



The Ronald O. Perelman Center for Political
Science and Economics (PCPSE)
133 South 36th Street
Philadelphia, PA 19104-6297

pier@econ.upenn.edu
<http://economics.sas.upenn.edu/pier>

PIER Working Paper 25-005

The (Un)Intended Consequences of Oil Sanctions Through the Dark Shipping of Sanctioned Oil

XIWEN BAI
Tsinghua University

JESÚS FERNÁNDEZ-VILLAYERDE
University of Pennsylvania

YILIANG LI
University of International
Business and Economics

LE XU
Shanghai Jiao Tong University

FRANCESCO ZANETTI
University of Oxford

February 10, 2025

The (Un)Intended Consequences of Oil Sanctions Through the Dark Shipping of Sanctioned Oil*

Xiwen Bai[†] Jesús Fernández-Villaverde[‡] Yiliang Li[§]
Le Xu[◇] Francesco Zanetti[¶]

February 10, 2025

Abstract

We examine the rise of dark shipping – oil tankers disabling AIS transceivers to evade detection – amid Western sanctions on Iran, Syria, North Korea, Venezuela, and Russia. Using a machine learning-based ship clustering model, we track dark-shipped crude oil trade flows worldwide and detect unauthorized ship-to-ship transfers. From 2017 to 2023, dark ships transported an estimated 7.8 million metric tons of crude oil monthly – 43% of global seaborne crude exports – with China absorbing 15%. These sanctioned flows offset recorded declines in global oil exports but create distinct economic shifts. The U.S., a net oil exporter, faces lower oil prices but benefits from cheaper Chinese imports, driving deflationary growth. The EU, a net importer, contends with rising energy costs yet gains from Chinese demand, fueling inflationary expansion. China, leveraging discounted oil, boosts industrial output, propagating global economic shocks. Our findings expose dark shipping’s central role in reshaping oil markets and macroeconomic dynamics.

JEL Classification: C32, C38, E32, Q43, R40.

Keywords: Dark shipping, oil sanction, satellite data, clustering analysis, LP.

*We are grateful to Hilde C. Bjørnland, Jianpeng Deng, Wei Li, Dongling Su, Pengfei Wang, Changhua Yu, and the participants at the AMLEDS webinar, as well as attendees at seminars held at Shanghai University of Finance and Economics and Peking University, for their valuable comments and suggestions.

† Bai: Tsinghua University, China. xiwenbai@mail.tsinghua.edu.cn. ‡ Fernández-Villaverde: University of Pennsylvania, U.S. jesusfv@econ.upenn.edu. § Li: University of International Business and Economics, China. yiliang_li@uibe.edu.cn. ◇ Xu: Shanghai Jiao Tong University, China. lexu1@sjtu.edu.cn. ¶ Zanetti: University of Oxford, U.K. francesco.zanetti@economics.ox.ac.uk.

1. Introduction

Dark shipping – the deliberate deactivation of AIS (Automatic Identification System) transceivers, which share a vessel’s location, heading, and speed to evade detection – has surged in recent years, largely due to Western sanctions restricting oil transport from Iran, Syria, North Korea, Venezuela, and Russia (Laudati and Pesaran, 2023; Triebert et al., 2024; Rodríguez, 2025). Designed to pressure sanctioned regimes by curbing their oil revenues (except in North Korea, where the goal is to restrict imports), these sanctions have led producers to offer premium rates to tanker operators transporting oil covertly, fueling the rise of dark shipping. Despite seaborne transportation accounting for over three-quarters of global crude oil (U.S. Energy Information Administration, 2024), the lack of systematic monitoring has hindered a comprehensive understanding of its impact on global trade, prices, and macroeconomic dynamics.

