MARKET INSIGHTS
Nov 2026

Capesize Voyage Speed Optimization:

A Reinforcement Learning Model for Fuel, Earnings, and CII



VESON NAUTICAL

BAYES BUSINESS SCHOOL
CITY ST GEORGE'S UNIVERSITY OF LONDON

Introduction

Voyage planning has grown considerably more complex in recent decades. Operators must simultaneously balance fluctuating bunker fuel prices, dynamic freight markets, port congestion, charter party obligations, and increasingly stringent emission regulations. Voyage optimisation determines both commercial viability and now environmental compliance. For operators of Capesize Bulkers, the stakes are high. The combination of long voyage distances and high bunker consumption means that even marginal speed optimisation can yield substantial economic gains and emissions reductions.

This analysis presents the results of collaborative research between Veson Nautical and Bayes Business School, in which we developed and evaluated a deep reinforcement learning model for Capesize Bulker speed optimisation. The research draws on 130,388 voyages that took place between July 2020 and June 2026, taken from VesselsValue's Trade AIS tracking. The model is trained on the voyages that took place before April 2025 and tested on voyages that followed, comparing its speed recommendations with the actual speeds sailed as observed in the data.

The model improved on operator decisions for fuel, earnings, and carbon intensity. On the 18,443 voyages used for testing, the model cut fuel consumption by an average of 47.2 metric tonnes per voyage (9.98%) and improved economic performance by $42,156 per voyage (22.6% improvement). Additionally, the speeds recommended by the model increased the share of annual emissions ratings that would score "A" under the IMO's Carbon Intensity Indicator (CII) from 76.1% to 90.5%.

This paper sets out how the model works, what we found, and what the results could mean for owners, charterers, and regulators.

The model

The model worked in two stages. It first calculated a starting speed for each voyage, and a reinforcement learning agent then adjusted that speed using a wider set of information (Figure 1).

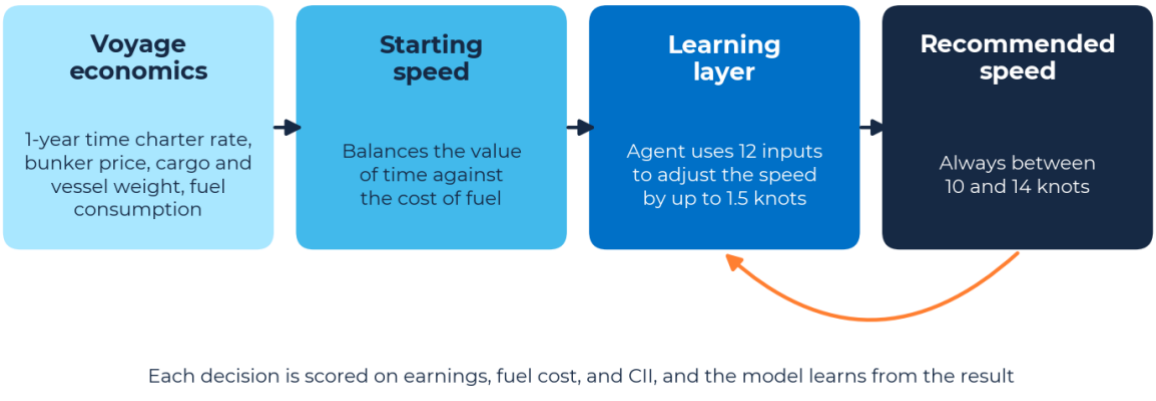


Figure 1: How the model reaches a speed recommendation

The starting speed was set at the point where sailing faster stopped being worth the extra fuel. Sailing faster finishes a voyage sooner and frees the vessel to earn hire elsewhere, but daily fuel consumption rises with the cube of speed, so each extra knot costs more than the

last. The calculation weighed the annual time charter rate against the fuel cost of extra speed, taken from the bunker price at departure and an estimate of how much fuel that vessel burns. Fuel burn also rises with the weight being moved, so the calculation included the cargo on board and the vessel's own weight, estimated at 25% of its deadweight tonnage.

