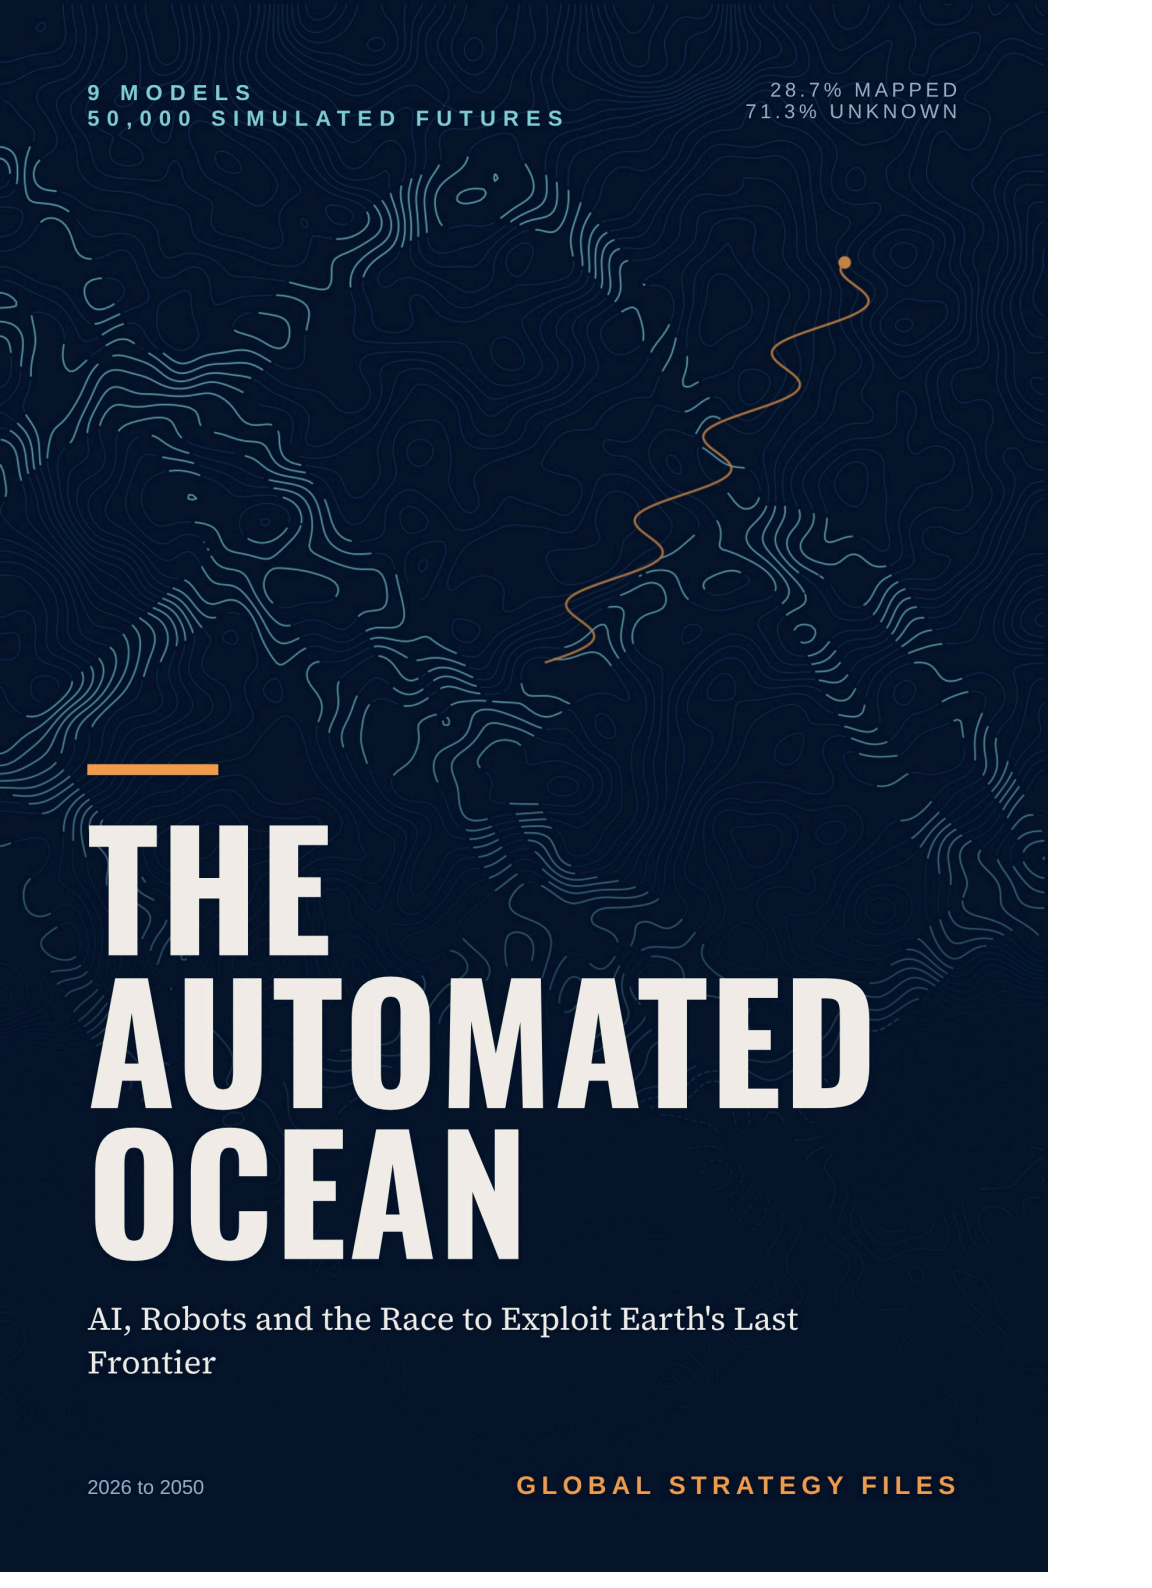
9 MODELS
50,000 SIMULATED FUTURES

28.7% MAPPED
71.3% UNKNOWN



THE AUTOMATED OCEAN

AI, Robots and the Race to Exploit Earth's Last Frontier

2026 to 2050

GLOBAL STRATEGY FILES

THE AUTOMATED OCEAN

AI, Robots and the Race to Exploit Earth's Last Frontier, 2026 to 2050

GLOBAL STRATEGY FILES

globalstrategyfiles.com

Copyright © 2026 Global Strategy Files. All rights reserved. No part of this publication may be reproduced, distributed or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in reviews and certain other noncommercial uses permitted by copyright law.

This book is a collective work of the contributors of Global Strategy Files. The Ocean Frontier Model and all projections are original work of Global Strategy Files. Projections are scenario estimates, not forecasts, and do not constitute investment advice.

First edition, November 2026.

globalstrategyfiles.com

We have better maps of Mars than of the floor of our own ocean. That is about to change. What we do with the knowledge is the question of the century.

Contents

Preface

Executive Summary

The Untapped Ocean: An Atlas in Nine Plates

Part I. The Last Frontier

Chapter 1. Seventy-One Percent Unknown

Chapter 2. The Machine Turn

Chapter 3. The Ocean Economy in 2026

Part II. Seeing the Ocean

Chapter 4. The Mapping Clock

Chapter 5. Robots of the Deep

Chapter 6. The Digital Twin and the Thinking Ocean

Chapter 7. The Ocean Data Economy

Part III. Feeding Ten Billion

Chapter 8. Fishing Under Algorithmic Watch

Chapter 9. Robotic Fish Farms Offshore

Chapter 10. Blue Food 2050

Part IV. Powering and Watering the Planet

Chapter 11. The Robotic Wind Dividend

Chapter 12. Waves, Tides, Heat and Carbon

Chapter 13. Drinking the Sea

Part V. The Deep

Chapter 14. The Nodule Rush

Chapter 15. Can Deep-Sea Mining Pay?
Chapter 16. The Pharmacy of the Abyss

Part VI. Living on and under the Sea

Chapter 17. Habitats and Floating Cities
Chapter 18. Ships Without Sailors

Part VII. Power and Conflict at Sea

Chapter 19. The Wires of the World
Chapter 20. Drone Navies and the Transparent Ocean
Chapter 21. The Polar Frontier
Chapter 22. The Blue Power Index
Chapter 23. Winners and Losers: Coastal Communities in the Automated Ocean

Part VIII. Limits and Futures

Chapter 24. The Rebound Trap
Chapter 25. Rules for a Machine Ocean
Chapter 26. Three Oceans for 2040
Chapter 27. What to Do Now

Annexes

Annex A. Methodology of the Ocean Frontier Model
Annex B. Principal Data Sources by Module
Annex C. Country Input Data for the Blue Power Index
Annex D. Country Profiles of the Blue Power Index
Annex E. Glossary
Annex F. Bibliography
Annex G. Year-by-Year Model Results
Annex H. A Chronology of the Automated Ocean
Annex I. Questions for Discussion and Further Research

List of Figures

Plate 1. Atlas of the untapped ocean: resources, routes and pressure points

Plate 2. The exploitation ladder: share of estimated potential already used or known

Plate 3. Gross in-situ value of the main metals in Clarion-Clipperton nodules

Plate 4. Estimated resource potential of ocean energy sources compared with world electricity demand

Plate 5. Fan chart of ocean economy gross value added under three scenarios, 2025 to 2050

Plate 6. Median gross value added by sector, 2025 to 2050

Plate 7. Adoption of seventeen ocean technologies and activities, 2026 to 2050

Plate 8. Blue Power Index profiles of the eight leading ocean nations

Plate 9. Decomposition of the growth of ocean value added between 2025 and 2050

Figure 3.1. Approximate gross value added of ocean economy sectors, 2025

Figure 4.1. The Mapping Clock: share of the seafloor mapped under four scenarios, 2017 to 2090

Figure 4.2. Sensitivity of the robotic surge scenario to vessel utilisation and swath width

Figure 4.3. Share of the seafloor mapped by ocean basin, end of 2024

Figure 5.1. How many times more expensive is the crewed method? 2026 and 2040

Figure 5.2. Cost paths of crewed and robotic methods for four categories of ocean work, 2010 to 2050

Figure 7.1. The ocean data value chain

Figure 8.1. Fish biomass and catch under four scenarios, 2026 to 2060

Figure 8.2. Biomass in 2050 as a function of technological creep and enforcement coverage