Detecting dark shipping is challenging. When a vessel disables its AIS transceiver, its location, heading, and speed become invisible to global tracking systems, creating data gaps that obscure its movements. Dark ships often exploit these gaps to transport oil from sanctioned countries or conduct unauthorized ship-to-ship transfers in maritime areas with weak regulatory oversight, such as international waters near Ceuta, off the coast of Gibraltar, and in West Africa (Lloyd’s List Intelligence, 2023; S&P Global Market Intelligence, 2023).¹ Consequently, identifying dark ships requires more than tracking AIS data; it demands a comprehensive approach that integrates vessel-specific suspicious characteristics, the likelihood of evasive activities during data gaps, and anomalies in navigational patterns.

While various attempts have been made to identify dark ships, many fall short. Existing methods often rely on incomplete data, lack clear criteria to distinguish dark shipping from routine maritime operations, or focus on individual ships or small vessel groups – failing to capture the larger tanker fleet.²

¹A ship-to-ship (STS) transfer is the process of transferring cargo, such as oil, between two vessels positioned alongside each other.

²The terms “dark ships” and “shadow fleet” are often used interchangeably but denote distinct phenomena. While “shadow fleet” typically refers to vessels bypassing the \$60 per barrel price cap on Russian oil exports (Levi et al., 2023; Johnson et al., 2023a; Cardoso et al., 2025), these ships operate within alternative systems without technically violating sanctions. In contrast, “dark ships” encompass a broader range of sanction evasion tactics, including vessels linked to Iran and Venezuela, which allow us to investigate a comprehensive set of sanctions beyond the case of Russia. These ships frequently breach sanctions through AIS manipulation or ship-to-ship transfers while still accessing G7/EU-related services. Differentiating these behaviors is essential for analyzing the impact of oil sanctions. Our investigation focuses on “dark ships.”

Additionally, some methods rely heavily on external sources, such as satellite imagery and proprietary data on ownership structures and insurance status, which may be inaccessible, unreliable, costly, or error-prone. The lack of fleet-wide approaches to identifying dark ships creates significant gaps in understanding the full scale of this covert activity and its broader economic and trade implications.

In this paper, we tackle these challenges by introducing a novel multi-attribute ship clustering model that employs a comprehensive three-level trip classification process to identify dark ships. Our model uniquely integrates a wide range of features, including vessel-specific characteristics, movement dynamics, and trip-level information from AIS data, enabling fleet-wide analysis of dark shipping activities. By doing so, we not only improve detection but also provide a framework to quantify the exports and imports of sanctioned oil and assess their aggregate implications.

A key contribution of our approach is the construction of two distinct measures that assess the likelihood of a ship either directly transporting oil from sanctioned countries or engaging in suspicious ship-to-ship transfers. The first measure focuses on a vessel’s proximity to ports in sanctioned countries during AIS gaps, while the second assesses the likelihood of dark ship-to-ship transfers by examining instances where two vessels are geographically close and have overlapping data gaps. For detecting dark ship-to-ship transfers, we propose a pioneering method that leverages geographical coordinates, heading, and timestamp information from historical AIS data, even when vessels have their AIS transceivers turned off. By analyzing vessel proximity and movement patterns during AIS gaps, our approach provides an alternative means of detecting unauthorized ship-to-ship transfers without relying on additional data sources such as satellite imagery (Rodger and Guida, 2022; Ballinger, 2024; Androjna et al., 2024).

Our identification results reveal significant differences in vessel-specific characteristics, travel patterns, and movement dynamics between tankers likely engaged in suspicious practices (“dark”) and those less likely (“white”). Compared to white ships, dark ships tend to be older, managed by operators with fewer vessels, and registered under obscure flags. Furthermore, dark ships exhibit higher levels of trip suspicion and irregular travel patterns, including more frequent idling, greater speed variability, and frequent detours from established routes. By our estimation, the dark ship fleet averaged 558 tankers during 2017–2023, comprising roughly one-fourth of the global crude oil tanker fleet. The fleet size and temporal dynamics of dark ships closely track major geopolitical and economic shifts over the sample period.

