$$

Neither s nor the ratio to / d depends on vessel type, so two vessels of different segments with the same call duration retain the same share R; within the model, segments differ only in the absolute size of the emission pool E. This is a property of the balance as written, not a general claim about vessel types, and the distinction matters. What the model shows is that vessel type does not enter the retained share once duration, fuel, engine, grid and overhead are held equal. It does not show that vessel type is irrelevant to shore power, and three of the quantities held equal are in fact segment-dependent. Boiler use is the clearest: Osipova and Carraro (2023, Figure 1) report that boilers account for 34% of the at-berth energy demand of container ships and as much as 70% for tankers, so the share of the measured pool that connection can address is not the same across segments. The anchorage emission rate is another, since the power ratios of IMO (2020), Table 17, differ by ship type and exceed unity for several dry cargo classes. Fuel and engine age are a third. The claim defended here is therefore the narrow one: segment enters the balance through the parameters, not through the formula. The three thresholds of Section 3.4 follow directly from equations (1) to (5). All symbols are defined in Table 3 above.

This bears directly on how the research question is framed. Earlier work treats shore power as viable for some segments and not others, and Pruyn and Willeijns (2022) conclude that it does not pay for tankers. The model here says something narrower and more useful: what decides the outcome is the length of the call and the carbon content of the grid, not the label on the ship. A segment effect reported in the literature is a call-duration effect seen through the segment that happens to have those calls.

The parameters of Table 3 are not equally consequential, and the differences matter for how the results should be read. Three points follow. First, the headline break-even is intrinsically a range, because the two parameters that set it, the fuel the auxiliary engine burns and the era in which it was built, take eight documented combinations rather than one. Second, the queueing extension is governed by three parameters (utilisation, berth count and anchorage emission rate) that are individually as influential as one another; none can be treated as a fixed background value, which is why all three are presented as axes rather than assumed. Third, the anchorage emission rate in particular is not a minor conversion detail: at a typical call duration it moves the retained share across a range of more than fifty percentage points, and its lower bound of zero, vessels shutting down entirely while waiting, is precisely the assumption under which the queueing penalty vanishes.

## 3.4 Deriving the three thresholds

The model produces three thresholds rather than three point estimates, and each is stated as a condition a port call either meets or does not. This section derives them from the balance of Section 3.3; the values they take on the 2024 fleet are established in Section 4.2.

The first condition concerns the grid. Replacing auxiliary-engine power with shore power reduces net $CO_2$ only where the electricity supplied is cleaner, per unit of energy, than the fuel it replaces. Setting the retained share of equation (5) to zero and solving for the grid intensity gives the break-even value directly, and that comparison depends on which fuel the auxiliary engine burns, because the fuel fixes both terms of the product: the specific fuel oil consumption of the engine and the emission factor of the fuel. The Fourth IMO GHG Study reports the first by fuel and build era (IMO, 2020, Table 19, p. 71) and the second by fuel (IMO, 2020, Table 21, p. 74). Neither is assigned a value here; the eight combinations those two tables define are carried as an axis. Stolz et al. (2021) derive the same threshold for a single configuration and obtain 784 $gCO_2$ per kWh of electricity, on a basis that includes a generator efficiency of 92%; the figures below are on a mechanical basis and therefore sit under that value by the generator factor, but the construction is identical.

$$Igrid^* = SFOC \times EF \qquad [gCO2 / kWh] \qquad (6)$$

The second is the minimum viable call duration. The grid claims a fixed share of the gross saving whatever the length of the call, while the connection overhead is a fixed duration whose relative weight falls as the call lengthens; setting the retained share to zero and solving for the duration gives the overhead divided by one minus the grid share. The quantity obtained is productive time at berth and not total occupancy, since a connecting vessel holds its berth for the call plus the overhead. The threshold is identical for every segment, because neither term contains a quantity specific to a vessel type.

$$d_{min} = t_o / (1 - s) \qquad \text{[h]} \qquad (7)$$

The third is the critical connection overhead. It is obtained by setting the minimum viable call duration equal to the two-hour boundary below which FuelEU Maritime exempts a call from the obligation to connect, then solving for the overhead. Section 4.2.3 evaluates it against the published range and draws out what it implies.

$$t_{o,crit} = d_{floor} (1 - s) \qquad \text{[h]} \qquad (8)$$

## 3.5 The queueing extension

Connection overhead does not only lengthen the individual call: it reduces berth throughput, which lengthens the queue of vessels waiting at anchorage, and waiting vessels burn fuel.

The berth is modelled as an M/M/c queueing system in Kendall's notation: Markovian (Poisson) arrivals, Markovian (exponential) service and c parallel servers. This is also known as the Erlang C or Erlang delay system. Writing $\lambda$ for the arrival rate and S for the service time of a call, the offered load and the berth utilisation are given by (9), and the mean waiting time at anchorage by the Erlang C delay formula (10). Shore power raises the service time by the connection overhead $t_o$, which raises utilisation proportionally; because waiting time rises non-linearly in utilisation, minutes of overhead can translate into hours of fleet-wide waiting near saturation (Figure 9, Appendix A). The additional waiting per vessel (11) enters the balance weighted by the anchorage emission rate $\alpha$, giving the retained share under congestion (12).

$$a = \lambda S, \quad \rho = a / c \qquad \text{[–]} \qquad (9)$$

$$W_q = C(c, a) \cdot S / [c(1 - \rho)] \qquad \text{[h]} \qquad (10)$$

$$\Delta W = W_q(S + t_o) - W_q(S) \qquad \text{[h]} \qquad (11)$$

$$Rqueue = 1 - s - to / d - \alpha \Delta W / d \qquad [-] \qquad (12)$$

Results are expressed in berth-hour equivalents, so that the absolute at-berth emission rate cancels from the ratio and never requires sourcing. Only the ratio of the anchorage emission rate to the at-berth rate is needed, and that ratio can be read from the Fourth IMO GHG Study, whose Table 17 (IMO, 2020, pp. 68-69) gives auxiliary boiler and auxiliary engine power output for every ship type and size class at berth and at anchor. Converting each output to fuel at its own specific consumption (IMO, 2020, Table 19, p. 70) and taking the ratio of the two modes gives the parameter directly; the fuel emission factor cancels, since both systems burn the same fuel. Some cargo types draw more power at anchor than at berth, so the axis is swept from 10 to 150%.

Three parameters govern this extension: berth utilisation, berth count, and the anchorage emission rate. None is published per port and all three are presented as axes, the third over the range just described. The single-berth case is the worst case and bounds the penalty from above, while parallel berths dampen it sharply.

## 3.6 Limitations

The connection overhead is not measured anywhere. No public dataset reports the time a vessel spends connecting and disconnecting, and the only published range comes from a single study (Zis, 2019, p. 85), which Section 4.2.3 sets out. A spread of that width is not measurement error but the absence of measurement, and it propagates unchanged into any quantity that contains the overhead. Threshold 1 is unaffected, since the break-even grid intensity does not contain it. Thresholds 2 and 3 both do, which is why neither is ever reported as a single value.

The most serious objection to the queueing extension is that connection may occur during existing idle time. If mooring, pilotage or customs formalities already impose dead time at berth, the overhead may not add to service time at all, in which case the effect would vanish. This cannot be resolved with publicly available data and is stated here as a boundary of the analysis.

The queueing model assumes that ships arrive at random and that service times follow an exponential distribution, which is what M/M/c means. Real traffic departs from both. Liner shipping runs partly to schedule, which reduces the randomness and would make waiting shorter than the model says. But arrivals also bunch, weather and tides close windows, and service times differ widely from ship to ship, all of which make waiting longer. The two sets of effects push in opposite directions and there is no basis here for saying which wins, so these figures are reference scenarios rather than bounds. Berth utilisation and berth count are not published port by port either, so the extension describes a class of terminals rather than measuring any one of them.

At-berth $CO_2$ in EU-MRV is not disaggregated by port of call. Individual ports therefore enter the study as illustrative traffic profiles drawn from public throughput figures, never as the basis of the computation.

Boiler emissions are not separated from auxiliary engine emissions, and the gross saving reported here should be read as an upper bound. Zis (2019, p. 87) states the mechanism directly: on a vessel equipped for shore power the auxiliary engines are switched off at berth while the auxiliary boilers continue to operate, keeping fuel temperatures and the main engine cylinders warm against damage from low-temperature contraction, so that the boilers become the only source of pollution left in the port. That function does not depend on the connection, since the vessel must remain ready to depart. Osipova and Carraro (2023, Figure 1) give the magnitude for the European fleet: boilers account for 44% of at-berth $CO_2$ in EU ports, from 34% of the at-berth energy demand of container ships to 70% for tankers, and they note that FuelEU Maritime and the AFIR scope the obligation to electricity demand, which excludes boiler power by construction. The power outputs of IMO (2020), Table 17, pp. 68-69, are consistent with both: for the dry cargo types the boiler output at berth is identical to the output at anchor, which is what a function unrelated to cargo work looks like, whereas for tankers it is several times higher at berth, where cargo heating and steam-driven pumping are added to it.

A second effect runs the other way, and it is smaller. To find the energy the grid must supply, the model divides the whole measured pool by the emission rate of an auxiliary engine. Zis (2019, p. 87) prescribes something narrower: the grid should cover only what the auxiliary engines would have produced. A boiler emits more $CO_2$ for each kilowatt-hour it delivers than an auxiliary engine does, so putting boiler emissions through the engine rate turns them into more grid energy than any ship would draw, and the grid cost comes out too high. That works against the overstatement described above without cancelling it, since the first effect is the larger. Neither correction is applied here, because the gross pool is measured fuel while the shares that would separate engine from boiler are modelled power for representative ship sizes, and multiplying one by the other would push a modelled quantity through every result. Separating them is left to further research, and it would not move the thresholds, in which the auxiliary share cancels between the saving and the overhead.

## 3.7 Reproducibility

Every figure in this study can be reproduced from two public files and the model workbook, an Excel file submitted with this thesis in which each parameter appears once and every downstream table is driven by formula rather than by pasted values. The EU-MRV dataset is downloadable from the THETIS-MRV public portal, which requires a free account but no institutional affiliation; the FuelEU Maritime emission factors are published in the regulation itself; grid carbon intensities are released under a Creative Commons licence. No proprietary, institutional or purchased data is used at any point, and no parameter in the model has been assigned a value that is not either measured or citable.