Figure 8.3. Catch in 2050 as a function of enforcement coverage, with technological creep of 4 percent a year

Figure 9.1. Aquaculture share of aquatic animal production since the 1950s

Figure 10.1. Aquatic animal food supply per person under four scenarios, 2022 to 2050

Figure 10.2. Offshore aquaculture output in 2050 by scenario

Figure 10.3. Main producers of aquatic animals, share of world production, 2022

Figure 10.4. Composition of aquatic animal production in the robotic offshore scenario, 2022 to 2050

Figure 11.1. Distribution of the reduction in the cost of offshore electricity attributable to robotics, 2035

Figure 11.2. Offshore wind capacity by country, end of 2025

Figure 11.3. Global offshore wind capacity, 2025 and GWEC outlook for 2035
Figure 11.4. Sensitivity of the floating wind cost reduction to each parameter
Figure 15.1. Distribution of the private net present value of a 12 Mt per year nodule project
Figure 15.2. Probability of a positive net present value by battery chemistry and regulatory delay
Figure 15.3. Probability of profitability with a strategic price premium
Figure 19.1. Time to restore all cables after a simultaneous regional attack
Figure 19.2. Median days to restore all cables by regional repair fleet size
Figure 22.1. The GSF Blue Power Index 2026: weighted contribution of each pillar
Figure 22.2. Ocean potential versus capability, 30 countries
Figure 22.3. Change in rank between 2026 and the 2040 projection
Figure 26.1. Probability of the three scenarios in 2040
Figure 26.2. Ocean economy gross value added by sector, 2025 and 2050 (median)
Figure 26.3. The scenario space: governance cooperation and climate stress in 3,000 simulated futures

List of Tables

Table 1.1. The four walls around the ocean and how machines breach them

Table 1.2. The state of ocean knowledge in 2026

Table 1.3. The nine modules of the GSF Ocean Frontier Model

Table 1.4. Milestones in measuring the ocean floor

Table 2.1. The main families of ocean robots in 2026

Table 2.2. A ladder of ocean autonomy

Table 2.3. The communications hierarchy of ocean robots

Table 3.1. Approximate structure of the global ocean economy in 2025 (GSF modelling baseline)

Table 3.2. The four flows of the ocean economy, latest data

Table 4.1. Share of the seafloor mapped to modern standards, Seabed 2030 announcements

Table 4.2. Scenarios of the Mapping Clock

Table 4.3. Results of the Mapping Clock

Table 4.4. The cost of finishing the global seafloor map

Table 4.5. Year in which 95 percent of the seafloor is mapped, by robotic productivity assumptions

Table 4.6. Uses of a complete seafloor map

Table 5.1. Indicative size of ocean robotics markets

Table 5.2. Leading ocean robotics ecosystems

Table 5.3. The Ocean Access Cost Index, crewed versus robotic methods

Table 5.4. Robotic cost advantage in 2040 under alternative assumptions (crewed cost divided by robotic cost)

Table 6.1. Building blocks of an ocean digital twin

Table 6.2. Channels of economic value from ocean forecasting and monitoring

Table 7.1. Main producers and buyers of ocean data

Table 8.1. Scenarios and results of the Fishing Under Algorithmic Watch model

Table 8.2. Minimum enforcement coverage required to keep biomass at 80% of B(MSY) in 2050

Table 8.3. Biomass in 2050 relative to B(MSY) under alternative biological assumptions

Table 9.1. Landmark offshore aquaculture systems

Table 9.2. Indicative economics of three farming systems (GSF estimates)

Table 10.1. Scenario assumptions of the Blue Food 2050 model

Table 10.2. Aquatic animal production and food supply, 2022 to 2050

Table 11.1. Results of the Robotic Wind Dividend model, 2035

Table 11.2. Sensitivity of the floating wind result (central estimate 11.2 percent)
Table 13.1. Key indicators of global desalination
Table 14.1. The three main deep-sea mineral resources
Table 14.2. Key parameters of the NORI-D pre-feasibility study (TMC, August 2025)
Table 15.1. Uncertain variables in the deep-sea economics model
Table 15.2. Results of the deep-sea economics model
Table 15.3. Probability of a positive net present value by scenario
Table 15.4. Break-even ecological value per square kilometre of seabed (profitable simulations only)
Table 15.5. Probability of positive NPV and median NPV by discount rate and median nickel price
Table 15.6. Effect of a strategic price premium on project economics
Table 16.1. Indicative estimates of marine biotechnology markets
Table 17.1. A realism scorecard for living at sea
Table 17.2. Reclamation versus floating construction for coastal expansion
Table 19.1. Selected incidents involving undersea cables and pipelines, 2022 to 2025
Table 19.2. Days until the last cable is restored after a simultaneous regional attack
Table 19.3. Median days to restore all cables, by regional fleet size and number of cuts
Table 20.1. Missions of military uncrewed maritime systems
Table 20.2. Asymmetries of the automated ocean at war
Table 21.1. The polar oceans and the automated ocean
Table 22.1. The seven pillars of the Blue Power Index
Table 22.2. The GSF Blue Power Index 2026 and projected 2040 ranking
Table 22.3. Countries most exposed to the frontier gap
Table 22.4. Rank under alternative weighting schemes, selected countries
Table 23.1. Prospects of five groups of ocean workers in the automated ocean
Table 24.1. Six systemic risks of the automated ocean
Table 24.2. A dashboard of ocean health for the automated age
Table 25.1. Key international frameworks for the automated ocean
Table 26.1. Scenario probabilities
Table 26.2. Projected gross value added of the ocean economy in 2050
Table 26.3. Median gross value added by sector, 2025 and 2050, $ bn
Table 26.4. Signposts for the three scenarios, 2026 to 2032
Table 26.5. Scenario probabilities under alternative driver assumptions
Table 27.1. An investor's map of the automated ocean

Table 27.2. A ten-year agenda for the automated ocean
Table A.1. Sector parameters of module M9
Table G.1. Share of the seafloor mapped by scenario, 2026 to 2050, %
Table G.2. Fish biomass relative to B(MSY) by scenario, 2026 to 2060
Table G.3. Catch relative to MSY by scenario, 2026 to 2060
Table G.4. Blue food projections by scenario
Table G.5. Cable repair queue: days until the last cable is restored (10th, 50th and 90th percentiles)
Table G.6. Deep-sea economics: distribution of key outputs (5th, 25th, 50th, 75th, 95th percentiles)

Abbreviations

AI	Artificial intelligence
AUV	Autonomous underwater vehicle
BBNJ	Biodiversity beyond national jurisdiction (High Seas Treaty)
BPI	GSF Blue Power Index
CCZ	Clarion-Clipperton Zone
EEZ	Exclusive economic zone
FAO	Food and Agriculture Organization of the United Nations
GEBCO	General Bathymetric Chart of the Oceans
GVA	Gross value added
GW	Gigawatt
GWEC	Global Wind Energy Council
IDRA	International Desalination and Reuse Association
IMO	International Maritime Organization
ISA	International Seabed Authority
LCOE	Levelised cost of electricity
LFP	Lithium iron phosphate
MASS	Maritime Autonomous Surface Ships
MSY	Maximum sustainable yield
Mt	Million tonnes
NPV	Net present value
O&M	Operations and maintenance
OECD	Organisation for Economic Co-operation and Development
ROV	Remotely operated vehicle
TMC	The Metals Company
UNCLOS	United Nations Convention on the Law of the Sea
USV	Uncrewed surface vessel
UUV	Uncrewed underwater vehicle
XLUUV	Extra-large uncrewed undersea vehicle
ZJ	Zettajoule

Preface

Why this book

Most writing about the future of the ocean falls into one of two moods. One celebrates new technologies and the wonders they reveal; the other warns of collapse. Both are useful and both leave decision makers without the numbers they need. How fast will the seafloor map be finished? Will artificial intelligence empty the sea of fish or help save it? How much food can robotic farms add? Is deep-sea mining profitable at realistic prices? Which countries are equipped to benefit from an ocean run by machines, and which risk watching others exploit their waters? How long would a region stay cut off if its cables were attacked?

This book tries to answer those questions with original models, built for the purpose and documented in full, rather than with impressions. It was written by the contributors of Global Strategy Files in the autumn of 2026, a year in which the High Seas Treaty entered into force, the International Maritime Organization adopted its first code for autonomous ships, the first deep-sea mineral reserves were declared, the ocean recorded its ninth consecutive year of record heat and a war closed one of the world's most important maritime chokepoints.

How to read it

The book is organised in eight parts. Part I sets out the problem and the baseline. Part II examines how the ocean is being

mapped, sensed, modelled and turned into data. Parts III and IV examine food, energy and water. Part V turns to the deep seabed and its minerals and molecules. Part VI looks at living and moving on the sea. Part VII examines power and conflict, including the polar oceans, the Blue Power Index and the distribution of gains and losses among coastal communities. Part VIII confronts limits, governance and the future, and ends with recommendations.

A visual atlas in nine plates, placed after the Executive Summary, shows the scale of the ocean's untapped wealth and the timeline of its possible exploitation. Each chapter opens with a short summary and closes with its key findings, so that readers short of time can follow the argument by reading those alone, starting with the Executive Summary. Readers interested in methods will find the equations and parameters of the nine modules of the Ocean Frontier Model in Annex A, the main data sources in Annex B, country data and profiles in Annexes C and D, and year-by-year model results in Annex G.

A note on numbers

The models in this book produce precise-looking numbers: 92 percent, 47 percent, 2040, 7.0 trillion dollars. They should be read as the outputs of transparent assumptions, not as predictions. Where results depend heavily on an assumption, the relevant chapter shows how they change when the assumption changes. The aim throughout is to bound the plausible, to reveal mechanisms and to identify the variables that decision makers can move.

All monetary values are in US dollars. Facts and figures are drawn from public sources current to November 2026 and are

cited in the text and in the bibliography.

Executive Summary

The question

Seven tenths of our planet is covered by ocean, and most of it has remained beyond human reach. Pressure, darkness, the radio silence of seawater and the cost of ships and crews have confined human activity to the surface, the coasts and a few industrial sites offshore. Artificial intelligence and robotics are removing these barriers. They do not make the ocean less hostile; they make human presence unnecessary. This book asks how far, how fast and with what consequences these technologies will allow humanity to live from the ocean and exploit it.