Each vessel's fuel estimate came from its own historical consumption records. Bunker prices came from daily prices for IFO380 fuel oil at six major hubs (Singapore, Algeciras, Busan, Fujairah, Houston, and Rotterdam), matched to the hub nearest to the departure port. A higher charter rate raised the starting speed, while a higher bunker price or a heavier cargo lowered it.

While the first stage applied the same calculation to every voyage, the second stage learned from experience. A single reinforcement learning agent was trained across the fleet. Each training episode corresponded to a vessel's year of voyages, and episodes were drawn in random order, resulting in the agent working through over two million voyage decisions, seeing a wide range of vessels, years and market conditions. The agent chose a speed for every voyage, measured the outcome, and gradually settled on the choices that produced the best results. The agent planned over a full year rather than a single voyage. A vessel-year averaged 11.9 voyages, so the model could see that a vessel sailing faster finishes sooner and fits more fixtures into the year.

Before each voyage, the agent looked at 12 inputs: the starting speed, loading condition, cargo as a share of deadweight tonnage, the C5TC rate, the bunker price, expected voyage length, how sharply fuel cost would rise with extra speed, the vessel's CII ratio so far that year and how far the CII ratio fell below the required level, remaining days in the year, voyages already completed, and revenue earned so far. Because the starting speed was one of these inputs, the agent could learn when the calculation tended to be too fast or too slow, and by how much.

The agent then moved the starting speed up or down by as much as 1.5 knots, and the final speed was kept between 10 and 14 knots. The model used that speed to calculate the voyage's fuel consumption, fuel cost, and earnings, and added the fuel burned and distance sailed to the vessel's total CII for the year.

Every decision was scored against time-charter earnings, fuel cost, and CII performance. In the main version of the model, which we called the balanced setting, these carried weights of 0.5, 0.3, and 0.2, respectively. A large penalty applied if the vessel's projected CII rating fell to D, and a penalty twice as large if it fell to E. The agent was trained using Proximal Policy Optimisation (PPO), a method that changes the agent's behaviour in small steps so that learning stays stable.

We trained the model on voyages from 2020 to April 2025 and tested it on 18,443 voyages from April 2025 onwards. Fuel use and earnings were calculated the same way for the model and for operators, so any difference in results was due to speed.

Our findings

On average, the model recommended speeds 0.24 knots below what operators sailed. That small reduction improved every headline measure, as seen in Table 1.

Table 1: Model compared with actual operator decisions (balanced setting)

Measure	Actual operators	Model	Change
Mean speed (knots)	11.26	11.02	−0.24
Mean fuel use per voyage (metric tonnes)	472.5	425.4	−47.2 (−9.98%)
Mean economic performance per voyage	$186,776	$228,932	+$42,156
Mean CII ratio (attained/required)	0.77	0.70	−0.07
Vessel-years rated A	76.1%	90.5%	+14.4 points

Source: Veson Nautical and Bayes Business School research.

The improvement was statistically significant ($p<0.001$), albeit modest. That was expected, given the comparison was against experienced operators. The gain was consistent, though, and at $42,156 per voyage, the cumulative effect across a fleet could be considerable.

Laden voyages accounted for most of the gain. Operators already sailed slower when laden, at 10.79 knots against 11.66 knots in ballast, but the model went further. It recommended a lower laden speed on 78.5% of voyages, by 0.73 knots on average. Fuel consumption depends on displacement as well as speed, so each knot costs more with cargo on board. On laden voyages, the model saved 13.05% of fuel and $81,315 per voyage.

In ballast, the model often pushed the other way. It suggested faster speeds than operators on 51.8% of ballast voyages, taking greater advantage of the lower fuel to bring the next fixture forward sooner. Operators may have had reasons the model could not see, such as port schedules, canal slots, or arrival windows. Even so, the model improved ballast results by $9,086 per voyage (Table 2).

Table 2: Results by loading condition

Loading condition	Operator speed (knots)	Model speed (knots)	Fuel saving	Economic improvement per voyage
Laden	10.79	10.06	13.05%	$81,315
Ballast	11.66	11.83	2.48%	$9,086

Source: Veson Nautical and Bayes Business School research.