# 4. Results and analysis

## 4.1 The emission pool and its concentration

Across the twelve cargo segments retained in the sample, the total $CO_2$ emitted at berth in EU ports amounts to 6.15 Mt per year. A large part of this pool is what shore power could address, though not all of it, since the share produced by auxiliary boilers is not displaced by an electrical connection (Section 3.6). The pool is markedly concentrated. Oil tankers and container ships account for almost half of it, 24.6% and 24.4% respectively and the four largest segments, adding Ro-pax and chemical tankers, together represent 76% of the total (Figure 1, Appendix A). The remaining eight segments share the final quarter, with the smallest four each below 2%.

Whether shore power is beneficial at those berths is a separate question, and it is the subject of the three thresholds that follow.

Two things follow from the shape of this pool. The first is that a fleet-wide policy is not needed to reach most of the tonnage: four segments carry three quarters of it, so an instrument aimed at them addresses the bulk of the problem while leaving eight segments and their a*dmin*istrative cost aside. The second cuts the other way. Concentration says where the emissions are, not where they can be removed, and the two need not coincide. Oil tankers head the list, yet they are the segment for which the largest share of at-berth fuel goes to boilers an electrical connection does not replace. The pool is therefore an upper bound on the prize and its ranking is not, on its own, a priority order.

## 4.2 The three thresholds

### 4.2.1 Threshold 1: break-even grid carbon intensity

The break-even follows from equations (1) to (3) and equation (6). Both terms of the auxiliary-engine emission rate are set by the fuel, so fuel type is not a detail of the calculation but one of its axes, and Table 4 gives the eight combinations of fuel and build era the Fourth IMO GHG Study defines. Two conversion efficiencies are deliberately omitted, following Stolz et al. (2021): the ship's alternator efficiency, which would raise the on-board figure, and shore-to-ship transmission losses, which would lower the delivered one. They act in opposite directions and are of comparable magnitude, so the comparison is made on a fuel-energy basis; grid transmission losses are omitted on the same ground, and both omissions are noted in Section 3.6. Across those eight combinations the break-even runs from 429 $gCO_2$/kWh for a recent LNG-fuelled auxiliary to 701 for a pre-1983 engine on heavy fuel oil, a spread of 272 $gCO_2$/kWh. The threshold is therefore a range before it is a number, and which end of it applies is decided on board rather than ashore.

Above the break-even value shore power increases net emissions; below it, it reduces them. For France, Belgium, the Netherlands and Germany the conclusion is unambiguous across all eight combinations: every one of those grids sits below even the lowest break-even estimate (grid intensities from Ember, 2025, and the European Environment Agency, 2025). The conditional case is Poland, the most carbon-intensive grid in the Union at 566 $gCO_2$/kWh in 2024, and it is conditional in a specific way. Setting the break-even equal to the Polish grid and solving for the auxiliary-engine consumption gives a continuous threshold of 566 / 3.114, or 181.8 g/kWh: any auxiliary burning more than that per kilowatt-hour is dirtier than the Polish grid and benefits from connecting, any auxiliary below it does not. Among the eight documented configurations only the two LNG-fuelled ones fall below the threshold; the six oil-fuelled configurations lie above it and gain from connecting even on the dirtiest grid in the Union (Figure 2, Appendix A). Two patterns run through Table 4, and both cut against intuition. The more recent the auxiliary engine, the lower the break-even, because a more efficient engine is a harder benchmark for the grid to beat. And the cleaner the fuel, the lower the break-even again: an auxiliary running on LNG breaks even at 429

$gCO_2$/kWh rather than 607, so it needs a substantially cleaner grid before connecting is worthwhile. Shore power is least attractive precisely where the vessel has already decarbonised its own power generation.

*Table 4. The eight combinations of auxiliary-engine fuel and build era documented by the Fourth IMO GHG Study, with the break-even grid carbon intensity and the minimum viable call duration each implies. SFOC from IMO (2020), Table 19, p. 71; emission factors from Table 21, p. 74. The minimum viable call duration is computed at the EU average grid of 213 $gCO_2$/kWh (Ember, 2025) and 90 minutes of connection overhead. No combination is treated as central. The last column but one gives the critical connection overhead, the value at which the minimum viable call duration equals the two-hour exemption boundary.*

| **Fuel** | **Build era** | *SFOC* (g/kWh) | *EF* (t $CO_2$/t) | **Break-even ($gCO_2$/kWh)** | **d min at EU average (h)** | **t crit (min)** | **Poland (566)** |
|---|---|---|---|---|---|---|---|
| HFO | pre-1983 | 225 | 3.114 | 701 | 2.16 | 84 | reduces |
| MDO | pre-1983 | 210 | 3.206 | 673 | 2.19 | 82 | reduces |
| HFO | 1984-2000 | 205 | 3.114 | 638 | 2.25 | 80 | reduces |
| MDO | 1984-2000 | 190 | 3.206 | 609 | 2.31 | 78 | reduces |
| HFO | 2001+ | 195 | 3.114 | 607 | 2.31 | 78 | reduces |
| MDO | 2001+ | 185 | 3.206 | 593 | 2.34 | 77 | reduces |
| LNG | 1984-2000 | 173 | 2.750 | 476 | 2.72 | 66 | increases |
| LNG | 2001+ | 156 | 2.750 | 429 | 2.98 | 60 | increases |

The result is more discriminating than the debate it enters. Shore power is often defended or attacked as though its value depended on the national grid, and for most of Europe that question is closed: every combination of fuel and build era documented by the IMO clears the French, Belgian, Dutch and German grids by a wide margin, so the electricity is not what decides the outcome there. Where the grid does decide, it decides against the cleanest vessels rather than the dirtiest. An auxiliary engine burning LNG generates power at 429 $gCO_2$/kWh equivalent, so it needs an unusually clean grid before the exchange pays, while a pre-1983 engine on heavy fuel oil gains on any grid in the Union. The counter-intuitive consequence is that the vessels with the worst on-board efficiency are the best candidates for connection, because the standard they set for the shore is the lowest.

### 4.2.2 Threshold 2: minimum viable berthing time

The second condition concerns the duration of the call, and Section 3.4 gives its derivation. On the EU average grid of 213 $gCO_2$/kWh (Ember, 2025) the grid share is about 33%, so roughly 67% of the gross saving is left to absorb the connection overhead.

At 90 minutes of overhead, roughly the highest point of the range (Zis, 2019), the retained share turns positive only beyond a berthing time of 2.25 hours (Figure 3, Appendix A). Below that duration the overhead consumes the entire saving and shore power becomes counterproductive.

Above it, the benefit rises steeply at first and then flattens: a four-hour call retains 25% of the gross saving, a twelve-hour call 50%, and a twenty-four-hour call 56%. The curve is identical for every segment: two vessels of different types with the same berthing time retain the same share, and the differences between segments in absolute tons reflect the size of their emission pool rather than any difference in the efficiency of shore power.

Because the retained share is the same for every segment, the net saving in tons is simply that share applied to each segment's emission pool. Table 5 reports it for the whole cargo fleet across the call-duration axis. The columns are the axis; the rows are measured.

Table 5. Net annual $CO_2$ saved by shore power, in kilotons per year, by cargo segment and call duration. Gross is the measured at-berth $CO_2$ from EU-MRV 2024. Each duration column applies the retained share at that call length, computed at 90 minutes of connection overhead, the EU average grid of 213 $gCO_2/kWh$, an auxiliary *SFOC* of 205 g/kWh and a fuel emission factor of 3.114 t $CO_2$ per tons, that is heavy fuel oil in an engine built between 1984 and 2000. A cleaner fuel raises the break-even and therefore lowers these figures. Negative values indicate calls at which connecting increases emissions.

| Ship type | Gross | 2 h | 4 h | 8 h | 12 h | 24 h | 36 h |
|---|---|---|---|---|---|---|---|
| Oil tanker | 1,512 | -126 | 440 | 724 | 818 | 913 | 944 |
| Container ship | 1,501 | -126 | 437 | 719 | 812 | 906 | 937 |
| Ro-pax ship | 974 | -81 | 284 | 466 | 527 | 588 | 608 |
| Chemical tanker | 664 | -56 | 193 | 318 | 359 | 401 | 415 |
| Bulk carrier | 496 | -42 | 145 | 238 | 269 | 300 | 310 |
| General cargo ship | 255 | -21 | 74 | 122 | 138 | 154 | 159 |
| Ro-ro ship | 221 | -18 | 64 | 106 | 120 | 133 | 138 |
| LNG carrier | 185 | -15 | 54 | 89 | 100 | 112 | 116 |
| Gas carrier | 127 | -11 | 37 | 61 | 69 | 77 | 80 |
| Vehicle carrier | 118 | -10 | 34 | 56 | 64 | 71 | 74 |
| Container/ro-ro cargo ship | 61 | -5 | 18 | 29 | 33 | 37 | 38 |
| Refrigerated cargo carrier | 33 | -3 | 10 | 16 | 18 | 20 | 20 |
| **Cargo fleet total** | **6,147** | **-514** | **1,791** | **2,943** | **3,327** | **3,712** | **3,840** |

Three things are visible in the table. The first is the sign. At a two-hour call every segment is negative, and the fleet increase its emission by 514 kt of $CO_2$ a year rather than saving any: the shortest call the regulation still compels, two hours, sits below the duration at which connection begins to pay. The obligation itself reaches only containerships and passenger ships, so among the segments in this table it bites on container ships and ro-pax vessels; the remaining rows show what the same arithmetic would imply were the obligation extended to them. The second is the shape. The saving rises steeply to about eight hours and then flattens, because the overhead is a fixed duration whose weight falls as the call lengthens while the grid share does not move. Between two and twelve hours the fleet total gains 3,842 kt; between twelve and thirty-six it gains only a further 512.

The third is the one that matters for policy. The ranking of segments is identical in every column, and it is identical to the ranking of the gross pool in Figure 1. That is not a coincidence but a consequence of the model: since the retained share does not depend on vessel type, each column is the same percentage applied to twelve different pools. Oil tankers and container ships together account for 1,631 kt of the 3,327 kt the fleet retains at a twelve-hour call, which is the same 49% share they hold of the gross pool. A port choosing where to install shore power is therefore choosing between pools, not between efficiencies, and the four leading segments carry three quarters of what is available at any call duration.

Read as a condition rather than as a table, this threshold is the one a port can act on. The grid intensity of a country changes slowly and is outside a port's control; the duration of a call is observed at every berth, every day, in data the port already holds. Threshold 2 therefore converts the question from one about national energy policy, on which a port has no leverage, into one about its own traffic, on which it has both information and choice. That is also what makes the two-hour figure in the regulation testable: it is a claim about the same quantity.