To answer, Global Strategy Files built the **Ocean Frontier Model**, nine original quantitative modules covering mapping, the cost of ocean work, fisheries, food, offshore wind, deep-sea mining, national ocean power, undersea cables and the ocean economy, with a scenario engine of 50,000 simulated futures. Every assumption is documented in Annex A. The results below are original to this book: to our knowledge, none of these estimates had been published in this form before.

Twelve findings

1. The seafloor map will not be finished by 2030, and at the current pace not this century. The mapped share of the seafloor rose from 6 percent in 2017 to 28.7 percent in 2026, but progress has slowed to about 1.3 points a year.

Accounting for diminishing returns, the status quo reaches 80 percent only around 2085. (Chapter 4)

2. A modest robotic fleet would finish the map by 2040, for less than 2 billion dollars. With 250 dedicated uncrewed survey vessels by 2040, 95 percent of the seafloor would be mapped by 2040; with 800, by 2035. Finishing the deep-ocean map would cost about 3.3 billion dollars with ships and 0.8 to 1.6 billion with uncrewed vessels. The obstacle is coordination, not money. (Chapter 4)

3. Robots are making ocean work several times cheaper. The GSF Ocean Access Cost Index shows robotic methods two to fourteen times cheaper than crewed methods in 2026, and four to forty times cheaper by 2040, across survey, inspection, intervention and sampling. (Chapter 5)

4. AI applied only to fishing efficiency would empty the sea. In the GSF bioeconomic model, AI that raises fishing power by 4 percent a year without better enforcement drives the stylised world fish stock to 13 percent of its optimal level by 2050 and catches to less than a third of the sustainable maximum. (Chapter 8)

5. The 92 percent rule. When fishing technology improves faster than 1.5 percent a year, at least 92 percent of fishing effort must be effectively monitored and constrained to keep stocks healthy. Applied to both efficiency and enforcement, AI keeps biomass near 80 percent of its optimal level and catches at 96 percent of the sustainable maximum. (Chapter 8)

6. Robotic offshore farming could raise blue food per person by up to half. Aquatic food supply rises from 20.7 kilo-

grams per person in 2022 to about 30 kilograms in 2050 with robotic offshore aquaculture (about 50 million tonnes produced offshore), and to nearly 32 kilograms with algae-based feeds, against 25.7 kilograms on the current trend. (Chapter 10)

7. Robotics is worth about 10 billion dollars a year to offshore wind. It cuts the cost of offshore electricity by a median of 7.6 percent for fixed-bottom and 11 percent for floating wind by 2035, worth about 10 billion dollars a year on a 420-gigawatt fleet. (Chapter 11)

8. Deep-sea mining is a coin toss with irreversible stakes. In 50,000 simulations of a 12-million-tonne-a-year nodule project, the probability of a positive private net present value is about 47 percent, falling to about 41 percent if nickel-free batteries dominate and 40 percent with seven years of regulatory delay. The project would disturb about 16,000 square kilometres of seabed. (Chapter 15)

9. The GSF Blue Power Index. The United States and China lead a ranking of 30 countries on seven pillars, followed by the United Kingdom, Japan, France, Australia and Norway. Indonesia has the largest gap between ocean potential and capability; by 2040, robotics rewards capability more than endowment. (Chapter 22)

10. Cutting cables would leave regions degraded for weeks. With current repair capacity, ten simultaneous cable cuts in one region take about 25 days to repair fully, forty cuts about 51 days and eighty cuts about 88 days. Robotic repair and 50

percent more capacity would roughly halve these times. (Chapter 19)

11. The ocean economy will roughly double to about 7 trillion dollars by 2050. Offshore oil and gas declines, while ocean data and robotics, offshore wind, desalination and biotechnology rise from about 11 percent to about a third of ocean value added. (Chapter 26)

12. The most likely future is a fragmented rush. The scenario engine assigns about 48 percent to the Fragmented Rush, in which machines arrive faster than rules; 34 percent to the Steered Ocean, in which technology is matched by governance; and 18 percent to the Saturated Sea, in which climate stress and weak rules degrade the ocean faster than technology compensates. (Chapter 26)

Key numbers at a glance

Indicator	Value	Chapter
Seafloor mapped, 2026	28.7%	4
Year 95% mapped with 250 uncrewed survey vessels	2040	4
Cost to finish the deep-ocean map with robots	$0.8 to 1.6 bn	4
Cost advantage of robotic over crewed ocean work, 2040	4.6x to 40.8x	5
Fish biomass 2050 with AI efficiency only	13% of optimal level	8
Minimum monitoring coverage at high technological creep	about 92%	8
Aquatic food per person 2050, robotic offshore scenario	29.9 kg (2022: 20.7 kg)	10
Offshore aquaculture output 2050, robotic scenario	about 50 Mt	10
Reduction in floating offshore wind cost from robotics, 2035	11% (median)	11
Robotic dividend for offshore wind	about $10 bn a year	11

Indicator	Value	Chapter
Probability a 12 Mt per year nodule project is profitable	about 47%	15
Days to restore 40 simultaneous cable cuts in one region	about 51	19
Ocean economy GVA 2050 (median)	about $7.0 tn (2025: $3.5 tn)	26
Probability of the Fragmented Rush scenario in 2040	about 48%	26

The central lesson

The technologies that make the ocean exploitable are the same technologies that make it governable. The satellite that guides a trawler to the fish can guide a coastguard to the trawler. The robot that maps a nodule field can map a protected area. Every model in this book points to the same conclusion: deployed only for extraction, AI and robotics accelerate the depletion of a warming ocean; deployed for monitoring and enforcement as well, they offer the first realistic chance to manage the sea as a whole.

The question is therefore not whether humanity will exploit the ocean to the maximum. The machines will make that possible. It is whether, for the first time, humanity will choose to exploit it to the optimum.

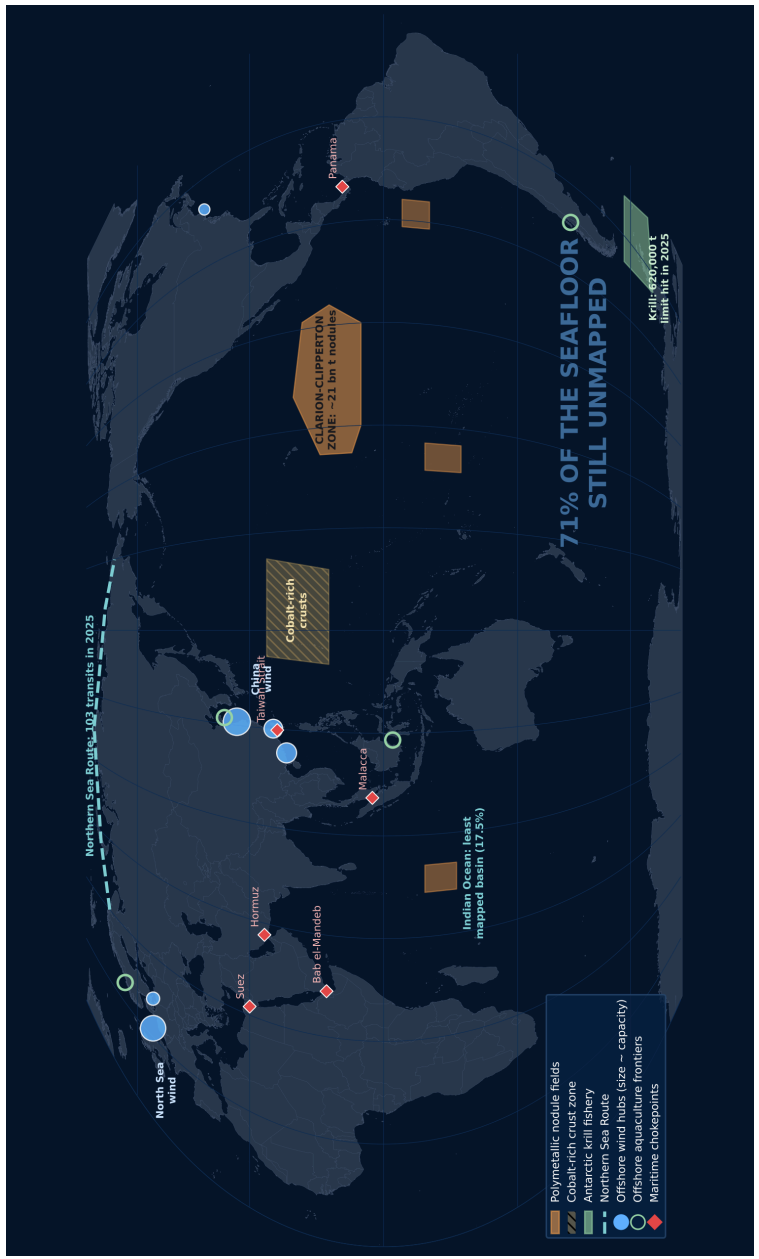
THE UNTAPPED OCEAN: AN ATLAS IN NINE PLATES

Why an atlas

Before the chapters and their models, this atlas shows in nine plates what the book is about: how much wealth the ocean holds, how little of it humanity has touched, and how the automated ocean could unfold between now and 2050. The plates combine published resource estimates with the outputs of the GSF Ocean Frontier Model. They are designed to be read on their own, as a visual summary of the argument.

The geography of an untapped frontier

Plate 1. Atlas of the untapped ocean: resources, routes and pressure points



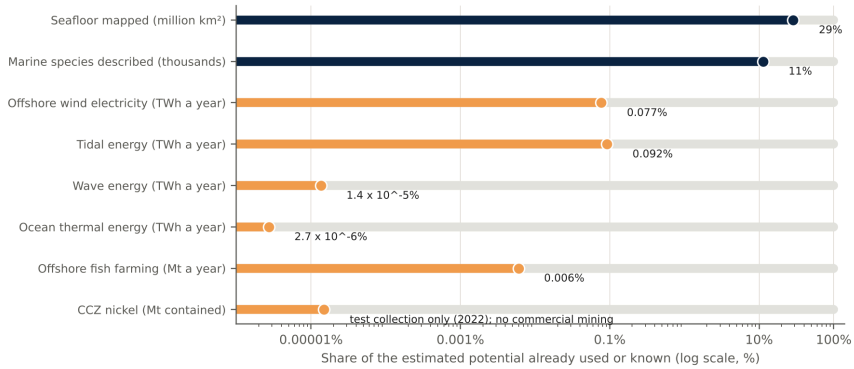
Sources: Hein et al. (2013) and ISA for nodule fields and crust zones (approximate outlines); GWEC (2026) for offshore wind; CCAMLR (2026) for krill; CHNL (2025) for the Northern Sea

Route; Niyazi et al. (2025) for mapping coverage; Global Strategy Files. Equal Earth projection centred on the Pacific. Outlines are schematic.