Thanks to our clustering model, we establish three empirical facts regarding oil sanctions and the dark shipping of sanctioned oil. First, from 2017 to 2023, an average of 7.8 million metric tons of crude oil per month were exported via dark ships from Iran, Syria, Venezuela, and Russia, accounting for approximately 43% of total world seaborne crude oil exports recorded in UN Comtrade data. Iran and Russia emerge as the primary contributors, with unauthorized oil shipments surging following key Western sanctions, such as the U.S. withdrawal from the Joint Comprehensive Plan of Action (JCPOA) in 2018 and the oil embargoes and price caps imposed on Russian oil in late 2022. Notably, after the G7 countries implemented a \$60 per barrel price cap on Russian crude oil exports in December 2022 (Hilgenstock et al., 2023; Johnson et al., 2023a,b; Kilian et al., 2024b), recorded global seaborne oil exports dropped sharply. However, sanctioned oil exports via dark shipping more than doubled, offsetting much of the reduction in supply and exerting downward pressure on global oil prices.³

Second, China is the largest importer of crude oil transported by dark ships from sanctioned countries. Between 2017 and 2023, China’s imports exceeded 97 million metric tons, accounting for 15% of total sanctioned oil exports. Other notable importers include South Korea, India, Egypt, Turkey, and Russia itself.⁴

Third, we analyze the broader economic effects of Western sanctions and the dark shipping of sanctioned oil on the U.S., the EU, and China using local projections (Jordà, 2005; Barnichon and Brownlees, 2019). Our findings show that rising intensity of oil sanctions reduces recorded world seaborne oil exports while simultaneously increasing exports of sanctioned oil via dark shipping. These shifts generate distinct economic consequences for the U.S. and EU, shaped by their oil trade positions and, crucially, their roles within the global supply chain involving China.

For the U.S., as a net oil exporter, the increased supply of sanctioned oil lowers oil prices, triggering a short-run contraction in production across both energy and non-energy sectors. However, the U.S. simultaneously benefits from lower import prices for goods from China – whose produc-

³Dark-shipped oil typically trades at a discount relative to global benchmarks. Since the onset of Western sanctions on Russian oil, the Urals free-on-board (FOB) price discount (relative to Brent crude) ranged between \$20 and \$35 per barrel for most of 2022 and early 2023, before narrowing later in 2023. The Eastern Siberia–Pacific Ocean (ESPO) FOB price discount has remained stable, ranging from \$5 to \$15 since mid-2022. Babina et al. (2023), Hilgenstock et al. (2023), and Kilian et al. (2024b) attribute these large discounts to oil market fragmentation caused by the embargo rather than the price cap. As Russia redirected crude exports to China and India, it had to lower FOB prices to compensate for higher shipping costs over longer distances. Moreover, as an oil export revenue-maximizer, Russia had no choice but to accept the lower prices demanded by India and China.

⁴Often, oil-exporting nations import oil, for instance, to maintain refinery operations at full capacity or because their refineries are optimized for slightly different crude varieties produced abroad.

tion costs are reduced by cheaper dark-shipped oil – fostering supply-driven growth characterized by deflationary pressure and output expansion. Conversely, for the EU, as a net oil importer, higher oil import prices due to sanctions suppress industrial production. Yet, rising demand for EU products from China – which expands production due to reduced oil costs – offsets this effect, leading to a net increase in output alongside demand-driven inflation in producer prices. For China, access to discounted oil substantially boosts industrial production, shaping economic dynamics in both the U.S. and the EU.

Importance of our results. Uncovering the previous facts is crucial, as policymakers must anticipate and mitigate the (un)intended consequences of oil sanctions and dark shipping if sanction regimes evolve. Without a thorough understanding of dark shipping and its broader implications, designing efficient and effective sanctions remains an impossible task.