The model responded far more strongly to bunker prices than operators did. Operator speeds barely moved as prices changed, with a correlation of −0.045, while the model's speeds were about six times more sensitive, at −0.26. Above $656 per metric tonne, the model averaged 10.63 knots against 11.14 knots for operators (Figure 2). One possible explanation is that operators were burning fuel bought weeks or months earlier, while the model used the price on the departure date. At those prices, the economic improvement reached $70,101 per voyage.

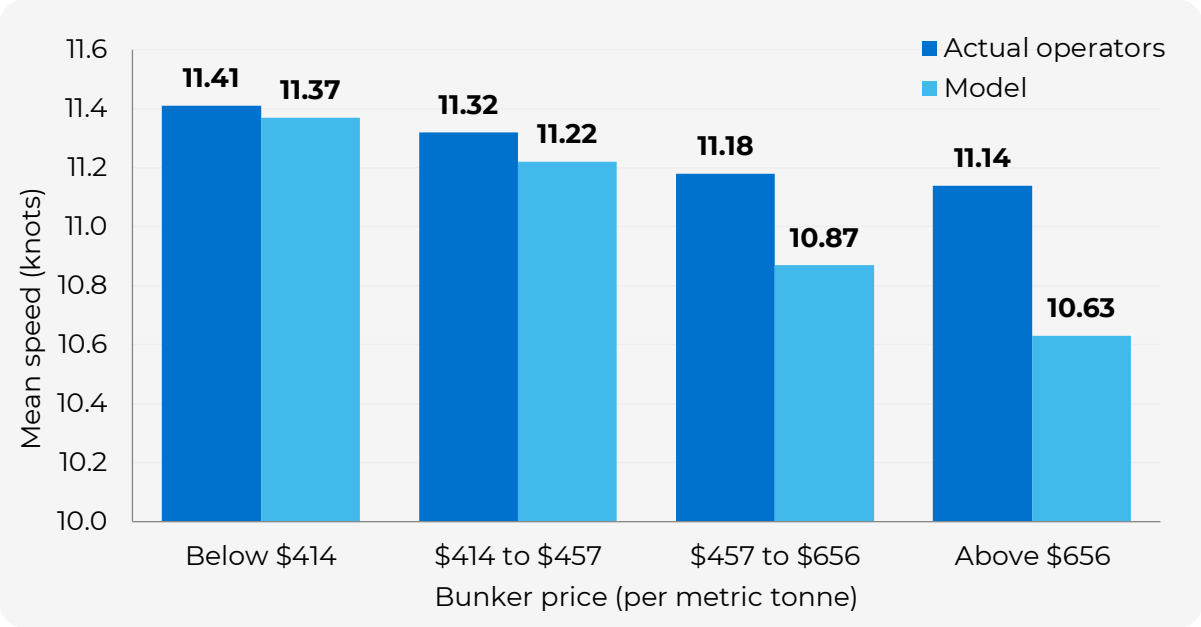


Figure 2: Mean speed by bunker price, model compared with actual operators. Source: Veson Nautical and Bayes Business School research.

Slower laden speeds also led to better CII ratings. We compared the projected CII rating of each vessel-year at the model's speeds with its rating at the speeds operators actually sailed. The share rated A rose from 76.1% to 90.5%, and D and E ratings, which would require a corrective action plan, fell from 50 vessel-years to five (Table 3).

Table 3: Projected annual CII ratings, 2,179 vessel-years

CII rating	Actual operators	Model
A	1,658 (76.1%)	1,973 (90.5%)
B	304 (14.0%)	148 (6.8%)
C	167 (7.7%)	53 (2.4%)
D	36 (1.7%)	5 (0.2%)
E	14 (0.6%)	0 (0.0%)

Source: Veson Nautical and Bayes Business School research.

The results held across every weighting we tested. We retrained the model with 11 different weightings, ranging from revenue only to CII only, and all 11 improved compared to operators, with economic gains ranging from $31,774 to $45,396 per voyage (Table A1).