The map is centred on the Pacific, the ocean that contains most of the untapped deep-sea wealth. The largest orange area is the Clarion-Clipperton Zone, where about 21 billion tonnes of polymetallic nodules lie on the abyssal plain. Smaller nodule fields lie in the Peru Basin, the Central Indian Ocean and around the Cook Islands; cobalt-rich crusts coat the seamounts of the western Pacific. The blue circles show where offshore wind is concentrated today: the North Sea and the coasts of China. Green rings mark the frontiers of offshore aquaculture, from Norway to the Yellow Sea, Indonesia and Chile. Red diamonds mark the chokepoints through which the wealth of the sea, and much of the world's trade, must pass. At the poles, the Northern Sea Route and the Antarctic krill fishery show where melting ice is opening new frontiers. Everywhere else, the dark ocean is a reminder of the central fact of this book: seven tenths of the seafloor remain unmapped.

How little has been touched

Plate 2. The exploitation ladder: share of estimated potential already used or known



Sources: Seabed 2030 (2026); WoRMS and Mora et al. (2011) for species (about 248,000 described of an estimated 2.2 million); IEA (2019) for offshore wind potential (420,000 TWh a year); IRENA and Ocean Energy Europe compilations for tidal, wave and ocean thermal potential; Gentry et al. (2017) for offshore fish farming potential (15 billion tonnes a year); Hein et al. (2013) for nickel contained in CCZ nodules; GSF estimates of current use. Log scale.

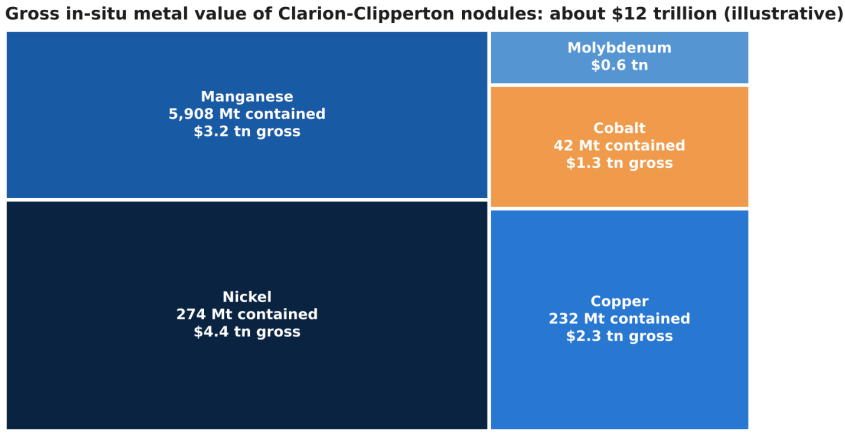
Plate 2 measures, on a logarithmic scale, the share of each ocean resource that humanity already uses or knows. The two best-known dimensions are knowledge itself: about 29 percent of the seafloor is mapped and roughly a tenth of marine species have been described. Every productive resource is far less developed. Offshore wind, the most successful ocean energy industry, produces less than a tenth of one percent of the electricity that the International Energy Agency estimates offshore wind could technically generate, more than 420,000 terawatt-hours a year, eighteen times today's global electricity demand. Tidal energy uses a similar sliver of its resource; wave and ocean thermal energy have barely started. Offshore fish farming produces less than one hundredth of one percent of the 15 billion tonnes a year that a 2017 study in *Nature Ecology & Evolution* estimated suitable ocean areas could theoretically yield. The nickel in the

Clarion-Clipperton nodules, roughly twice the land-based reserve base according to Hein and colleagues, has been touched only by a 3,000-tonne test collection.

The point is not that all this potential should be exploited. Most of it should not. The point is that the ocean's frontier is not marginal: it is, in almost every dimension, larger than everything humanity has developed so far.

The metal endowment of a single abyssal plain

Plate 3. Gross in-situ value of the main metals in Clarion-Clipperton nodules



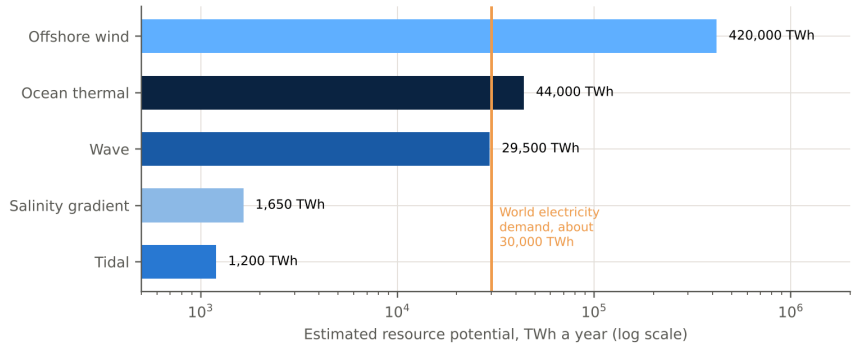
Sources: Hein et al. (2013) for nodule tonnage (21,100 million dry tonnes) and mean grades (manganese 28%, nickel 1.3%, copper 1.1%, cobalt 0.2%, molybdenum 587 g/t); illustrative prices of $545 per tonne of manganese (as silicate), $16,000 nickel, $9,800 copper, $30,000 cobalt and $45,000 molybdenum. Gross in-situ value is not an economic value: it ignores recovery, costs and price effects.

Plate 3 multiplies the estimated tonnage of nodules in the Clarion-Clipperton Zone by their average metal grades. The result is staggering: about 5,900 million tonnes of manganese, 274 million tonnes of nickel, 232 million tonnes of copper and 42 million tonnes of cobalt, with a gross in-situ value of roughly 12 trillion dollars at recent prices. Nickel alone accounts for more than a third of that value.

This number must be read with great care, and Chapter 15 explains why. Gross in-situ value is not wealth that can be captured: only a fraction of the nodules can be recovered, extraction and processing are costly, and bringing even a small part of this tonnage to market would depress prices. The economic model of Chapter 15 finds that a single large project has roughly a coin toss's chance of being profitable. The plate shows the size of the prize that drives the deep-sea rush; the model shows how uncertain the prize really is.

The energy of the sea

Plate 4. Estimated resource potential of ocean energy sources compared with world electricity demand



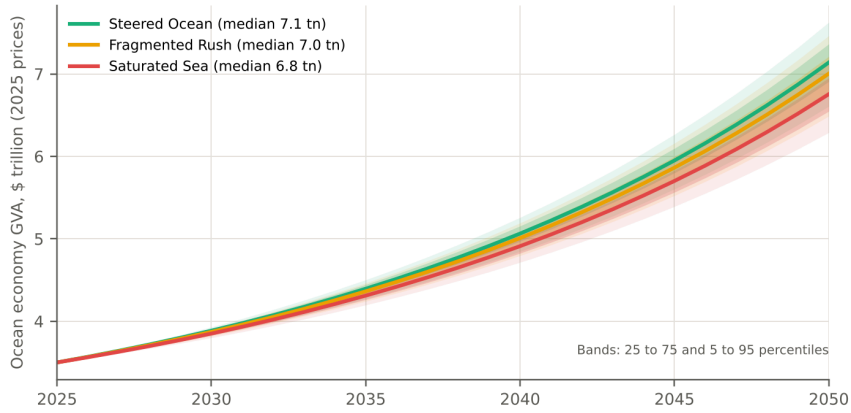
Sources: IEA, Offshore Wind Outlook 2019 (offshore wind, technical potential); IRENA compilations as reported by ENSEC COE (2025) for ocean thermal, wave, salinity gradient and tidal energy (theoretical or resource potential); world electricity demand rounded. Log scale. Potentials are not directly comparable: they combine technical and theoretical estimates.

The ocean is the largest renewable energy reservoir on the planet. Offshore wind alone has a technical potential many times global electricity demand. Ocean thermal energy and waves each exceed it on paper. Salinity gradients and tides are smaller but highly predictable. Plate 4 also shows why offshore wind dominates today: its resource is the largest and its technology the most mature. The others are frontiers in

the strict sense: resources that are known to exist at enormous scale but that no industry yet knows how to harvest economically. Chapter 12 explains why robotics, by lowering the cost of building and maintaining machines in hostile water, is the key that could unlock them.

Three trajectories for the ocean economy

Plate 5. Fan chart of ocean economy gross value added under three scenarios, 2025 to 2050

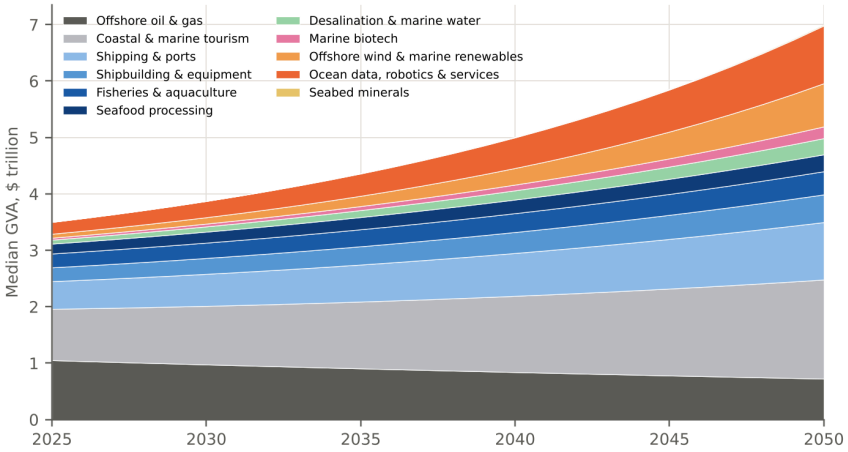


Source: GSF Ocean Frontier Model, module M9, 20,000 simulations; median and percentile bands by scenario.