Related literature. Our analysis is related to multiple strands of research. First, it builds on recent studies leveraging high-frequency satellite data to analyze global transport patterns, including port call behavior (Bai et al., 2022; Komaromi et al., 2022; Brancaccio et al., 2024), maritime traffic flows (Prochazka et al., 2019; Brancaccio et al., 2020; Cerdeiro et al., 2020; Li et al., 2022, 2023; Zhang et al., 2023), and responses to supply chain disruptions and geopolitical events (Bai et al., 2024; Notteboom et al., 2024). Additionally, recent studies have applied deep learning and satellite imagery to detect unauthorized ship-to-ship transfers. However, these approaches often face limitations due to incomplete AIS data or the costly requirements of proprietary satellite imagery (Rodger and Guida, 2022; Androjna et al., 2024; Ballinger, 2024; Bernabé et al., 2024). Compared to this literature, we develop a multi-attribute ship clustering model that relies solely on AIS data, offering a fleet-wide approach to detect dark ships and quantify sanctioned oil trade flows.

Second, our work connects to the literature examining the interactions between the oil industry and macroeconomic dynamics, employing structural vector autoregressions (SVARs) to estimate the impact of demand and supply shocks on oil prices and the real economy (Kilian, 2009, 2014; Baumeister and Kilian, 2016; Baumeister and Hamilton, 2019; Känzig, 2021; Kilian et al., 2024b; Verduzco-Bustos and Zanetti, 2025), microdata to explore industry developments and their aggregate implications (Asker et al., 2019; Bjørnland et al., 2021; Bornstein et al., 2023), and structural models to disentangle the channels through which oil shocks propagate in the economy (Boden-

stein et al., 2011; Lippi and Nobili, 2012; Melek et al., 2017; Kilian et al., 2024a). Our contribution to this literature is the quantification of sanctioned oil exports transported by dark ships and the identification of the dynamic causal effects of dark shipping on prices and production in the U.S. and EU economies.

Third, we contribute to the literature on international trade and macroeconomic dynamics triggered by sanctions, particularly those imposed on Russia following its invasion of Ukraine (Itskhoki and Mukhin, 2022; Laudati and Pesaran, 2023; Sturm, 2023; de Souza et al., 2024; Ghironi et al., 2024; Li et al., 2024b).⁵ While existing studies acknowledge the role of dark shipping in sustaining Russia’s economy by preserving its oil revenue despite Western sanctions (Babina et al., 2023; Hilgenstock et al., 2023; Johnson et al., 2023a,b; Kilian et al., 2024b; Cardoso et al., 2025), the opacity of sanctioned oil trade has hindered quantitative assessments of its impact on global oil trade, prices, and macroeconomic conditions. Our contribution is to develop a scalable framework that not only detects dark shipping across fleets but also quantifies the exports and imports of dark-shipped oil, providing policymakers with a tool to assess the economic ramifications of sanctions and optimize sanction designs.

The remainder of the paper is organized as follows: Section 2 introduces our multi-attribute ship clustering model; Section 3 presents the identification results; Section 4 examines the impact of oil sanctions and dark shipping on the global oil market, while Section 5 analyzes their macroeconomic effects on the U.S. and EU economies. Finally, Section 6 concludes. An online appendix provides the full algorithm description, pseudocode, estimation procedures for local projections, and additional figures and tables.

2. A Clustering Model for Dark Ship Identification

In this section, we develop a multi-attribute ship clustering model to identify vessels likely involved in transporting sanctioned oil from Iran, Syria, North Korea, Venezuela, and Russia between 2017 and 2023. We begin by providing background information on the dark shipping industry and the challenges of vessel identification. Next, we introduce the satellite data used in our analysis. The core of this section explains the motivation for applying an unsupervised learning-based clustering

⁵Felbermayr et al. (2024) review recent research on the economics of sanctions, presenting both stylized facts derived from the Global Sanctions Database (GSDB) and quantitative evidence on the effects of sanctions on economic outcomes.

model to distinguish dark ships from white ships, followed by a brief description of the algorithm.