### 4.2.3 Threshold 3: the critical overhead and the two-hour boundary

The minimum viable berthing time depends on the connection overhead. Zis (2019, p. 85) provides the only published range: terminal operators at the Port of Gothenburg report that connection and disconnection can be completed in under ten minutes, but the paper notes that this applies at an optimal setting with no complications, and that one to one and a half hours is more realistic once the coordination required between ship and port crews is accounted for. Martínez-López et al. (2021) assume one hour of connection and disconnection lag time in their own assessment, which corroborates the lower end of that realistic range. Setting the minimum viable berthing time at the EU average grid carbon intensity equal to the two-hour exemption boundary and solving for the overhead gives critical values between 84 minutes with an oil-fuelled (pre-1983 era) and 60 minutes with an LNG-fuelled (2001+ era), the overhead at which the boundary and viability coincide.

This third threshold is of a different kind from the first two, and the difference should be stated plainly. Thresholds 1 and 2 are properties of the technology: they follow from the fuel, the engine and the grid, and they would hold whatever the law said. Threshold 3 is more a property of the regulation, obtained by setting the minimum viable duration equal to the boundary the legislator chose and solving for the only quantity left free. It says less about ships and more about the fit between a legal text and the physics it presumes.

What it says is that the fit cannot be verified, and Table 4 shows why the problem is worse than a single number suggests. The critical overhead is not one value but eight, one per fuel and build era, running from 60 minutes for a recent LNG-fuelled auxiliary to 84 for a pre-1983 engine on heavy fuel oil. That span very nearly coincides with the 60 to 90 minutes the only published source describes as realistic. The pivot is therefore not a point inside the plausible range but a range that overlaps it almost entirely, which has a consequence the single-value reading hides: whatever the true overhead turns out to be, it will sit above the pivot for some vessels and below it for others. A fixed boundary of two hours cannot be correct for the whole fleet at once. It compels the cleanest vessels to connect on calls that lose $CO_2$ while releasing the dirtiest from calls that would have gained, and no measurement of the overhead, however precise, would remove that.

## 4.3 The queueing effect

The queueing penalty is governed by three parameters, none published per port and each individually decisive, so the model sweeps all three. Tables 6, 7 and 8 report one sweep each, sharing the call duration as their row axis. Each varies one parameter along its columns while holding the other two at the values of Table 3. Every entry is the share of the gross saving retained once anchorage waiting is counted, so a negative entry means shore power increases net emissions and SAT means the overhead has pushed berth occupancy to or beyond capacity.

The first is the number of berths operating in parallel, which sets how much of the added service time a terminal can absorb before vessels begin to wait.

Table 6. Share of the gross saving retained, by call duration and number of parallel berths, at 75% berth utilisation, an anchorage emission rate of 50% of the at-berth rate, 90 minutes of connection overhead, a grid of 213 $gCO_2/kWh$, *SFOC* 205 g/kWh and emission factor 3.114.

| **Call duration** | **1** | **2** | **3** | **4** | **5** | **8** | **10** | **15** |
|---|---|---|---|---|---|---|---|---|
| 2 h | SAT | SAT | SAT | SAT | SAT | SAT | SAT | SAT |
| 4 h | SAT | SAT | SAT | SAT | SAT | SAT | SAT | SAT |
| 6 h | -746% | -348% | -215% | -149% | -110% | -50% | -31% | -5% |
| 8 h | -286% | -116% | -59% | -31% | -14% | 11% | 19% | 30% |
| 12 h | -100% | -21% | 6% | 19% | 26% | 38% | 42% | 47% |
| 16 h | -42% | 9% | 26% | 35% | 40% | 47% | 50% | 53% |
| 24 h | 2% | 32% | 42% | 47% | 50% | 55% | 56% | 58% |
| 36 h | 26% | 45% | 51% | 54% | 56% | 59% | 60% | 61% |

Parallel berths dampen the penalty but do not remove it. At a twelve-hour call, a single-berth terminal retains minus 100%, which is to say that connecting destroys the saving and as much again. Two berths bring it to minus 21%, three cross into positive territory at 6%, five reach 26%, ten reach 42% and fifteen 47%. The returns fall away quickly: the step from one berth to three is worth more than the step from three to fifteen.

The top rows matter more than that gradient. A four-hour call saturates at every column of the table, from one berth to fifteen. This is not a limitation of the sample but a property of the model: the inflated utilisation is the utilisation multiplied by one plus the ratio of overhead to call duration, and the number of berths appears nowhere in that expression. Capacity changes how long vessels wait; it does not change whether the system is stable. A terminal handling short calls at high utilisation cannot build its way out of the problem, and must shorten the connection overhead instead.

The second is berth utilisation, which is the parameter the model is most sensitive to and the one a port has least direct control over in the short run.

Table 7. Share of the gross saving retained, by call duration and berth utilisation, at five parallel berths, an anchorage emission rate of 50%, 90 minutes of connection overhead, a grid of 213 $gCO_2/kWh$, *SFOC* 205 g/kWh and emission factor 3.114.

| Call duration | 40% | 50% | 60% | 70% | 75% | 80% | 85% | 90% |
|---|---|---|---|---|---|---|---|---|
| 2 h | -29% | -105% | SAT | SAT | SAT | SAT | SAT | SAT |
| 4 h | 25% | 16% | -12% | -291% | SAT | SAT | SAT | SAT |
| 6 h | 39% | 35% | 24% | -17% | -110% | SAT | SAT | SAT |
| 8 h | 46% | 44% | 37% | 17% | -14% | -133% | SAT | SAT |
| 12 h | 53% | 52% | 48% | 39% | 26% | -4% | -132% | SAT |
| 16 h | 57% | 56% | 53% | 47% | 40% | 23% | -28% | -539% |
| 24 h | 60% | 59% | 58% | 54% | 50% | 42% | 20% | -80% |
| 36 h | 62% | 62% | 61% | 59% | 56% | 51% | 40% | -3% |

The boundary between the numbers and the SAT region is the result. A four-hour call has already fallen to minus 291% at 70% utilisation and saturates at 75%. A twelve-hour call retains 52% at 50% utilisation, 26% at 75%, turns negative at 80% (minus 4%), collapses to minus 132% at 85% and saturates at 90%. Only calls of twenty hours or more stay positive across most of the range.

Read against Table 5, this is the tension at the centre of the managerial argument. The segments holding the largest emission pools are handled at the busiest terminals, and the busiest terminals are precisely where the connection overhead does most damage. Berth capacity and shore power are not independent investments.

The third is the emission rate of a vessel waiting at anchorage, expressed as a share of its rate at berth. It is not published for any European port, but it can be derived from IMO (2020), Table 17, which reports boiler and auxiliary power in both modes for every cargo class. It is the parameter that decides whether the queueing effect exists at all.

Table 8. Share of the gross saving retained, by call duration and anchorage emission rate as a share of the at-berth rate, at five parallel berths, 75% berth utilisation, 90 minutes of connection overhead, a grid of 213 $gCO_2/kWh$, *SFOC* 205 g/kWh and emission factor 3.114. The axis runs past 100% because some cargo types draw more power at anchor than at berth.

| Call duration | 10% | 20% | 40% | 50% | 80% | 100% | 120% | 150% |
|---|---|---|---|---|---|---|---|---|
| 2 h | SAT | SAT | SAT | SAT | SAT | SAT | SAT | SAT |
| 4 h | SAT | SAT | SAT | SAT | SAT | SAT | SAT | SAT |
| 6 h | 11% | -19% | -79% | -110% | -200% | -261% | -321% | -412% |
| 8 h | 35% | 23% | -2% | -14% | -51% | -76% | -101% | -138% |
| 12 h | 49% | 43% | 32% | 26% | 10% | -1% | -12% | -29% |
| 16 h | 54% | 50% | 43% | 40% | 29% | 22% | 15% | 4% |
| 24 h | 58% | 56% | 52% | 50% | 44% | 40% | 36% | 30% |
| 36 h | 61% | 60% | 57% | 56% | 53% | 50% | 48% | 44% |

The response is exactly linear in this parameter, because the anchorage rate multiplies the waiting term and enters the balance nowhere else. For a twelve-hour call each additional tenth costs 5.5 percentage points, carrying the retained share from 49% at the bottom of the axis, through 26% at half the at-berth rate, to minus 1% at parity and minus 29% at the top. A reader holding a different value can therefore interpolate directly rather than re-running the model. Shorter calls are steeper: a six-hour call moves from 11% to minus 412% across the same range.

The right-hand columns deserve attention. A ship that emits as much at anchor as at berth, or more, is not a hypothetical case. For several dry cargo types the power IMO reports at anchor exceeds the power at berth, because their boiler load is the same in both modes while their auxiliary demand rises. For those segments the queueing penalty does not reduce the saving, it reverses it, and the columns past 100% are where they sit. The studies in Section 2 that treat shore power as a simple substitution at the berth place every ship implicitly in the leftmost column, and the width of this table is what the queueing extension adds.

Taken together, the three sweeps say something the static literature cannot. In every other part of this study the worst case is a call that saves nothing; here the worst case is a call that costs more than it saves and, beyond saturation, a terminal that stops functioning. The difference matters for how the result should be used. A negative retained share is an emissions problem and can be traded against other measures; saturation is an operational failure and cannot. It is also the only place in this thesis where the model returns no number at all, and that silence is the result rather than a gap in it.

## 4.4 Sensitivity to the open axes

The model assigns no value to any parameter that no public source fixes, and the sections above report each of those parameters one at a time. That discipline is only useful if the reader can see what each one does. Figure 6 collects all seven axes in a single view, sweeping each across its plausible range while the others are held at a reference case.

Two patterns emerge. The first is that the parameters differ in shape, not only in weight. Grid intensity, connection overhead and the anchorage emission rate enter the balance linearly, so a given error in any of them produces a proportional error in the result. Call duration, berth count and berth utilisation are non-linear, and steeply so at their unfavorable ends: the penalty at a saturated berth is not a bigger version of the penalty at a busy one, it is an order of magnitude worse. The practical consequence is simple. A port assessing shore power can be approximate about its electricity and must be precise about its congestion.

The second pattern is a ranking, and it runs against the emphasis of the existing literature. Figure 7 sweeps each axis across its full plausible range and records how far the retained share moves, ordering the parameters by that amplitude.