Plate 5 shows how the ocean economy could grow from 3.5 trillion dollars in 2025. In all three scenarios it roughly doubles by 2050, but the paths diverge after the mid-2030s. The Steered Ocean, in which automation is matched by effective governance, reaches a median of about 7.1 trillion dollars; the Fragmented Rush about 7.0 trillion; the Saturated Sea, in which climate stress and weak rules erode the ocean's productivity, about 6.8 trillion. The differences in measured output look modest. The differences in unmeasured wealth, fish stocks, reefs, biodiversity and coastal protection, are far larger, as Chapters 8 and 24 show.

The changing anatomy of the ocean economy

Plate 6. Median gross value added by sector, 2025 to 2050

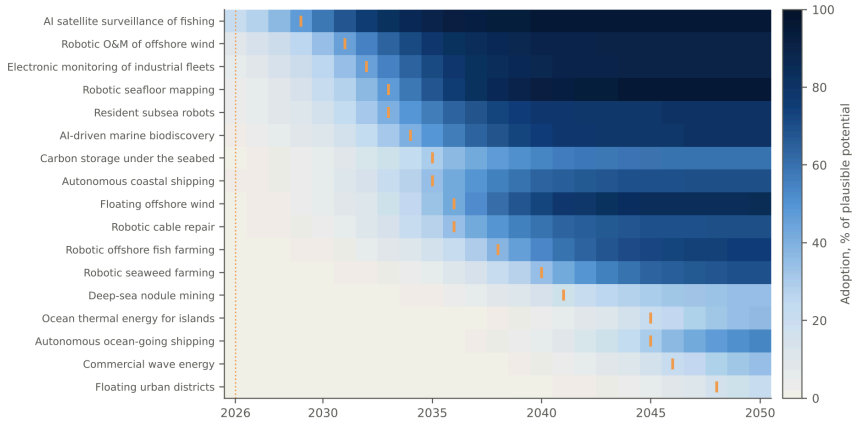


Source: GSF Ocean Frontier Model, module M9, median paths by sector.

Plate 6 shows the same growth decomposed by sector. At the bottom, offshore oil and gas shrinks steadily with the energy transition. Tourism and shipping continue to grow. At the top, the sectors born of the automated ocean, ocean data and robotics services, offshore wind, desalination and marine biotechnology, swell from a thin band in 2025 to about a third of the total by 2050. Seabed minerals, despite their enormous in-situ value, remain a hairline at the top of the chart, a reminder that physical endowment and economic value are not the same thing.

The timeline of a frontier that has not yet opened

Plate 7. Adoption of seventeen ocean technologies and activities, 2026 to 2050



Source: Global Strategy Files expert assessment. Each row is a logistic adoption curve expressed as a share of plausible potential; orange marks indicate the year of half adoption. Illustrative.

Plate 7 is a calendar of the ocean frontier. Each row represents an activity; colour intensity shows how far it has progressed toward its plausible potential in each year. The first wave, already under way, is about seeing and monitoring: AI surveillance of fishing, robotic maintenance of wind farms, electronic monitoring of fleets and robotic seafloor mapping, all reaching half of their potential before the mid-2030s. The second wave, around the late 2030s, is about producing: floating wind, offshore fish farms, seaweed farming, robotic cable repair. The third wave, in the 2040s, includes the activities that remain uncertain, limited or contested: deep-sea mining, ocean thermal and wave energy, autonomous ocean-going ships and floating urban districts. The plate expresses a central conclusion of this book: the automated ocean begins with knowledge and control, and only later turns to large-scale extraction. The order matters, because the rules written during the first wave will govern the second and third.

The shape of ocean power

Plate 8. Blue Power Index profiles of the eight leading ocean nations

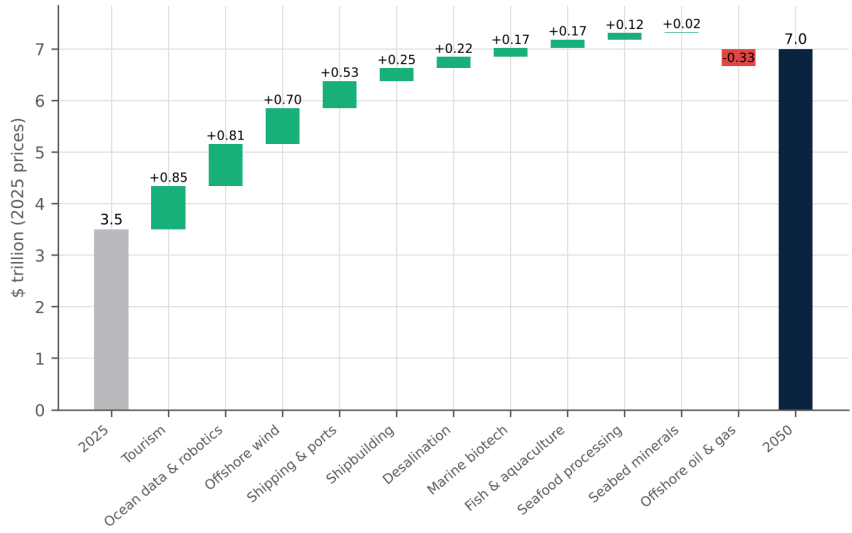


Source: GSF Ocean Frontier Model, module M7. Seven pillars scored 0 to 100: space, food, energy and water, connectivity, machines, force and rules.

Plate 8 shows that ocean power comes in different shapes. The United States combines vast maritime space with the strongest machines and naval force. China's profile leans toward food, energy and industrial capability rather than space. The United Kingdom, France and Australia combine large maritime zones with strong technology and governance. Japan and Norway are machine powers. These shapes matter for the automated ocean, because robotics rewards capability more than endowment (Chapter 22).

Where the next three and a half trillion dollars come from

Plate 9. Decomposition of the growth of ocean value added between 2025 and 2050



Source: GSF Ocean Frontier Model, module M9, median sector values.

Plate 9 decomposes the median increase in ocean value added between 2025 and 2050, from about 3.5 to 7.0 trillion dollars. Tourism and shipping remain large contributors, but the two largest additions after tourism come from the automated ocean itself: ocean data and robotics services and offshore wind, each adding between 0.7 and 0.8 trillion dollars. Offshore oil and gas subtracts about a third of a trillion. The next ocean economy will be built less on what is pumped from beneath the sea than on what is measured, generated and managed in it.

PART I. THE LAST FRONTIER

Chapter 1. Seventy-One Percent Unknown

> **In brief.** The ocean covers seven tenths of the planet, yet in 2026 only 28.7 percent of its floor was mapped to modern standards. Four walls, pressure, darkness, the radio silence of seawater and cost, kept humans out. AI and robotics breach them not by making the ocean less hostile but by making human presence unnecessary. This chapter sets out the book's central paradox: the same machines that make the ocean exploitable also make it governable.

1.1 A planet named for the wrong surface

We call our planet Earth, but seven tenths of its surface is water. The ocean covers about 361 million square kilometres, holds 97 percent of the planet's water, absorbs more than nine tenths of the excess heat trapped by greenhouse gases and supplies roughly half of the oxygen we breathe. It feeds three billion people who depend on fish for a significant share of their animal protein. It carries about four fifths of world trade by volume. More than 99 percent of the data that crosses between continents travels through fibre cables lying on its floor. And yet, measured by the simplest test of knowledge, which is whether we know the shape of the ground beneath it, the ocean is still mostly a blank.

When the Nippon Foundation and the General Bathymetric Chart of the Oceans launched the Seabed 2030 project at the United Nations Ocean Conference in 2017, only 6 percent of the world's seafloor had been mapped to modern standards. By the spring of 2026 the figure had reached 28.7 percent, after a record year in which almost five million square kilometres of new data were added. That progress is real and it is accelerating. But it means that in 2026, more than seventy percent of the floor of our own planet remains unmapped at a resolution we would consider elementary on land. We have better maps of the surface of Mars.

This book is about what happens when that blank begins to fill, and about who will fill it. Its central argument is simple. For the whole of human history, the ocean resisted exploitation not because it lacked resources but because it was too hostile, too dark, too deep and too expensive for human beings to work in. Artificial intelligence and robotics are removing that resistance. They do not make the ocean less hostile. They make human presence unnecessary. And once human presence is no longer the constraint, the economics, the geopolitics and the ecology of seventy percent of the planet change at the same time.

The question that runs through every chapter is the one a reader might ask on first hearing this argument: how far will AI and robots allow humanity to live from the ocean and exploit it to the maximum? The answer this book gives, built on nine original models and fifty thousand simulated futures, is more interesting than either the optimists or the pessimists expect. The machines will open almost the whole ocean to economic use within a generation. But the same machines will also, for the first time, make it possible to see exactly

what is being taken, where and at what cost. Whether the automated ocean becomes a managed commons or a plundered frontier will not be decided by engineers. It will be decided by the rules that states choose to write and enforce in the next fifteen years.

1.2 Four walls around the ocean

Why has the ocean resisted for so long? It is worth being precise, because each obstacle explains a specific feature of the technologies that are now overcoming it.

The first wall is pressure. Water pressure rises by about one atmosphere every ten metres. At 4,000 metres, the typical depth of the abyssal plains where polymetallic nodules lie, a structure must withstand roughly 400 times the pressure at the surface, around 40 megapascals, the weight of a small car pressing on every square centimetre. At the bottom of the Challenger Deep, close to 11,000 metres, the pressure exceeds 1,000 atmospheres. Human divers working in saturation can operate for weeks at a few hundred metres at most, at great cost and real risk. Below that, every human visit requires a pressure hull, and the world's fleet of crewed submersibles capable of reaching the deep sea can be counted on two hands.

The second wall is darkness. Sunlight fades exponentially in seawater. Below about 200 metres, the twilight zone begins; below 1,000 metres there is no usable sunlight at all. Even artificial light carries only a few metres through turbid water. This is why the ocean is mapped with sound rather than images. Multibeam echosounders fire fans of acoustic pulses at the seafloor and measure their return time; the width of the strip they map, the swath, grows with depth, which means

paradoxically that the shallow coastal seas are the slowest to map, while the deep plains can be swept in wide bands.