2.1. Background

In recent years, the dark shipping industry has become a covert means of evading Western oil-related sanctions on Iran, Syria, North Korea, Venezuela, and Russia. This issue has drawn significant attention from both the media and policymakers ([U.S. Department of State, 2020](#); [Triebert et al., 2023, 2024](#); [European Parliament, 2024](#)). However, a key challenge in assessing the scope and impact of dark shipping is the lack of reliable oil export data from sanctioned nations. For instance, Russia ceased publishing detailed oil export data after its 2022 invasion of Ukraine. Furthermore, countries importing sanctioned oil often withhold accurate import data due to the covert nature of these transactions.

Identifying vessels engaged in the dark shipping of sanctioned oil poses significant challenges. These ships do not report their covert activities, and before advanced satellite-based AIS data were available (to be described soon), real-time monitoring was nearly impossible. Even with AIS data, dark ships frequently disable their AIS transceivers, obscuring their location, heading, and speed from global tracking systems. These frequent and unexplained data gaps, combined with intentional practices such as AIS manipulation or “spoofing,” make it difficult to track their movements and identify violations.

For example, the anonymously owned crude oil tanker *Roma* (IMO: 9182291), which we will later classify as a dark ship, exhibited frequent AIS transmission gaps throughout 2023. During this period, it primarily operated in international waters between the Strait of Hormuz near Iran and Southeast Asia, with over 40% of its recorded time lacking AIS observations.⁶

So far, existing research has relied on alternative data sources, including satellite imagery ([Androjna et al., 2024](#); [Ballinger, 2024](#); [Bernabé et al., 2024](#)) and proprietary information on vessel ownership and insurance, to identify individual vessels or small groups engaged in dark shipping. For instance, anonymous ownership or opaque corporate structures often signal potential dark shipping activity ([S&P Global Market Intelligence, 2023](#)). Similarly, the absence of known protection and indemnity (P&I) insurance or coverage outside the International Group of 12 P&I

⁶A New York Times report by [Triebert et al. \(2023\)](#) documents an instance of AIS spoofing involving the crude oil tanker *Cathay Phoenix* (IMO: 9249324), which our method also classifies as a dark ship. In February 2023, the vessel transmitted a false location signal west of Japan while actually loading oil 250 miles north at Russia’s Kozmino port. This journey likely violated U.S. oil sanctions on Russia.

clubs (e.g., the American Club), which insures 95% of the global tanker fleet, serves as another red flag (Levi et al., 2023; Lloyd’s List Intelligence, 2023). However, these external data sources are often restricted, costly, and unavailable to the public. Moreover, while identifying individual dark ships is useful, it does not reveal the broader scale of the clandestine oil trade, thereby limiting its effectiveness in assessing dark shipping’s aggregate economic impact on the U.S. and EU economies.

2.2. AIS Data

To construct global indicators for dark shipping, we use satellite-based AIS data, comprising approximately 330 million crude oil tanker records. From this dataset, we curate an annual sample of around 2,160 tankers for detailed analysis.⁷ Each AIS entry includes the International Maritime Organization (IMO) ship identification number, timestamp, draft, speed, heading, and geographical coordinates. The AIS system processes over 2,000 reports per minute, updating as frequently as every two seconds, providing extensive global coverage of crude oil tanker movements from 2017 to 2023.

These data, decoded from satellite tracking systems installed on vessels, apply to tankers over 300 gross tons engaged in international voyages, as mandated by the IMO (Heiland et al., 2019; Bai et al., 2024). The recorded position, speed, and heading allow us to track ship movements and detect when vessels deactivate their AIS transceivers – key indicators of potential dark shipping, including the transport of sanctioned oil or ship-to-ship transfers conducted without AIS signals.

2.3. The Challenges of Labeling Ships

Identifying dark ships is a labeling problem: Is ship x a dark or white ship? (Xu and Tian, 2015; Sarker, 2021). The sheer volume of AIS records – numbering in the millions – naturally lends itself to the application of machine learning methods to automate this labeling.