The parameter on which this study was most exposed, the fuel and build era of the auxiliary engine, turns out to matter least. Across all eight combinations the Fourth IMO GHG Study defines, the retained share moves by twenty-two percentage points. The parameters that dominate are operational, and none of them is published in Europe: call duration and berth utilisation each move the result by around two hundred percentage points, connection overhead by more than a hundred and seventy, and berth count by close to a hundred and fifty. Grid carbon intensity, which the shore-power literature treats as the central question, ranks fifth of seven.

Two of those axes can be read off directly, because together they fix the second threshold. Table 9 gives the minimum viable call duration for every combination of connection overhead and grid carbon intensity.

Table 9. Minimum viable call duration in hours, by connection overhead and grid carbon intensity in $gCO_2/kWh$, at *SFOC* 205 g/kWh and emission factor 3.114. Below the duration shown, connecting increases net emissions. A cleaner fuel, such as LNG at an emission factor of 2.750, lengthens every entry.

| Overhead (min) | 50 | 150 | 250 | 350 | 450 | 550 | 600 |
|---|---|---|---|---|---|---|---|
| 15 | 0.27 | 0.33 | 0.41 | 0.55 | 0.85 | 1.81 | 4.16 |
| 30 | 0.54 | 0.65 | 0.82 | 1.11 | 1.69 | 3.61 | 8.32 |
| 45 | 0.81 | 0.98 | 1.23 | 1.66 | 2.54 | 5.42 | 12.48 |
| 60 | 1.08 | 1.31 | 1.64 | 2.21 | 3.39 | 7.22 | 16.64 |
| 75 | 1.36 | 1.63 | 2.05 | 2.77 | 4.24 | 9.03 | 20.80 |
| 90 | 1.63 | 1.96 | 2.47 | 3.32 | 5.08 | 10.84 | 24.96 |
| 120 | 2.17 | 2.61 | 3.29 | 4.43 | 6.78 | 14.45 | 33.27 |
| 180 | 3.25 | 3.92 | 4.93 | 6.64 | 10.17 | 21.67 | 49.91 |

The table is the operational form of Threshold 2. A port reads down its own overhead and across its own grid intensity. At ninety minutes of overhead, the minimum viable call is 1.63 hours on a grid of 50 $gCO_2/kWh$, 2.47 hours at 250, and 10.84 hours at 550. Halving the overhead to thirty minutes brings the same 250 $gCO_2/kWh$ case down to 0.82 hours. The two-hour boundary therefore sits too high over a large part of the table and too low in the upper right of it, and which region a given port falls into depends on a quantity that is not measured.

This bears on where further work should go. The question of whether shore power reduces emissions has largely been debated as a question about electricity generation, and Figure 7 suggests that debate is settled for most of Europe and secondary everywhere else. What is not settled, and what no public dataset addresses, is how long vessels actually stay at berth, how heavily those berths are used, and how much fuel a vessel burns while waiting for one. Those three quantities decide the answer, and all three are measurable at modest cost.

The ranking in Figure 7 also explains why the thresholds are reported as conditions. If the influential parameters were the sourced ones, the model could return a value and defend it. They are not: the four axes that move the result most are operational quantities no European source publishes, and three of them are known to the port and to nobody else. A single figure computed on assumed values for those axes would be precise and unfounded. A condition computed across their ranges is less satisfying and can be checked, which is the trade this study accepts throughout.

# 5. Conclusion

Shore power reduces net $CO_2$ for a European cargo call when three conditions hold at once: the grid behind the plug is cleaner per unit of energy than the auxiliary engine it replaces, the call is long enough to absorb the time lost to connecting, and the berth is not so busy that this lost time pushes other vessels into anchorage waiting. These are not three findings but three ways for one balance to fail, which is why the answer is a set of conditions rather than a verdict. Chapter 4 established the magnitudes. This chapter sets out what the work contributes to the literature, what it offers to policy, where it is open to criticism, and what should be done next.

## 5.1 Contributions to the literature

The first contribution is the integration itself. Table 1 showed that the three quantities governing the emission balance are each treated seriously somewhere in the literature and never together: grid carbon intensity where the generation mix is the object of study, connection time where the concern is the schedule, congestion where waiting is measured from AIS data. No strand carries all three through to a condition a port could apply. This study closes that gap by writing the balance once and solving it for the point at which it changes sign, which yields three thresholds a call must clear: a break-even grid intensity, a minimum viable call duration, and a critical connection overhead. Stating the result as thresholds rather than as percentages is what makes it portable, because a threshold survives a change of fuel, grid or port while a percentage does not. One consequence of writing the balance this way is worth recording: vessel segment does not enter it. Equations (3) and (5) contain no quantity specific to a ship type, so two vessels of different segments making the same call at the same berth keep the same share of their gross saving. Segment governs the size of the pool, not the efficiency with which shore power addresses it, which narrows rather than contradicts the segment-level conclusions of Pruyn and Willeijns (2022).

The second contribution is the congestion channel, and it is the one the literature had not quantified. Connection overhead does not merely lengthen the call that incurs it: it reduces berth throughput, and because waiting rises non-linearly in utilisation, minutes of overhead become hours of fleet-wide waiting near saturation. Modelling the berth as an M/M/c system and charging the additional waiting at the anchorage emission rate turns that mechanism into a number. The result is the least intuitive of the study: at a twelve-hour call and 75% berth utilisation a single berth loses more than the entire saving, five berths retain 26% and ten retain 42%, and past roughly 80% utilisation the balance turns negative even though every individual call clears the duration test. Akakura (2023) and Rhodes et al. (2025) measure waiting and its emissions but attribute none of it to a connection procedure; the berth-allocation literature models the trade-off but only inside terminal-specific optimisations. Neither yields a condition that transfers. This one does.

The third contribution is the parameter architecture, and it is what makes the first two reusable. Every quantity in the model belongs to exactly one of three categories: measured from EU-MRV, sourced from a citable authority, or carried across an explicit range because no public value exists. No fourth category is admitted, so an assumption cannot enter the model without becoming visible as a range. Section 3.7 sets out what that buys in reproducibility; what it buys here is transferability. Because the categories are declared rather than the values fixed, the framework moves to another fleet, another region or another regulatory regime by substituting the measured pool and re-reading the same three thresholds, and a reader who disagrees with an axis can interpolate a different answer without rebuilding anything. That construction has a cost, since results computed on open axes are bounds rather than estimates. A model calibrated on terminal data answers one question at one port; this one answers the same question wherever the public data exist.

## 5.2 Policy contributions

For a port authority, the three thresholds combine into a decision rule that needs no data the port does not already hold. Figure 8 plots the retained share over the two quantities that govern it, call duration and grid carbon intensity, with the break-even boundary drawn on top. A port reads across at its own grid intensity and obtains the call duration below which connecting does harm: for an oil-fuelled auxiliary built after 2000 with 90 minutes of overhead is 1.59 hours in France and 2.88 hours in Germany, and on the Polish grid an LNG-fuelled vessel of the same age never breaks even at any duration. Berths whose traffic sits mostly above that duration are candidates and berths below it are not, whatever segment they serve. Combined with the concentration result, four segments holding 76% of the pool, this gives a two-stage priority rule: go where the tons are, then check that the calls there are long enough to capture them. The congestion results add a third filter that cuts against the first two, since the terminals holding the most tons are also those where the overhead does the most damage to throughput. A port operating above 80% berth utilisation should treat additional berth capacity as a precondition for shore power rather than as a separate project.

For a vessel operator the same logic applies to the retrofit decision, which is governed by the operator's own call-duration profile rather than by the class of the vessel, and which can be computed from its own port logs.

For the regulator, the operative result is the critical overhead, which cuts both ways. Above the critical overhead call needs more than two hours to be worth connecting while the obligation already applies from two hours, so the regulation compels connection over a band of losing calls, for example: at ninety minutes of overhead, from 2 to 2.25 hours. Below it the error reverses, and calls that would gain sit inside the exemption and are released from the obligation. Eighty minutes is therefore one of the overheads at which the two-hour figure is right, and it falls inside the 60 to 90 minutes Zis (2019, p. 85) describes as realistic, so the literature cannot say whether the boundary is set too high or too low. Three recommendations follow. Connection overhead should be measured and published, since it is the one quantity that fixes the pivot and no public source reports it. The exemption should then be recomputed rather than fixed: the viable duration is the overhead divided by one minus the grid share, so it moves with the grid intensity of the port and with the fuel and build era of the auxiliary engine as well as with the overhead itself. A single figure in hours cannot track three axes at once, which is why it is wrong on one side or the other almost everywhere; a short table or formula over those axes, of the kind set out in Table 9, would tell a master when connecting pays and when it does not. And the mandate should recognise berth congestion, because requiring connection at a saturated terminal can raise fleet emissions even where every call clears the duration test.

## 5.3 Critique

The most important limitation is what the model does not contain. It is an emission balance and carries no cost, no tariff and no capital expenditure, so it says nothing about whether shore power is a good investment. That was deliberate, since a call which increases emissions cannot be rescued by a favourable tariff and the prior question had to be settled first, but it means the results identify where connecting is physically worthwhile and not where it is financially justified. The literature reviewed in Section 2.2 answers the second question and this study does not; the two are complementary and neither substitutes for the other.

Three further boundaries constrain the claims directly. The connection overhead drives the regulatory finding and rests on a published range rather than a measurement, which is why that finding is stated as a condition and never as a value. The congestion result assumes that the overhead adds to service time rather than overlapping with idle time the call already contains; no public source settles how much overlaps, and if the overlap were substantial the penalty would shrink or disappear. And EU-MRV reports at-berth $CO_2$ as one figure without separating boiler from auxiliary engine, so the gross saving is an upper bound wherever boiler use is high, which Osipova and Carraro (2023) put at 44% of all at-berth $CO_2$ in EU ports.

Two features of the method deserve the same scrutiny. The design is a product of its data constraint rather than a free choice: three earlier directions failed on access to berth-level data, and the fleet-level model with open axes was the response. That constraint became the methodological identity of the work, but it also fixed what the study could not examine, and a reader with terminal access would design it differently. The parameter architecture also had to be corrected twice during development. An early version computed segment results from per-vessel averages, which mix vessels with entirely different EU exposure and describe no real vessel; and the auxiliary fuel was carried at the heavy fuel oil value without that choice being declared, so a parameter labelled as sourced in fact contained a hidden assumption. Both were corrected, the second by making fuel an axis. They are recorded here because they are the same mistake twice: a value had been chosen for a good reason, the reason was not written down, and the parameter then looked sourced when it was assumed. The three-category architecture exists to make that impossible, and neither error was caught by the architecture itself. Both were found by re-reading the model against the data it claimed to use, which is the check the architecture cannot perform on its own.