The third wall is the radio silence of seawater. This is the least appreciated obstacle and arguably the most decisive. Electromagnetic waves, which carry radio, satellite links, GPS and Wi-Fi, are absorbed by salt water within metres. A machine below the surface cannot receive positioning from satellites, cannot stream video to an operator and cannot receive instructions, except through a physical cable or through acoustic links whose bandwidth is measured in kilobits per second, roughly what a dial-up modem offered in the 1990s. A remotely operated vehicle tethered to a ship can be steered in real time; a free-swimming vehicle cannot. It must navigate by inertial sensors and acoustic beacons, decide for itself what to do when it meets the unexpected, and store its data until it surfaces. Autonomy, in other words, is not a luxury in the ocean. It is the physical precondition for any untethered work.

The fourth wall is cost. All three physical barriers translate into money. A deep-water research vessel costs tens of thousands of dollars per day to operate, before scientific equipment. A dynamically positioned construction vessel capable of deploying heavy remotely operated vehicles can cost well over 100,000 dollars per day. Weather windows shorten campaigns; transit times to remote areas consume days before any work begins. Ocean knowledge has therefore historically

advanced along narrow ship tracks, wherever a vessel happened to pass, leaving vast areas between them untouched.

Table 1.1. The four walls around the ocean and how machines breach them

Barrier	Physical constraint	Consequence for human work	How AI and robotics breach it
Pressure	+1 atmosphere per 10 m; about 400 atmospheres at 4,000 m	Divers limited to a few hundred metres; crewed submersibles rare and costly	Uncrewed pressure-tolerant vehicles; no life support, smaller hulls, lower cost
Darkness	Sunlight negligible below 200 m to 1,000 m	Visual inspection impossible at range; mapping depends on sound	Acoustic sensing, sonar processing and computer vision trained on low-light imagery
Radio silence	Electromagnetic waves absorbed within metres	No GPS, no live control, no real-time data from untethered vehicles	Onboard autonomy, inertial and acoustic navigation, edge AI that decides without a human
Cost	Ship day rates of 30,000 to 150,000 dollars and more	Sparse sampling, short campaigns, knowledge along ship tracks only	Uncrewed surface vessels, long-endurance gliders, swarms, resident subsea robots

Sources: physical constants from standard oceanography; day rates are indicative industry ranges compiled by Global Strategy Files.

1.3 What we know, and what we do not

The extent of our ignorance is uneven, and the unevenness matters strategically. A 2025 study in *Frontiers in Marine Science* reviewing progress between 2019 and 2024 found that the North Atlantic and North Pacific were the best mapped basins, with over a third of each mapped by the end of 2024, while the Indian Ocean remained the least mapped, with only 17.5 percent coverage. In other words, the oceans that bor-

dered the industrial powers of the twentieth century are the best known; the oceans that border the fastest growing populations of the twenty-first century are the least known.

The gaps go well beyond bathymetry. We do not know how many species live in the deep sea; estimates of the undescribed fraction routinely exceed 80 percent for some groups. We have only fragmentary time series of oxygen, acidity and temperature below 2,000 metres. We have almost no direct observation of the ocean under sea ice. We know the location of the main fishing grounds, but, as satellite and machine learning analysis has shown, we did not until recently know who was fishing them: a study published in *Nature* in January 2024 found that about three quarters of the world's industrial fishing vessels were not publicly tracked.

Table 1.2. The state of ocean knowledge in 2026

Domain	What is known	What is largely unknown
Seafloor shape	28.7% mapped to modern standards (spring 2026)	71.3% of the seafloor, especially the Indian Ocean, Southern Ocean and coastal shallows
Water column	Upper 2,000 m sampled by autonomous profiling floats	Deep ocean below 2,000 m, polar seas under ice, fine-scale mixing
Life	Commercial species and their main stocks	Most deep-sea species, microbial diversity, genetic resources
Human activity	Vessels transmitting identification signals	About three quarters of industrial fishing vessels, a quarter of transport and energy vessels (2017 to 2021 data)
Resources	Main nodule fields, major oil and gas basins	Most hydrothermal sulphide deposits, cobalt crusts, biological resources

Sources: Seabed 2030; Niyazi et al. (2025), *Frontiers in Marine Science*; Paolo et al. (2024), *Nature*; Global Strategy Files.

1.4 Why the ignorance is ending now

Three things changed in the decade before this book was written, and together they explain why the 2020s are the decade in which the ocean begins to open.

The first is the collapse in the cost of sensors and computing. A multibeam echosounder, an inertial navigation unit and a processor capable of running a neural network now fit inside a vehicle a few metres long. A decade ago, the same package required a ship.

The second is the maturing of autonomy. Machine learning allows a vehicle to classify what it sees, to recognise a pipeline, a cable, a species or an anomaly in the sediment, and to change its plan without waiting for a human. Uncrewed surface vessels now carry out survey work commercially. In June and July 2024, a 20-metre Saildrone Surveyor completed a 26-day deep-water cable route survey in the North Atlantic for Meta, mapping more than 4,500 kilometres of seabed without a port call or outside assistance. The survey industry reports that uncrewed surface vessels can use about 95 percent less fuel than traditional survey ships. When a robot replaces a ship, it does not just cut emissions. It removes the crew, the galley, the accommodation, the safety margins, and therefore most of the cost.

The third change is demand. Offshore wind, subsea cables, aquaculture, desalination, critical minerals and naval competition all require knowledge of the ocean that did not matter much when the sea was merely something to sail across. Offshore wind capacity reached 92.5 gigawatts at the end of 2025 and the Global Wind Energy Council expects about 420 gigawatts by 2035. Every one of those turbines needs a seabed

survey before it is built and continuous inspection after. Demand for ocean work is therefore growing faster than the supply of ships and sailors able to do it. The gap is being filled by machines.

1.5 From crossing the ocean to operating it

For most of history, humanity's relationship to the ocean has been one of transit and harvest. Ships crossed it. Fishers harvested its surface layers. Navies controlled its lanes. The ocean floor and the water column below a few hundred metres were, economically speaking, irrelevant.

The automated ocean is different in kind. Its defining feature is persistent presence: machines that stay at sea for months, sensors that report continuously, digital models that simulate the ocean in near real time, and production facilities (fish farms, wind farms, desalination intakes, subsea processing stations, eventually mining systems) that operate far from shore with few or no people on site. The relationship shifts from crossing the ocean to operating it.

That shift has four consequences that structure this book.

1. **Seeing.** The ocean becomes measurable. Mapping, monitoring and forecasting improve by orders of magnitude, and the cost of knowledge falls (Part II).
2. **Producing.** Food, energy, fresh water and molecules can be produced offshore at industrial scale, largely by robots supervised from land (Parts III and IV).
3. **Extracting.** The deep seabed becomes technically accessible to mining, and the question shifts from whether it can be done to whether it should be, and who decides (Part V).

4. **Controlling.** Fleets of autonomous vehicles become instruments of national power, both to protect undersea infrastructure and to threaten that of others (Part VII).

Two further parts complete the picture. Part VI examines what it means to live on and under the sea, from floating cities to ships without sailors. Part VIII confronts the limits: a warming ocean under unprecedented thermal stress, the rebound effect by which efficiency gains increase total extraction, and the governance frameworks that will determine whether the automated ocean is steered or simply seized.

1.6 The method of this book

Books about the ocean tend to fall into two genres. The first is descriptive and celebratory, cataloguing new technologies and wonders of the deep. The second is alarmist, cataloguing threats and collapses. Both are useful, but neither answers the questions a decision maker actually faces. How fast will the map be finished? How much more fish will robotic aquaculture produce, and will AI-assisted fishing empty the sea? How much cheaper will offshore electricity become? Is deep-sea mining profitable at realistic metal prices? Which countries are positioned to capture the value of the automated ocean, and which risk watching others capture it in their own waters? How long would a region stay disconnected if its cables were cut?

To answer these questions, Global Strategy Files built the **Ocean Frontier Model**, an integrated set of nine quantitative modules, documented in full in Annex A:

Table 1.3. The nine modules of the GSF Ocean Frontier Model

Module	Name	Question	Method
M1	Mapping Clock	When will the seafloor be fully mapped, and at what cost?	Trend analysis of Seabed 2030 data; logistic fleet scenarios; cost model of ships versus uncrewed vessels
M2	Ocean Access Cost Index	How fast is the cost of ocean work falling?	Learning curves (Wright's law) applied to four categories of ocean work
M3	Fishing Under Algorithmic Watch	Does AI empty or save fish stocks?	Bioeconomic surplus production model with technological creep and enforcement
M4	Blue Food 2050	How much aquatic food per person in 2050?	Scenario projections of capture, aquaculture and offshore farming
M5	Robotic Wind Dividend	What is robotics worth to offshore wind?	Monte Carlo on cost of electricity, operations and availability
M6	Deep-Sea Economics	Can nodule mining pay, and at what ecological price?	Monte Carlo net present value with metal prices, battery chemistry and regulatory delay
M7	Blue Power Index	Which nations are equipped for the automated ocean?	Composite index of 30 countries on seven pillars
M8	Cable Repair Queue	How long would a region stay cut off?	Queueing simulation of the global repair fleet under attack scenarios
M9	Ocean Economy 2050 and Scenarios	What will the ocean economy be worth, and which future is most likely?	Sector growth model and scenario engine with 50,000 draws

The models are deliberately transparent. Every parameter is listed with its source or, where the evidence is thin, with the judgement that underlies it. Where a result depends heavily

on an assumption, the book says so and shows how the result changes when the assumption changes. The aim is not to predict the future with false precision, but to put numbers on the questions that matter, to bound the range of plausible outcomes and to identify the variables that decision makers can actually move.

1.7 The paradox at the heart of the automated ocean

One finding runs through almost every module, and it is worth stating at the outset because it is the key to the whole book.

The technologies that make the ocean exploitable are the same technologies that make it governable. An uncrewed vessel that maps a nodule field for a mining company can map a marine protected area for a regulator. A computer vision system that counts salmon in a pen can count the catch on a trawler's deck. A satellite and machine learning pipeline that tells a fishing fleet where the tuna are can tell a coastguard where the fishing fleet is. For most of history, the ocean was ungovernable because it was invisible. That excuse is disappearing.

The models in this book show what follows. Where AI is applied only to extraction, the results are stark: in the bioeconomic model of Chapter 8, AI-boosted fishing efficiency without matching enforcement drives fish biomass in 2050 down to 13 percent of the level that maximises sustainable yield. Where the same AI is applied to monitoring and enforcement as well, biomass stabilises at about 80 percent of that level and catches remain close to the maximum the ocean can sustain. The difference between those two futures is not technological. It is a policy choice.