Scaled up to every vessel in the dataset, the results were substantial. Applying the balanced model to all 1,687 vessels at 12 voyages a year would have saved about 950,000 metric tonnes of fuel and 2.97 million tonnes of CO_2, worth around $853 million a year. This estimate assumes full adoption, 2025 to 2026 market conditions, and no weather or congestion delays, so it is an illustration rather than a forecast.

Implications

All these savings come from speed decisions alone, with no new equipment, fuels, or trades. For owners and operators, the model could work as decision support at the point of departure, using data that a chartering desk already holds. It could also serve as a baseline

for tracking realised savings over time. The most immediate lesson concerns laden voyages, where the results suggest the extra fuel cost of speed is often underestimated.

For regulators, the model offers a way to test proposed CII changes before they take effect. Tightening the annual reduction factor and retraining the model shows how a fleet is likely to respond, and results can be split by route, since a Brazil to China iron ore voyage has different economics from a shorter one from Australia. In our results, speed changes alone raised the share of A-rated vessel-years from 76.1% to 90.5%. That suggests that the CII reference lines in force during the test period can be met through operations alone, although the annual reduction factor continues to tighten to 2030. The weighting tests also showed which signal does the work. Removing the CII term from the model's objective raised its mean speed by 0.19 knots, more than removing the fuel cost term, and the hard penalty for a D or E rating cost only $952 per voyage, about 2% of the gain. The CII framework restrained speed more than the fuel price did, even in a period of very high bunker prices.

The model can also be set up for the stakeholder using it. The earnings term uses the C5TC index, so an operator trading on spot rates could substitute their own earnings series and retrain. The weights can then be set to match the charter structure and who bears which cost. Under a time charter, the charterer pays for bunkers while the owner earns hire rate, so the two would weigh the fuel term differently on the same voyage.

Conclusion

The research found that a reinforcement learning model, built on a calculated starting speed, improved on the real speed decisions Capesize Bulker operators made. Across 18,443 test voyages, fuel use fell by 9.98%, economic performance improved by $42,156 per voyage, and the share of A-rated vessel-years rose from 76.1% to 90.5%.

The findings have limitations. The test period overlapped the 2026 Hormuz crisis, when bunker prices were unusually high. This likely raised the measured economic gain, although the fuel saving held between 8.8% and 11.2% in every bunker price quartile. Earnings were based on the C5TC index, not actual fixture rates, and the data did not include weather, port congestion, or hull condition. Each vessel's fuel consumption estimate came from the same data used to test the model, and the ten-knot floor ruled out slower speeds.

Within those limits, the findings were consistent. Most of the value came from laden voyages, where operators already slowed down but, according to the model, not by enough. The model also recommended faster speeds on over a third of voyages, often in ballast, where an earlier arrival was worth more than the extra fuel. Planning across a full year was essential for the model to assign a proper value to time. Across every weighting tested, the model's speeds used less fuel and earned more than operators did at the speeds they actually sailed.

Appendix

Table A1 shows results for all 11 weightings, with weights listed in the order revenue, fuel cost, and CII.

Table A1: Results by weighting

Weighting	Weights	Mean speed (knots)	Fuel saving	Economic improvement per voyage
Balanced	0.5 / 0.3 / 0.2	11.02	9.98%	$42,156
Revenue-weighted	0.7 / 0.2 / 0.1	11.02	10.33%	$44,521
Fuel-weighted	0.3 / 0.5 / 0.2	10.96	10.50%	$45,396
CII-weighted	0.3 / 0.3 / 0.4	11.03	9.65%	$39,392
Balanced, no CII penalty	0.5 / 0.3 / 0.2	11.06	10.15%	$43,108
No revenue	0 / 0.5 / 0.5	11.05	10.20%	$43,562
No fuel	0.5 / 0 / 0.5	10.99	10.27%	$44,100
No CII	0.5 / 0.5 / 0	11.21	9.62%	$39,989
Revenue only	1 / 0 / 0	11.36	8.16%	$31,774
Fuel only	0 / 1 / 0	11.19	9.14%	$39,019
CII only	0 / 0 / 1	11.03	9.85%	$42,574

Source: Veson Nautical and Bayes Business School research.