## 5.4 Further work

Four extensions follow, ordered by how much they would change and how little they would cost. The first is to measure the connection overhead. It is the parameter on which the regulatory finding turns, no public source reports it, and a terminal timing connection and disconnection across a season of calls would convert the widest axis in this study into a measured quantity. Such a campaign should record ship type and size alongside each timing, since the number of cable connections, the height of the freeboard and the crew procedures all differ by vessel, and the overhead is unlikely to be a single figure for the whole fleet. It is the cheapest of the four extensions and more of the conclusions depend on it than on anything else.

The second is berth occupancy and queue length. Ports hold both, neither is published in a comparable form, and without them the congestion results remain scenarios rather than estimates.

The third is the split between auxiliary engine and boiler in at-berth fuel: IMO (2020), Table 17, pp. 68-69, already publishes the two power outputs separately by ship type and operating mode, and applying that separation at vessel level would remove the largest single uncertainty in the size of the addressable pool. Neither of these requires new instrumentation; both require a reporting decision.

The fourth is validation. The berth-allocation literature holds the terminal-level schedules this study lacks, and testing the M/M/c bounds against an observed berth schedule would establish whether the congestion effect operates in practice or is absorbed by existing idle time. That test would settle the one assumption on which the least intuitive result of this thesis depends. Extending the scope beyond cargo traffic, to the passenger and cruise segments excluded in Section 3.2, would follow naturally once the overhead is measured, since those segments differ in call duration and hotel load rather than in the structure of the balance.

# References

Abu Bakar, N. N., Bazmohammadi, N., Çimen, H., Uyanik, T., Vasquez, J. C., & Guerrero, J. M. (2022). Data-driven ship berthing forecasting for cold ironing in maritime transportation. *Applied Energy, 326*, 119947. https://doi.org/10.1016/j.apenergy.2022.119947

Abu Bakar, N. N., Bazmohammadi, N., Vasquez, J. C., & Guerrero, J. M. (2023). Electrification of onshore power systems in maritime transportation towards decarbonization of ports: A review of the cold ironing technology. *Renewable and Sustainable Energy Reviews, 178*, 113243. https://doi.org/10.1016/j.rser.2023.113243

Akakura, Y. (2023). Analysis of offshore waiting at world container terminals and estimation of $CO_2$ emissions from waiting ships. *Asian Transport Studies, 9*, 100111. https://doi.org/10.1016/j.eastsj.2023.100111

Alamoush, A. S., Ballini, F., & Ölçer, A. I. (2021). Revisiting port sustainability as a foundation for the implementation of the UN SDGs. *Journal of Shipping and Trade, 6*(1), 1–52. https://doi.org/10.1186/s41072-021-00101-6

Bignucolo, F., Visentin, M., De Pieri, D., Augello, C., & Faggioni, N. (2025). Technical and economic feasibility of cold ironing in Italy. *Energies, 18*(22), 5950. https://doi.org/10.3390/en18225950

Canepa, M., Ballini, F., Dalaklis, D., Frugone, G., & Sciutto, D. (2023). Cold ironing: Socio-economic analysis in the port of Genoa. *Logistics, 7*(2), 28. https://doi.org/10.3390/logistics7020028

Comer, B., & Carvalho, F. (2023, July 7). *IMO's newly revised GHG strategy: What it means for shipping and the Paris Agreement*. International Council on Clean Transportation. https://theicct.org/marine-imo-updated-ghg-strategy-jul23/

Ember. (2025). *European electricity review 2025*. https://ember-energy.org/latest-insights/european-electricity-review-2025/

European Commission. (n.d.). *Questions and answers on Regulation (EU) 2023/1805 on the use of renewable and low-carbon fuels in maritime transport, and amending Directive 2009/16/EC*.

Retrieved August 8, 2026, from https://transport.ec.europa.eu/transport-modes/maritime/decarbonising-maritime-transport-fueleu-maritime/questions-and-answers-regulation-eu-20231805-use-renewable-and-low-carbon-fuels-maritime-transport_en

European Commission. (2023). *Regulation (EU) 2023/1805 on the use of renewable and low-carbon fuels in maritime transport (FuelEU Maritime)*. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32023R1805

European Environment Agency. (2025). *Greenhouse gas emission intensity of electricity generation in Europe*. https://www.eea.europa.eu/en/analysis/indicators/greenhouse-gas-emission-intensity-of-1

Eurostat. (2026). *Maritime transport statistics: Short sea shipping of goods*. European Commission.https://ec.europa.eu/eurostat/statisticsexplained/index.php?title=Maritime_transport_statistics_-_short_sea_shipping_of_goods

Glavinović, R., Krčum, M., Vukić, L., & Karin, I. (2023). Cold ironing implementation overview in European ports: Case study of Croatian ports. *Sustainability, 15*(11), 8472. https://doi.org/10.3390/su15118472

Hu, Z. (2022). Carbon emission reduction of shore power from power energy structure in China. *Frontiers in Marine Science, 9*, 1077289. https://doi.org/10.3389/fmars.2022.1077289

IMO. (2020). *Fourth IMO greenhouse gas study 2020*. International Maritime Organization. https://www.imo.org/en/OurWork/Environment/Pages/Fourth-IMO-Greenhouse-Gas-Study-2020.aspx

IMO. (2025). *IMO net-zero shipping talks to resume in 2026* [Press briefing]. International Maritime Organization. https://www.imo.org/en/mediacentre/pressbriefings/pages/imo-net-zero-shipping-talks-to-resume-in-2026.aspx

Innes, A., & Monios, J. (2018). Identifying the unique challenges of installing cold ironing at small and medium ports: The case of Aberdeen. *Transportation Research Part D: Transport and Environment, 62*, 298–313. https://doi.org/10.1016/j.trd.2018.02.004

Kelmalis, A., Dimou, A., Lekkas, D. F., & Vakalis, S. (2024). Cold ironing and the study of RES utilization for maritime electrification on Lesvos Island port. *Environments, 11*(4), 84. https://doi.org/10.3390/environments11040084

Le, S.-T. (2024). Research on drivers and barriers to the implementation of cold ironing technology in zero emissions port. *Environmental Health Insights, 18*, 11786302241265090. https://doi.org/10.1177/11786302241265090

Martínez-López, A., Romero, A., & Orosa, J. A. (2021). Assessment of cold ironing and LNG as mitigation tools of short sea shipping emissions in port: A Spanish case study. *Applied Sciences, 11*(5), 2050. https://doi.org/10.3390/app11052050

Notteboom, T., Van der Lugt, L., Van Saase, N., Sel, S., & Neyens, K. (2020). The role of seaports in green supply chain management: Initiatives, attitudes, and perspectives in Rotterdam, Antwerp, North Sea Port, and Zeebrugge. *Sustainability, 12*(4), 1688. https://doi.org/10.3390/su12041688

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. https://theicct.org/publication/shore-power-eu-oct23/

Peddi, K. P., Ricci, S., & Rizzetto, L. (2024). Reduction potential of gaseous emissions in European ports using cold ironing. *Applied Sciences, 14*(15), 6837. https://doi.org/10.3390/app14156837

Pruyn, J., & Willeijns, J. (2022). Cold ironing: Modelling the interdependence of terminals and vessels in their choice of suitable systems. *Journal of Shipping and Trade, 7*(1), 17. https://doi.org/10.1186/s41072-022-00119-4

Rhodes, R., Leiphardt, C., Young, H. S., Morten, J., Hayes, B., Dillon, J., Louttit, W., Powell, M., & McCauley, D. J. (2025). Investigation of a port queuing system on $CO_2$ emissions from container shipping. *Marine Pollution Bulletin, 218*, 118151. https://doi.org/10.1016/j.marpolbul.2025.118151

Stolz, B., Held, M., Georges, G., & Boulouchos, K. (2021). The $CO_2$ reduction potential of shore-side electricity in Europe. *Applied Energy, 285*, 116425. https://doi.org/10.1016/j.apenergy.2020.116425

Williamsson, J., Costa, N., Santén, V., & Rogerson, S. (2022). Barriers and drivers to the implementation of onshore power supply: A literature review. *Sustainability, 14*(10), 6072. https://doi.org/10.3390/su14106072

Yuksel, O., Bayraktar, M., & Seyhan, A. (2025). Environmental and economic analysis of cold ironing using renewable hybrid systems. *Clean Technologies and Environmental Policy, 27*, 3489–3517. https://doi.org/10.1007/s10098-024-03065-w

Zis, T. P. V. (2019). Prospects of cold ironing as an emissions reduction option. *Transportation Research Part A: Policy and Practice, 119*, 82–95. https://doi.org/10.1016/j.tra.2018.11.003

# Appendix A. Figures

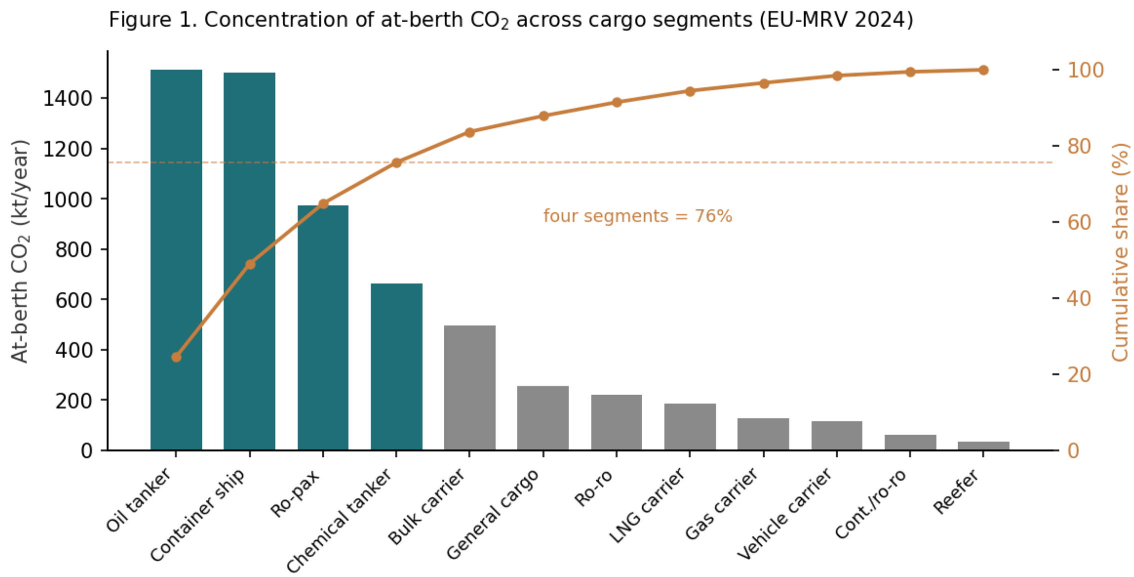


*Figure 1. Concentration of at-berth $CO_2$ across the twelve cargo segments (own analysis of EU-MRV 2024). Bars show annual at-berth $CO_2$ per segment; the line shows the cumulative share. The four leading segments account for 76% of the pool.*