⁷About one-third of these tankers are Very Large Crude Carriers (VLCCs), while the rest are classified as Aframax, Suezmax, Dirty Panamax, and others. We focus on the dark shipping of crude oil rather than petroleum products because the latter are diverse, making the aggregation of transportation volume difficult when deriving trade flows for a single product. Furthermore, unlike crude oil tankers, which can only transport crude oil, petroleum products require more specialized vessels (e.g., those with well-maintained coatings). However, these vessels can opportunistically switch to carrying crude oil (Li et al., 2022), further complicating the task of calculating trade flows. Consequently, despite the frequent involvement of dark ships in transporting sanctioned petroleum products, we leave their identification for future research.

When designing such a machine learning algorithm, we face the challenge that no training set exists for classifying dark ships. By nature, these vessels conceal their identities, and commercial entities do not disclose their identification results, which are based on proprietary, opaque, or error-prone data. Additionally, dark ships do not always exhibit suspicious behavior; they can temporarily operate as “white ships,” minimizing detection risk during periods of heightened scrutiny. These complexities render supervised and semi-supervised models unsuitable.

Lastly, our objective extends beyond identifying specific malign vessels. We aim to label the global fleet of dark ships and quantify the exports and imports of sanctioned oil based on the trips undertaken by these vessels.

Thus, we need a computationally efficient algorithm capable of grouping data without labeled examples. As we will see, K -means clustering is well-suited for this task due to its simplicity and effectiveness (MacQueen, 1967; Lloyd, 1982; Wu, 2012). This approach relies solely on AIS data for detection, avoiding the computational and financial burdens of satellite imagery analyses. Appendix A provides a detailed explanation of how we build our own centroid-based ship clustering algorithm, but we summarize the key ideas in the next subsection.

2.4. Overview of the Algorithm

Our point of departure is the domain knowledge that dark ships often exhibit similar behavioral patterns at a granular level. More concretely, we follow the [U.S. Department of State \(2020\)](#), which identifies two widely recognized indicators of deceptive shipping practices. The first is disabling or manipulating AIS transceivers to obscure movements. The second is engaging in ship-to-ship transfers, especially in high-risk areas for sanction violations or unauthorized activities.

With this domain knowledge, we design a three-level trip classification process that builds on existing methods in the literature while introducing novel features and clustering steps. Specifically, we integrate (i) direct involvement in shipping sanctioned oil (e.g., calling at sanctioned-country ports), (ii) AIS data gaps that may hide suspicious port visits or ship-to-ship transfers, and (iii) anomalies in kinematic movements (such as high speed variability or frequent detours).⁸ By incorporating these three levels, we go beyond the common practice of estimating the proba-

⁸This multi-level approach ensures that simply calling at a sanctioned port does not automatically render a vessel “dark.” Instead, we distinguish between vessels that may legitimately carry sanctioned cargo (e.g., at or below the G7 price cap for Russian crude) and those exhibiting deceptive or evasive practices, such as turning off AIS when near sanctioned ports or engaging in suspicious ship-to-ship transfers.

bility that a vessel is “dark” based solely on suspicious characteristics (e.g., old age or registration under a high-risk flag of convenience).⁹

Level 1 (sanctioned port calls). In the first step, we identify trips that start or end at a port in any of the five sanctioned countries (Iran, Syria, North Korea, Venezuela, and Russia) after sanctions have been imposed on that country. These trips are flagged as high risk, reflecting their direct involvement in shipping sanctioned oil.


Level 2 (AIS data gaps). Next, we examine prolonged gaps in a vessel’s AIS transmissions – defined here as gaps beyond the vessel-specific 99th percentile of time differences between consecutive AIS signals – and analyze their start and end locations.