Glossary

AER: Annual Efficiency Ratio. The carbon intensity measure used for Bulkers under CII. It compares a vessel's CO_2 emissions with its deadweight tonnage and the distance it sails in a year.

Ballast voyage: A voyage without cargo, usually to reposition a vessel for its next loading port.

C5TC: Capesize 5 Time Charter Average. A benchmark index of daily earnings for Capesize Bulkers, published by the Baltic Exchange.

Capesize Bulker: A large dry bulk vessel, typically of 100,000 DWT and above, used mainly to carry iron ore and coal. Capesize Bulkers are too large to pass through the Panama Canal, and the name refers to routes around the Cape of Good Hope and Cape Horn. The vessels in this study range from 120,300 to 216,700 DWT.

CII: Carbon Intensity Indicator. An annual A to E rating introduced by the International Maritime Organization in 2023. Vessels rated D for three consecutive years, or E in a single year, must submit a corrective action plan.

DWT: Deadweight tonnage. The total weight a vessel can carry, including cargo, fuel, and stores.

Heuristic: A rule-based calculation that gives a reasonable answer quickly but does not adapt. In this research, the heuristic is the formula used to calculate the starting speed.

IFO380: Intermediate Fuel Oil 380. A heavy fuel oil widely used as a benchmark for marine bunker prices.

Laden voyage: A voyage carrying cargo.

PPO: Proximal Policy Optimisation. A reinforcement learning method that updates its decision rules in small, controlled steps so that learning stays stable.

Reinforcement learning: A type of machine learning in which an agent learns by making decisions, observing the results, and receiving a reward or penalty.

Slow steaming: Deliberately sailing below a vessel's design speed to reduce fuel consumption and emissions.

Acknowledgements

This white paper presents the findings of recent Bayes Business School graduate Michael Kasinis, whose research project *A Hybrid Heuristic-Deep Reinforcement Learning Framework for Capesize Bulk Carrier Speed Optimisation* (supervised by Onassis Associate Professor Ioannis Moutzouris) was completed under the Veson Nautical Research Support Programme, in collaboration with Veson Nautical analyst Peter Delf and the wider analyst team. This programme, a partnership between Veson Nautical and Bayes Business School, provides students of the MSc in Shipping, Trade and Finance and the MSc in Energy, Trade and Finance courses with direct exposure to real-world industry projects, extensive datasets, and support from Veson's analyst team. Michael's work has been shortlisted for the Maritime UK Master's Competition in November 2026.

This work was supported by the UK Department for Transport, as part of the UK Shipping Office for Reducing Emissions (UK SHORE) Programme and the UK Engineering and Physical Sciences Research Council (EPSRC), United Kingdom [grant number EP/Y024605/1].

About the Authors



Michael Kasinis,
Assistant Commercial Manager, Columbia Shipmanagement

Michael is an Assistant Commercial Manager at Columbia Shipmanagement. He holds a BSc degree in Mathematics, Operational Research and Statistics from Cardiff University, which included an industrial placement in the Risk Division of Lloyds Banking Group. He also holds an MSc degree in Shipping, Trade and Finance from Bayes Business School in London.

His MSc dissertation, completed with Veson Nautical data, has been selected to be a finalist in the 2026 Maritime Masters programme.



Peter Delf,
Valuation Analyst, Veson Nautical

Peter is a Valuation Analyst within the Valuation & Analytics team at Veson Nautical specialising in the Bulker and Offshore sectors. The valuation analyst team is responsible for monitoring and maintaining the valuations created by the VesselsValue algorithm to ensure they are accurate, reliable and up to date. The team also provides consultancy report writing, bespoke value creation, presentation and market commentary, feature premium analysis, energy efficiency calculations and valuation accuracy.

Peter joined Veson Nautical in September 2024 and holds a BA degree in History from Durham, University.

VESON NAUTICAL

BAYES BUSINESS SCHOOL
CITY ST GEORGE'S UNIVERSITY OF LONDON

 | © 2026 Veson Nautical. All rights reserved.