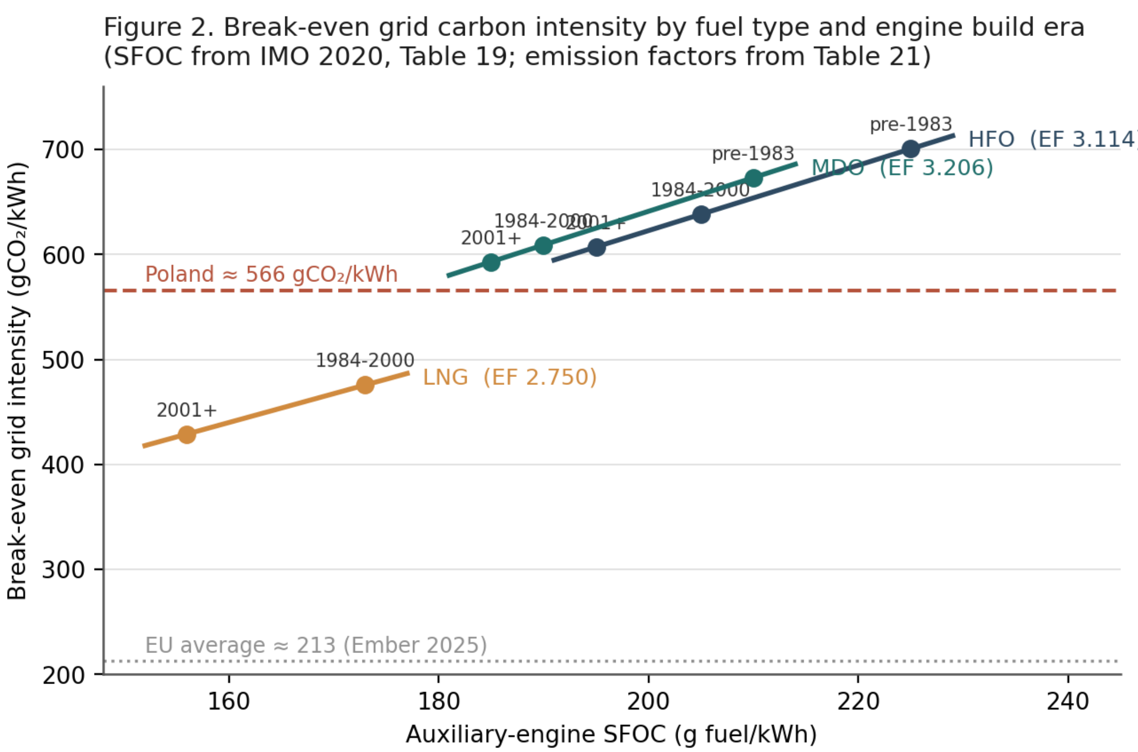

Figure 2. Break-even grid carbon intensity by fuel type and engine build era. Each line is one fuel, its slope given by that fuel's emission factor; the marked points are the *SFOC* values the Fourth IMO GHG Study reports for auxiliary engines (IMO, 2020, Table 19, p. 71, and Table 21, p. 74). The Polish grid lies above five of the eight combinations and below three.

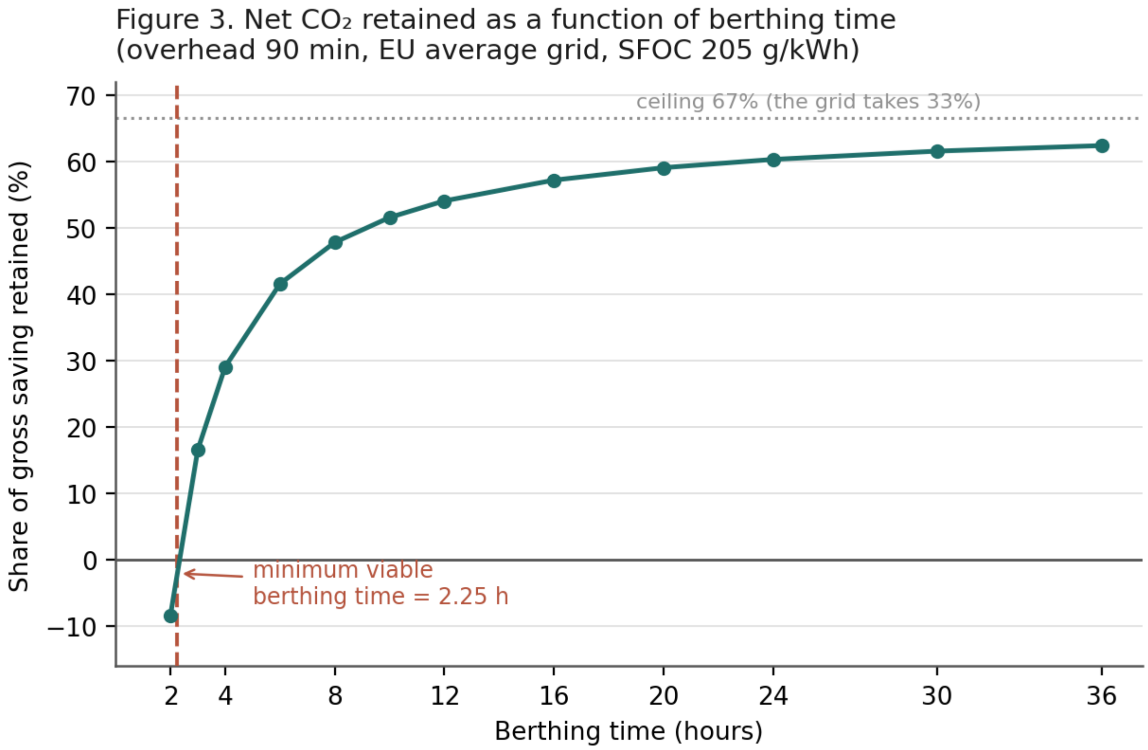


Figure 3. Share of the gross saving retained as a function of berthing time, at 90 minutes of overhead, the EU average grid of 213 $gCO_2$/kWh and the illustrative *SFOC* of 205 g/kWh. Below 2.25 hours the retained share is negative: shore power increases emissions.

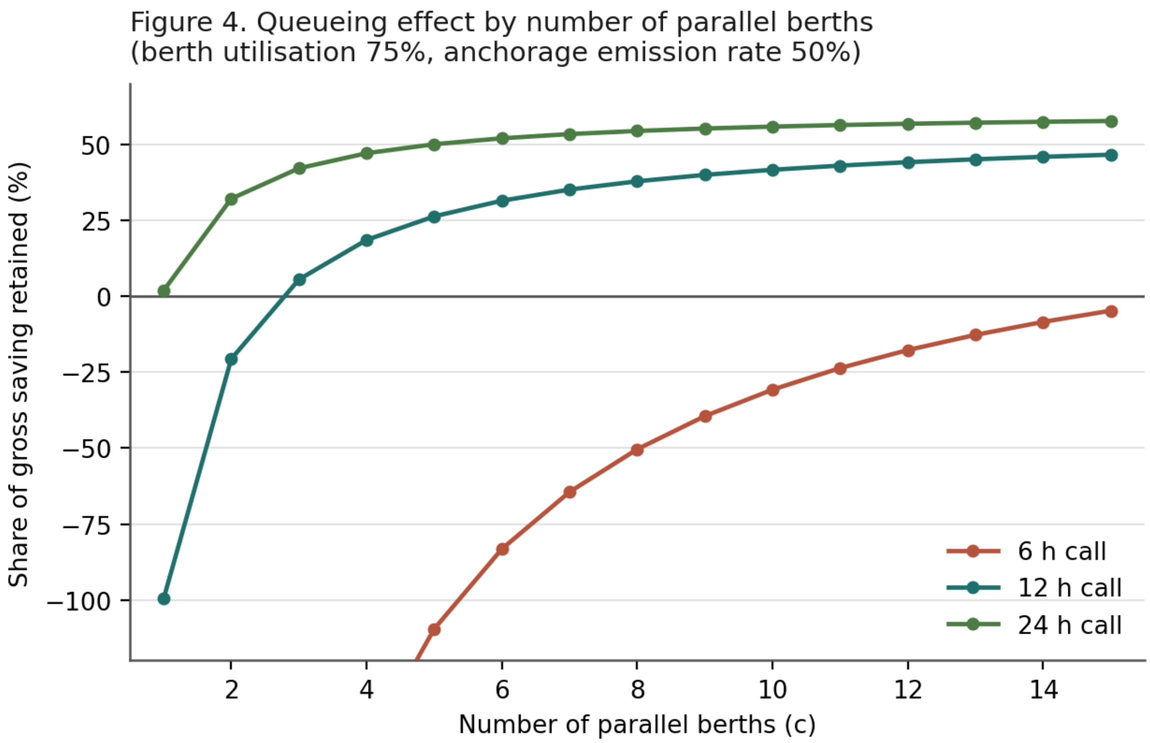


*Figure 4. Queueing effect by number of parallel berths, at 75% berth utilisation and an anchorage emission rate of 50%, for three call durations. The single-berth case is the worst case; parallel berths dampen the penalty sharply.*