The title of this book, *The Automated Ocean*, is therefore descriptive rather than celebratory. Automation is coming to the ocean with near certainty; the models leave little doubt about that. What remains open is the purpose to which it is put. The question is not whether humanity will exploit the ocean to the maximum. It is whether it will learn, for the first time, to exploit it to the optimum.

1.8 A short history of ocean ignorance

The scale of what remains unknown is easier to grasp when set against the history of how the ocean has been measured.

For most of history, the depth of the sea was measured with a weighted line lowered until it touched bottom. The British research expedition of HMS *Challenger*, between 1872 and 1876, made just under 500 deep soundings across all the oceans, each taking hours. It revealed that the deep ocean floor was not a featureless basin but a landscape of plains, ridges and trenches, and it discovered polymetallic nodules, the deposits that are at the centre of Chapter 14 a century and a half later.

The echo sounder, developed in the 1920s, replaced the line with sound. A ship could now measure depth continuously while under way, and the mid-ocean ridges and great fracture zones were mapped in outline over the following decades, a body of work that underpinned the theory of plate tectonics. But each echo sounder measured only a single point beneath the ship. Between ship tracks, the map was interpolation.

The multibeam echosounder, which came into civilian use from the 1970s and 1980s, was the decisive advance. By firing a fan of beams, it mapped a strip of seafloor several times as

wide as the water depth. Modern standards of seafloor mapping, the standards used by Seabed 2030, are essentially multibeam standards.

In the 1990s, satellites added a different kind of map. Altimeters measuring the height of the sea surface could detect the tiny bulges caused by the gravitational pull of seamounts and ridges, allowing scientists to estimate seafloor shape everywhere at once. These satellite-derived maps are global but coarse, resolving features several kilometres across at best. They are the reason we can say that we have "a map" of the whole ocean floor, and the reason specialists insist that we do not: a satellite gravity map would miss most hills, canyons, wrecks and cables, and could not support the design of a wind farm foundation or the planning of a mining operation.

Table 1.4. Milestones in measuring the ocean floor

Period	Method	What it revealed	Limitation
Before 1920	Weighted lines	Isolated depth points; the deep sea is not uniform	Hours per sounding; a few hundred deep soundings worldwide
1920s to 1960s	Single-beam echo sounder	Mid-ocean ridges, trenches, continental margins	One point beneath the ship
1970s onward	Multibeam echosounder	Detailed swaths of seafloor	Coverage limited to ship tracks
1990s onward	Satellite altimetry and gravity	Global seafloor shape at coarse resolution	Features smaller than several kilometres missed
2017 onward	Seabed 2030 compilation, crowdsourcing, uncrewed vessels	Global grid of all available multibeam data	71 percent still unmapped in 2026

Source: Global Strategy Files.

1.9 Why ignorance is a strategic variable

It might seem that ignorance of the ocean is merely a scientific gap, to be closed in due course. The argument of this book is that it is also a strategic variable, one that shapes who can exploit the sea and who cannot.

Knowledge of the ocean is not evenly distributed. The navies of a few states hold detailed bathymetric, acoustic and oceanographic data on large areas, including other states' waters, which they do not share. Oil and gas companies hold high-resolution surveys of the basins they have explored. Cable companies hold detailed route surveys. Fishing companies know where the fish are. Most coastal developing states know far less about their own waters than these actors do.

In such a world, the party that knows more negotiates better. A coastal state that does not know whether its seabed contains valuable minerals cannot judge the value of an exploration licence. A fisheries agency that cannot see the vessels in its zone cannot price access fairly or detect theft. A government that cannot model its coastal waters cannot plan aquaculture or protect its coastline. The automated ocean will either reduce these asymmetries, if data are shared and capacity is built, or entrench them, if knowledge remains concentrated in the hands of those who own the machines. This theme returns in Chapters 4, 6, 22 and 25.

1.10 Key findings of Chapter 1

- The ocean covers 71 percent of the planet, but in spring 2026 only 28.7 percent of its floor had been mapped to modern standards, up from 6 percent in 2017.

- Four barriers explain the ocean's long resistance: pressure, darkness, the radio opacity of seawater and cost. Of these, radio silence makes autonomy the decisive technology.
- Falling sensor and computing costs, maturing machine autonomy and fast rising demand from offshore industries are converging in the 2020s to open the ocean.
- The relationship between humanity and the sea is shifting from crossing and harvesting to persistent operation by machines supervised from land.
- The same technologies enable both exploitation and governance. Which use prevails is the central question of the coming generation.

Chapter 2. The Machine Turn

> **In brief.** Ocean access has passed from the age of the diver to the age of the vessel and is entering the age of the autonomous machine. This chapter describes the families of ocean robots, what artificial intelligence adds to them, the energy and communications limits they face, and why autonomy changes which ocean activities are economically possible at all.

2.1 Three ages of ocean access

The history of human access to the ocean below its surface can be divided into three ages, each defined by who, or what, does the work.

The age of the diver lasted from antiquity until the middle of the twentieth century. Pearl divers, sponge divers and later helmet divers in canvas suits worked within a few dozen metres of the surface. Their reach defined the economically useful ocean: a thin coastal fringe and the surface layer exploited by fishing nets.

The age of the vessel began after the Second World War. In January 1960, the bathyscaphe *Trieste* carried two men to the bottom of the Challenger Deep. In 1964, the research submersible *Alvin* entered service and went on to make thousands of dives, including the exploration of hydrothermal

vents in the late 1970s that transformed biology. In 1966, the United States Navy used an early tethered robot, CURV, to recover a hydrogen bomb lost off the Spanish coast near Palomares, demonstrating that machines on cables could work where humans could not. Over the following decades, the offshore oil and gas industry turned the remotely operated vehicle, the ROV, into a standard industrial tool. A ship on the surface, a cable, a robot with arms and cameras at the bottom, a team of pilots in a control cabin: this triad built the subsea oil industry of the North Sea, the Gulf of Mexico, Brazil and West Africa.

The age of the autonomous machine is the one now beginning. Its distinguishing feature is not that machines go deep (they have done so for sixty years) but that they no longer need a ship and a crew to accompany them. An autonomous underwater vehicle, an AUV, is launched, carries out its mission and returns or surfaces on its own. A glider drifts and dives for months on the energy of a few batteries. An uncrewed surface vessel crosses an ocean basin under remote supervision from an operations centre on land. A resident subsea robot lives permanently on the seabed next to the infrastructure it maintains, docking to recharge and upload data.

The difference between the second and third ages is economic before it is technical. In the age of the vessel, the marginal cost of ocean work was the day rate of a ship and its crew. In the age of the autonomous machine, it is increasingly the cost of the robot itself, amortised over years of operation, plus a small team supervising many units at once. This is the transition that the Ocean Access Cost Index in Chapter 5 attempts to measure.

2.2 A taxonomy of ocean robots

The family of ocean machines is diverse, and the distinctions matter because each type has a different economic role.

Table 2.1. The main families of ocean robots in 2026

Type	How it works	Typical depth	Typical endurance	Main uses	Autonomy level
Remotely operated vehicle (ROV)	Tethered to a ship by an umbilical cable; piloted by humans	To 6,000 m and beyond	As long as the ship stays	Intervention, construction, repair, sampling	Low to medium (assisted piloting)
Autonomous underwater vehicle (AUV)	Free swimming, battery powered, preprogrammed or adaptive missions	To 6,000 m	Hours to a few days	Seabed mapping, pipeline inspection, mine countermeasures, science	Medium to high
Underwater glider	Changes buoyancy to glide up and down; very low power	To 1,000 m (some deeper)	Weeks to months	Oceanographic sampling, acoustic monitoring	Medium
Profiling float	Drifts with currents, dives and surfaces on a cycle	To 2,000 m (deep floats to 6,000 m)	Several years	Global temperature and salinity monitoring	Low (fixed cycle)
Uncrewed surface vessel (USV)	Surface craft, wind, solar or diesel powered, remotely supervised	Surface, with hull sonars	Weeks to a year	Bathymetry, metocean data, communications relay, patrol	Medium to high
Resident subsea vehicle	Lives on the seabed, docks to charge and upload	Field depth	Months between servicing	Inspection and light intervention at oil, gas and wind sites	High
Extra-large uncrewed undersea vehicle (XLUUV)	Submarine-sized autonomous vessel	Military depths	Weeks to months	Surveillance, mine laying, payload delivery	High

Type	How it works	Typical depth	Typical endurance	Main uses	Autonomy level
Aerial drone over sea	Flying platform launched from ship, shore or turbine	Above surface	Minutes to hours	Turbine blade inspection, surveillance, search and rescue	Medium to high

Source: Global Strategy Files, compiled from manufacturer and industry documentation.

Two developments within this taxonomy deserve emphasis.

The first is the **convergence of surface and subsurface systems**. A surface vessel can now act as a mothership for underwater drones, launching and recovering them, relaying their acoustic signals to satellites and recharging them. This combination largely solves the radio silence problem described in Chapter 1: the surface vehicle bridges the water and the sky. A single human operator on land can supervise a surface vessel, several underwater vehicles and an aerial drone working together on the same task.

The second is **residency**. For decades, every subsea operation began with mobilising a ship. Resident systems invert this logic. The robot is already there, waiting in a docking station on the seabed, and can be tasked from shore at any time. The offshore oil and gas industry pioneered this approach, and offshore wind is now adopting it. Residency transforms ocean work from a sequence of expensive expeditions into a continuous service.

2.3 What artificial intelligence adds

It is common to speak loosely of "AI-powered" ocean robots. It is more useful to be specific about the functions that machine

learning and related techniques now perform, because each one removes a different human bottleneck.

Perception. Sonar images, low-light video and acoustic recordings are noisy and hard to interpret. Deep learning models trained on labelled examples can now detect and classify objects in this data: a crack in a weld, a free span under a pipeline, a buried cable, a species of coral, a mine on the seabed, a vessel in a satellite radar image. What once required a human analyst to watch hours of video can be done in real time on the vehicle.

Navigation and localisation. Without GPS, an underwater vehicle must estimate its own position from inertial sensors, Doppler velocity logs and acoustic beacons. Errors accumulate over time. Techniques that match what the vehicle sees to a map of the terrain, and that build the map while navigating, allow long missions with bounded error.