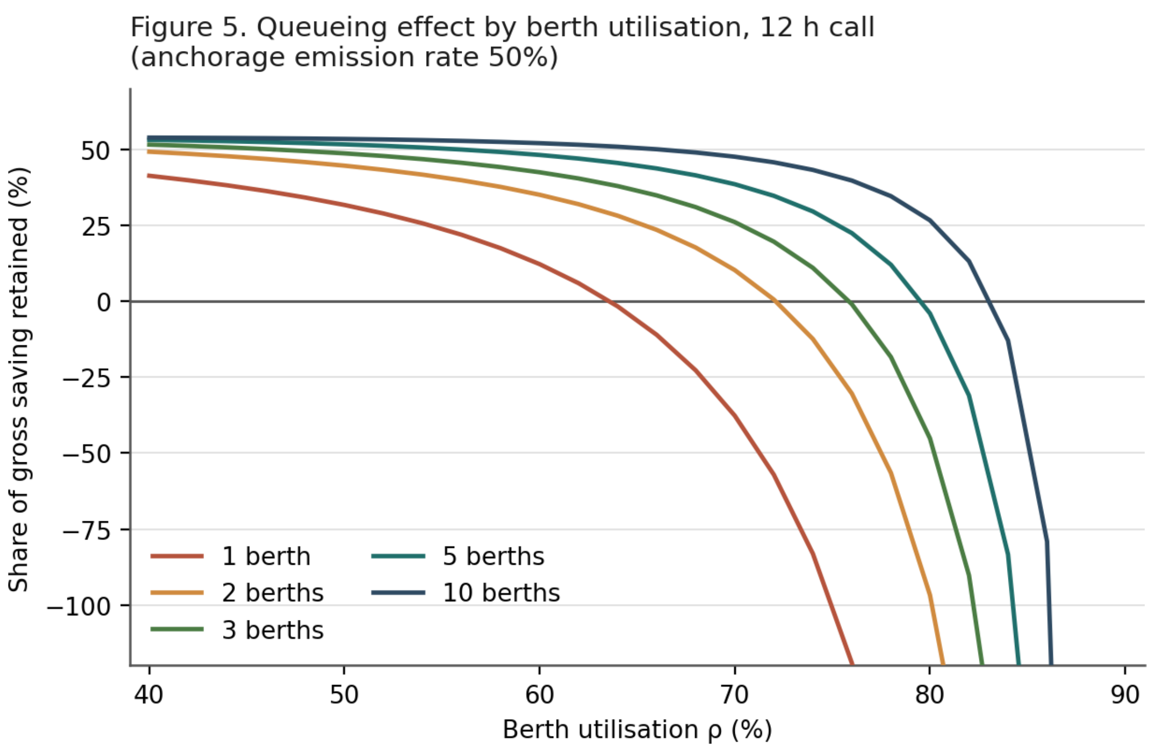

Figure 5. Retained share of the gross saving at a twelve-hour call, as a function of berth utilisation, for one to ten parallel berths, with the anchorage emission rate held at 50%. More berths raise the curve; higher utilisation lowers it. The single-berth case turns negative early, while five berths stay positive until just below 80% utilisation and ten until roughly 83%.

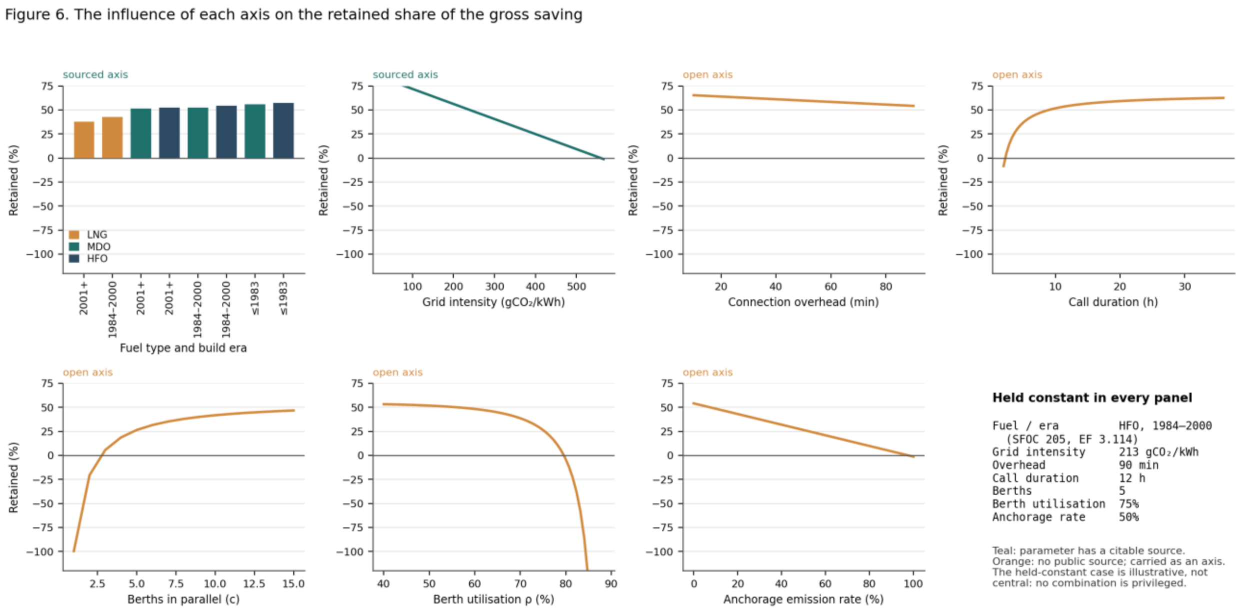


*Figure 6. The influence of each axis on the retained share of the gross saving. Each panel sweeps one parameter across its plausible range while the others are held at the illustrative case shown, which is illustrative and not central: no combination of fuel and build era is privileged. Teal marks parameters with a citable source; orange marks parameters carried as open axes because no public source fixes them.*

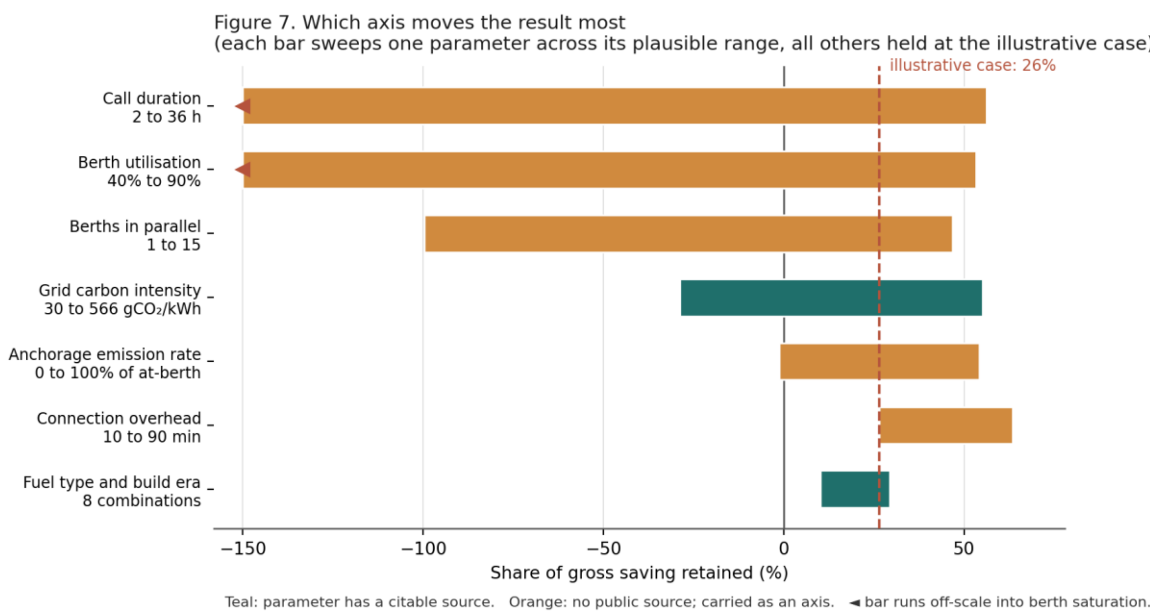

*Figure 7. Amplitude of each axis. Each bar spans the retained share obtained by sweeping one parameter across its plausible range with all others held at the illustrative case, and the bars are ordered by amplitude. Arrows mark bars that run off-scale into berth saturation. Fuel type and build era, the only axis drawn entirely from a citable source, has the smallest influence of the seven.*

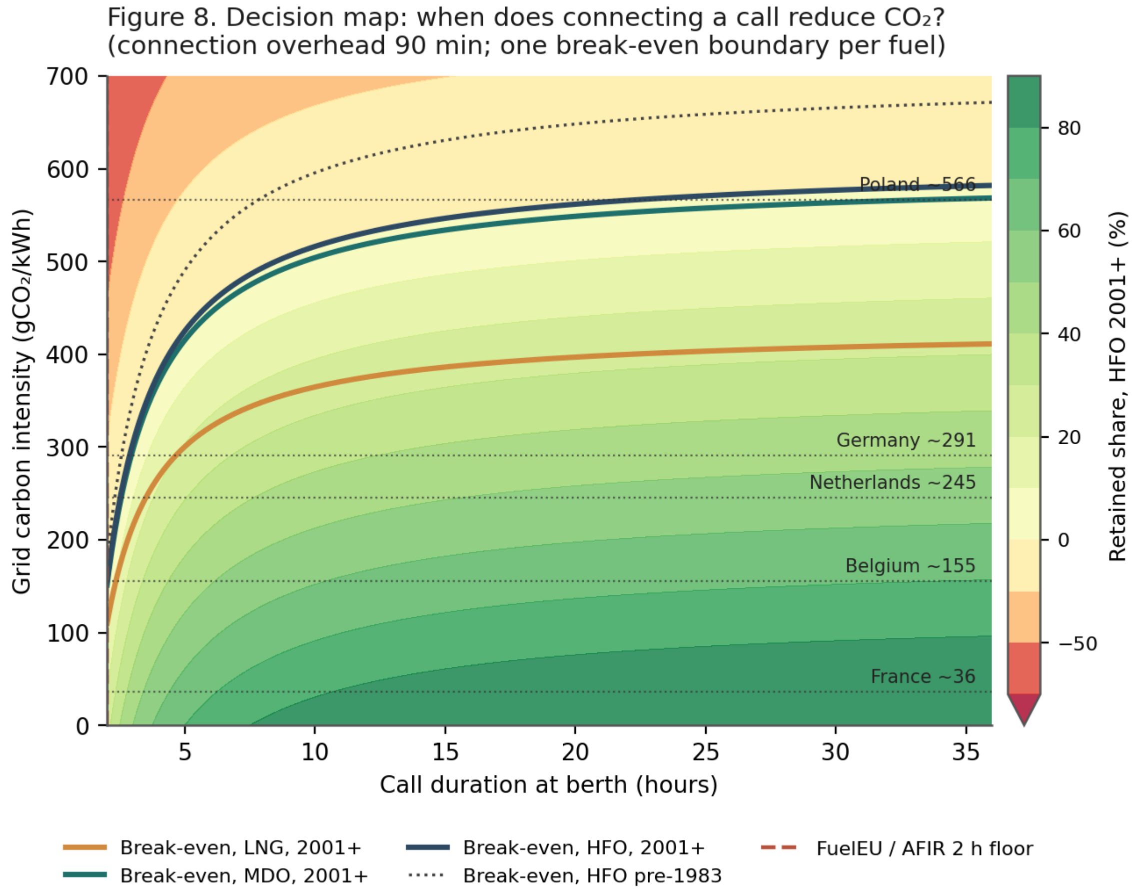


Figure 8. Decision map for a port call, synthesising Thresholds 1 and 2. The retained share is shown over call duration and grid carbon intensity at 90 minutes of connection overhead, shaded for an oil-fuelled auxiliary built after 2000. Each coloured contour is the break-even boundary for one fuel; the dotted contour is the most favourable case in the fleet. The map synthesises results presented in Section 4 and introduces no new data.

Figure 9. How ninety minutes of connection overhead becomes hours of fleet-wide waiting

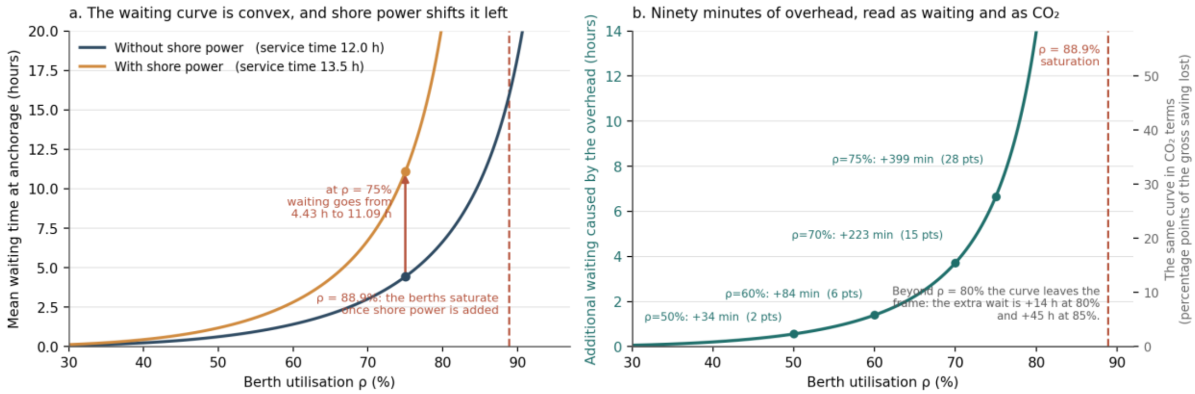


Figure 9. How ninety minutes of connection overhead becomes hours of fleet-wide waiting (own computation, five parallel berths, twelve-hour call, anchorage emission rate 50%). Panel a plots the mean waiting time at anchorage from the Erlang C delay formula (10) against berth utilisation, first for a twelve-hour service time and then for the thirteen-and-a-half-hour service time that shore power produces. The curve is convex, so the same ninety minutes of overhead costs more waiting the busier the terminal is, and the berths saturate at 88.9% utilisation once the overhead is added, against 100% without it. Panel b plots the difference between the two curves, read on the left as hours of extra waiting per vessel and on the right as the percentage points of the gross saving that the waiting consumes. The two scales are the same curve: the right-hand one is the left-hand one multiplied by the anchorage emission rate and divided by the call duration. At 50% utilisation the overhead costs 34 minutes of waiting and 2 points of the saving; at 75% it costs 6.7 hours and 28 points; at 85% it costs 45 hours and the saving is gone. Values are from the Queueing Effect sheet of the workbook.

# Appendix B. Statement on the use of generative AI

This appendix is provided in accordance with the LSM guidelines for a responsible use of artificial intelligence in a Master's thesis. It states the tool used, the stages at which it was used, the nature of the prompts given, and the verification applied to its output.

## Tool and period of use

Claude (Anthropic) was used between the start of the project and its submission. No other generative system was used. No confidential, sensitive or proprietary content was submitted to it: every input was either public data, published literature, or text written for this thesis.

## Stages at which it was used

The assistant was used at four stages. In building the model, it was asked to express the net balance in closed form, to derive the minimum viable call duration from it, and to implement the Erlang C delay formula in a spreadsheet, Excel having no built-in function for it. In sourcing, it was asked to locate primary references for parameters that were not yet properly referenced, and to check whether the figures attributed to those references actually appear in them. In producing the figures, it was asked to generate the charts from the workbook values. In writing, it was asked to draft sections and to reorganise the structure of the discussion, on the basis of results and arguments that had already been established.

It was not used to formulate the research question, to select the research design, or to decide which conclusions the results support. Those decisions were taken before the assistant was involved, or against its suggestions.

## Nature of the prompts

Prompts were of four kinds. Derivational prompts asked for an algebraic step, for example to solve the retained-share expression for the call duration at which it reaches zero. Implementation prompts asked for a formula or a chart, for example to rebuild the Erlang C probability from the Poisson functions available in Excel. Verification prompts asked whether a stated figure appears in a named source, for example whether the Fourth IMO GHG Study reports a specific *SFOC* for auxiliary engines. Drafting prompts asked for a section to be written from a set of results and a stated argument.

## Verification of the output, and errors identified

The assistant produced errors, and identifying them formed a substantial part of the work. Five are recorded here because they illustrate the level of scrutiny applied.

It attributed the figure for shipping's share of global $CO_2$ emissions to a source that does not contain it; the figure was replaced with the value reported in the Fourth IMO GHG Study. It proposed a value for EU average grid carbon intensity citing Ember, but the published figure differs; checking the report directly gave the value now used. It proposed the Californian At Berth Regulation as the source of an upper bound on connection overhead, when that regulation states a compliance ceiling rather than an observed duration; the claim was downgraded and the source replaced. It carried heavy fuel oil as the fuel type without declaring it, so that a parameter labelled as sourced in fact contained an undeclared assumption; fuel type is now an axis. And it proposed a connection overhead range attributed to a reference that did not exist in the bibliography; the range was eventually located in Zis (2019, p. 85) after obtaining the full text.

In one instance the correction ran the other way. A reformulation of the net saving that made the threshold explicit was proposed by the author, found on checking to be dimensionally incorrect, and adopted in corrected form; that expression now appears in the model and in Section 3.3. Every numerical result in this thesis was recomputed independently of the assistant's output and reconciled with the workbook before being reported.

## Statement

The analysis, the choice of question and method, the judgements about what the results support, and the responsibility for any remaining error are the author's own. The assistant was used as a tool, and its output was treated as unreliable until verified.

# Appendix C. The queueing computation, step by step

Excel provides no Erlang C function, so the probability of waiting is rebuilt from the Poisson distribution: the mass at c servers divided by one minus the utilisation, over the cumulative mass to c minus one plus that same quantity. The three tables in Section 4.3 inline these steps in a single formula, which makes them compact but not auditable. Table C1 therefore sets the computation out one step at a time for the reference case, so that each intermediate quantity can be checked independently.

*Table C1. The queueing computation for the reference case: a twelve-hour call, five parallel berths, 75% berth utilisation, a connection overhead of 90 minutes and an anchorage emission rate of 50% of the at-berth rate. Waiting times are in hours per vessel.*

| Step | Value | Note |
|---|---|---|
| Call duration d (h) | 12 | |
| Service time without shore power S0 (h) | 12 | |
| Service time with shore power S1 (h) | 13.5 | d + overhead |
| Utilisation without rho0 | 0.75 | |
| Utilisation with rho1 | 0.844 | rho x S1/S0 |
| Offered load a0 | 3.75 | rho0 x c |
| Offered load a1 | 4.219 | rho1 x c |
| Erlang C without | 0.462 | probability of waiting |
| Erlang C with | 0.641 | |
| *Wq* without (h) | 4.433 | |
| *Wq* with (h) | 11.085 | |
| Extra waiting per vessel (h) | 6.652 | delta W |
| Retained share | 0.264 | |

The retained share on the last line is the entry that appears at twelve hours and five berths in Table 6, at twelve hours and 75% utilisation in Table 7, and at twelve hours and 50% in Table 8, which is the consistency check between the three sweeps. The single berth case is simpler still: at c equal to one the Erlang C probability collapses to the utilisation itself and the waiting time reduces to the standard M/M/1 result, the utilisation times the service time over one minus the utilisation, which can be evaluated by hand.

**Summary**

International shipping accounts for about 2.9% of global $CO_2$ emissions from human activity, and much of the pollution that goes with it is concentrated in ports, where berthed ships run auxiliary engines to power their systems. Onshore power supply lets a berthed ship switch those engines off and draw electricity from the quay instead. From 2030, FuelEU Maritime requires containerships and passenger ships above 5 000 gross tonnage to connect in the main European ports.

Connecting a ship does not eliminate the emissions it would have produced. It relocates them. This thesis asks under what circumstances that relocation reduces net $CO_2$ for the European cargo fleet, and when it does not. The model is deterministic and anchored on the verified at-berth emissions of 10 601 cargo vessels reported under EU-MRV for 2024. Every parameter is either measured, taken from a citable source, or swept across a published range; none is fixed at a convenient value.

Three conditions must hold at once. The shore electricity must be cleaner per unit of energy than the fuel it replaces, which places the break-even between 429 and 701 $gCO_2/kWh$ depending on the fuel and build era of the auxiliary engine. The call must be long enough to absorb the time spent connecting and disconnecting, which sets a minimum viable duration of 2.25 hours on the European average grid. And the berth must not be so heavily used that this added time pushes vessels into anchorage waiting: beyond roughly 80% berth utilisation the balance turns negative, and mandatory connection at a saturated terminal raises fleet emissions instead of lowering them.

The three conditions tighten together rather than trading off, so shore power is most effective on the vessels that have done least to decarbonise their own generation, calling at length in uncongested berths on clean grids. The two-hour exemption boundary in the regulation is correctly placed at exactly one value of the connection overhead, 80 minutes, and no public source measures that quantity.

**UNIVERSITÉ CATHOLIQUE DE LOUVAIN**
**Louvain School of Management**
Place des Doyens, 1 bte L2.01.01, 1348 Louvain-la-Neuve
Boulevard Emile Devreux 6, 6000 Charleroi, Belgique
Chaussée de Binche 151, 7000 Mons, Belgique
**www.uclouvain.be/lsm**