Adaptive mission planning. A preprogrammed vehicle follows a lawnmower pattern regardless of what it finds. An adaptive vehicle changes its plan: it returns to inspect an anomaly at higher resolution, follows a plume of discoloured water to its source, or skips an area already covered by a partner vehicle. This is the step from automation to autonomy.

Fleet coordination. When several vehicles work together, someone must allocate tasks, avoid collisions and fuse data. Increasingly, this is done by algorithms, allowing one human to supervise many machines. The ratio of operators to vehi-

cles is the single most important determinant of the economics of robotic ocean work.

Prediction and digital twins. At the scale of the ocean basin, AI models trained on decades of observations and simulations now produce forecasts of currents, temperature and waves far faster than traditional physical models. Mercator Ocean International's GLONET system, piloted within the European Digital Twin of the Ocean, illustrates the approach (Chapter 6).

Decision support. Finally, AI translates data into decisions: where to place a wind turbine, when to harvest a salmon pen, which vessel to inspect for illegal fishing, when to dispatch a repair robot. This is where the value of all the other functions is realised.

Table 2.2. A ladder of ocean autonomy

Level	Description	Human role	Example in 2026
0	Manual	Human performs the task directly	Diver, crewed submersible
1	Remote control	Human pilots the machine in real time through a cable	Work-class ROV on a construction vessel
2	Assisted	Machine stabilises and holds position; human directs	ROV with automatic station keeping
3	Supervised autonomy	Machine executes a planned mission; human monitors and intervenes	Survey AUV, uncrewed survey vessel
4	Adaptive autonomy	Machine changes its own plan based on what it detects	Inspection AUV that revisits anomalies; adaptive gliders
5	Collective autonomy	Several machines coordinate among themselves; human sets goals	Early swarm trials; resident systems in pilot projects

Source: Global Strategy Files, adapted from autonomy level frameworks used in automotive and maritime regulation.

Most commercial ocean work in 2026 sits at levels 1 to 3. The economic transformation analysed in this book corresponds to the shift of a large share of that work to levels 3 to 5 between now and 2040.

2.4 The market that is forming

The commercial market for these machines is still small compared with the industries it serves, but it is growing fast. Market research firms disagree widely on its size, because definitions vary (some include military programmes, some only civilian vehicles, some include services and software). MarketsandMarkets estimated the global market for unmanned underwater vehicles at about 6.9 billion dollars in 2026, with a projected rise to around 19 billion dollars by 2031, which would imply growth above 20 percent a year. Whatever the exact figure, the direction is clear and consistent across sources: double-digit growth, driven by defence, offshore energy and survey.

Three features of this market shape the analysis that follows.

First, it is **concentrated in a small number of countries**. The leading manufacturers of AUVs, ROVs and survey USVs are headquartered in the United States, Norway, the United Kingdom, France, Germany, Japan, China and a handful of others. Countries with vast ocean territories but little industrial capacity in marine robotics will have to buy these capabilities or rent them as services, which raises questions of data ownership and sovereignty explored in Chapter 22.

Second, it is **dual use by nature**. A vehicle that maps a seabed for a wind farm can map it for a submarine; a vehicle that in-

spects a cable can cut one. Export controls, defence procurement and civilian markets are intertwined.

Third, it is **shifting from products to services**. Survey companies increasingly sell data rather than vessels: the customer buys square kilometres mapped or kilometres of pipeline inspected, delivered from a remote operations centre. This is the same shift that cloud computing brought to information technology, and it lowers the entry barrier for customers while concentrating capability in a few providers.

2.5 Why autonomy changes the economics, not just the technology

It is tempting to treat robotics as a set of tools that make existing ocean work somewhat cheaper. The models in this book suggest a more radical conclusion: autonomy changes which activities are economically possible at all.

Consider three examples developed in later chapters.

A deep-water seabed survey from a ship costs, in the Mapping Clock model, about 12.9 dollars per square kilometre at today's day rates and coverage rates. An uncrewed surface vessel costs about 6.2 dollars per square kilometre today and, with learning effects as fleets grow, perhaps half of that within a decade. At ship costs, finishing the global map is a multibillion-dollar international project requiring dozens of vessels for decades. At robotic costs, it becomes something that a coalition of governments and companies could plausibly complete within a decade (Chapter 4).

An offshore fish farm in exposed waters is uneconomic if it requires a permanent crew to feed fish, inspect nets and monitor health in rough seas. With computer vision feeding,

robotic net inspection and remote supervision, it becomes viable, and with it a source of food production that the models suggest could reach some 50 million tonnes a year by 2050 (Chapter 10).

A floating wind farm hundreds of kilometres offshore is expensive to maintain with crew transfer vessels that cannot sail in high waves. Robotic inspection and maintenance reduce downtime and costs, and the wind model shows a median reduction of about 11 percent in the cost of floating offshore electricity attributable to robotics by 2035 (Chapter 11).

In each case, the effect of the machine is not simply to cut cost at the margin. It is to move a whole category of activity from the impossible to the possible. That is what is meant by the machine turn.

2.6 The energy limit

The single greatest physical constraint on ocean robots is energy. An underwater vehicle must carry all the energy it uses; it cannot refuel at sea or draw power from the sun. The endurance of a battery-powered AUV is measured in hours to days, which is why the most enduring underwater platforms, gliders and floats, are those that use almost no energy at all, moving by buoyancy rather than propellers.

Engineers are attacking this limit in four ways.

Better batteries. Lithium-based batteries in pressure-tolerant housings have extended AUV missions considerably, and improvements in energy density continue. But batteries re-

main heavy, and each kilogram of battery is a kilogram that must be carried and buoyed.

Docking and recharging. Rather than carry energy for a whole mission, a vehicle can return to a docking station on the seabed or on a surface platform to recharge. This is the logic of resident systems. A docking station connected to an offshore wind farm, an oil field or a seabed cable can supply effectively unlimited energy to a robot that lives nearby.

Harvesting energy at the surface. Uncrewed surface vessels can use solar panels, wind and wave motion. The wind-driven survey vessels used in Atlantic and Pacific missions can stay at sea for months. Wave-powered gliders convert the motion of the sea into forward thrust.

Hybrid and fuel cell systems. Larger vehicles, especially military XLUUVs, use hybrid diesel-electric systems that snorkel to recharge or, in some designs, fuel cells that offer much greater energy density than batteries.

These developments change the economics of ocean work in a specific way. The robots that matter most for persistent presence are not those with the most impressive sensors but those that can stay longest without human intervention. The fall in costs described in Chapter 5 comes as much from endurance as from cheaper hardware.

2.7 The communications limit

The second constraint is communication. Below the surface, the main options are acoustic modems, which transmit data as sound at rates from a few hundred bits to a few tens of kilobits per second over kilometres; optical systems, which can transmit at high rates with blue-green light but only over

tens of metres in clear water; and physical connections through docking stations or tethers.

At the surface, satellite connectivity has improved dramatically. Low-Earth-orbit constellations now provide broadband links to vessels almost anywhere on the ocean, at costs that were unthinkable a decade ago. A surface vessel equipped with such a link can relay data from underwater vehicles to operators on land in near real time, and receive updated instructions. The combination of an uncrewed surface vessel with a satellite link above and an acoustic link below is becoming the standard architecture of remote ocean operations.

The practical consequence is a hierarchy of autonomy. Surface vessels, with good communications, can be supervised closely. Underwater vehicles near a surface relay can be supervised loosely. Underwater vehicles on long, deep missions far from any relay must be autonomous in the full sense. The deeper and further a task, the more intelligence must be placed on the machine itself.

Table 2.3. The communications hierarchy of ocean robots

Position	Link	Typical data rate	Degree of human supervision
Surface vessel	Satellite broadband	Megabits per second	Continuous remote monitoring
Underwater vehicle near a surface relay	Acoustic modem	Hundreds of bits to tens of kilobits per second	Periodic updates, mission changes
Underwater vehicle near a docking station	Optical or wired at dock	High, but only when docked	Data upload between missions
Underwater vehicle on a long deep mission	None or occasional acoustic	Negligible	Fully autonomous between surfacings

Source: Global Strategy Files.

2.8 Key findings of Chapter 2

- Human access to the deep ocean has passed from the age of the diver to the age of the vessel and is now entering the age of the autonomous machine.
- The family of ocean robots ranges from tethered ROVs to long-endurance gliders, uncrewed surface vessels, resident subsea robots and submarine-sized military drones.
- AI contributes perception, navigation, adaptive planning, fleet coordination, prediction and decision support. The operator to vehicle ratio is the key economic variable.
- The market for underwater vehicles is growing at double-digit rates, concentrated in a few countries, dual use by nature and shifting from products to services.
- Autonomy does more than lower costs: it moves whole categories of ocean activity from impossible to feasible.

End of the free excerpt

The complete book, with twenty-four chapters in eight parts, the nine modules of the GSF Ocean Frontier Model, thirty original figures and tables of results, the Blue Power Index for thirty countries with country profiles, three scenarios for 2040 and ten rules for governing the automated ocean, is available as a PDF at globalstrategyfiles.com.

THE AUTOMATED OCEAN

Seven tenths of our planet has stayed beyond human reach. Pressure, darkness, the radio silence of seawater and the cost of ships kept us on the surface. Artificial intelligence and robots are removing those walls. What will they let us take from the sea, and at what price?

Built on nine original models and 50,000 simulated futures, this book puts numbers on the automated ocean. At today's pace the seafloor map will not be finished this century; 250 survey drones would finish it by 2040 for under two billion dollars. AI that only makes fishing more efficient would leave fish stocks at 13 percent of their optimal level by 2050, unless 92 percent of fishing is monitored. Robotic offshore farms could raise blue food per person by up to half. A deep-sea nodule mine is a coin toss for its investors and an irreversible loss for the abyss. Forty cut cables would leave a region degraded for seven weeks.

With the Blue Power Index of thirty nations, three scenarios for 2040 and ten rules for a machine ocean.

ABOUT THE AUTHORS

Global Strategy Files is an independent analysis platform devoted to geopolitics, strategic foresight and the forces reshaping international security and the world economy. Its contributors combine expertise in international relations, defense, economics, energy, technology and quantitative modelling. *The Automated Ocean* is a collective work of the site's contributors.

globalstrategyfiles.com

GLOBAL STRATEGY FILES