One type of suspicious trip occurs when a vessel is geographically close to a port in one of the five sanctioned countries, allowing it to visit during a prolonged AIS data gap. This scenario is illustrated generically in Panel 1a of Figure 1 and with real-life satellite imagery in Panel 1b. The latter shows that *Roma* (IMO: 9182291) – previously identified as “dark” using our algorithm – was loading crude oil from Kharg Island in Iran in August 2022 while its AIS transceiver was turned off.¹⁰


The second type of suspicious trip involves potential unauthorized ship-to-ship transfers. We identify such trips by analyzing the geographical coordinates and timestamps of two vessels before and after their respective data gaps, particularly when these gaps overlap. This case is depicted in Panel 1c and illustrated with real-life satellite imagery in Panel 1d. The latter reveals that *Abyss* (IMO: 9157765) engaged in an unauthorized ship-to-ship transfer with *Shanaye Queen* (IMO: 9242118) in January 2022 near Bahrain. Both vessels were identified as dark ships by our algorithm and were alleged to have actively participated in loading and transferring sanctioned Iranian oil (Meade, 2023; Lloyd’s List Intelligence, 2024b).

⁹According to [Lloyd’s List Intelligence \(2023\)](#), vessels engaged in dark shipping are often older and registered under high-risk flags of convenience. A flag state defines the jurisdiction under which a vessel is registered or licensed, akin to its nationality. While on the high seas, vessels are subject to international law and the regulations of their flag state. Dark ships frequently choose flags of convenience to evade stricter oversight. For example, Panama, Liberia, and the Marshall Islands – three major flags of convenience – are the most represented in the global fleet. In 2022, these three flag states accounted for 44.3% of the world’s cargo capacity ([UNCTAD, 2022](#)).

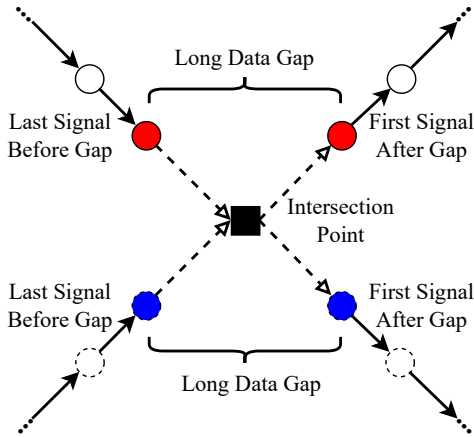
¹⁰For this figure, our algorithm first identifies *Roma* as a dark ship in 2022, after which we search for real-life satellite imagery to verify this classification.




(a) Direct Port Visit: Mechanism



(b) Direct Port Visit: Satellite Imagery



(c) Ship-To-Ship Transfer: Mechanism



(d) Ship-To-Ship Transfer: Satellite Imagery

Figure 1: Deceptive Shipping Practices During Prolonged AIS Data Gaps

Notes. Panel 1a illustrates a vessel returning to a nearby port in a sanctioned country during a prolonged AIS data gap. Panel 1b presents satellite imagery of such a visit by the identified dark ship *Roma* (IMO: 9182291), with dashed white arrows indicating the vessel's trajectory during the data gap. Panel 1c shows a ship-to-ship transfer in which both vessels disable their AIS transceivers, creating overlapping data gaps. The intersection point is estimated using the last known coordinates and headings before the AIS signals disappear. Panel 1d provides satellite imagery of a dark ship-to-ship transfer involving two identified dark ships, *Abyss* (IMO: 9157765) and *Shanaye Queen* (IMO: 9242118), with dashed white arrows representing the trajectories of both vessels.

Level 3 (kinematic anomalies). In the last step, we measure speed and route irregularities during each trip by calculating the vessel’s average speed, the standard deviation of its speed, and a detour factor that compares its actual route to the direct route. Trips with unusually low speeds, high speed variability, or circuitous detours receive higher suspicion scores, as these patterns often reflect deceptive maneuvers or potential spoofing of AIS signals.¹¹

After classifying each trip across three levels, we assign a categorical suspicion score (0, 0.5, or 1) based on whether it triggers flags for sanctioned port calls, AIS gaps, and/or kinematic anomalies. A trip receives a score of 1 if it starts or ends in a sanctioned port (Level 1). Otherwise, if both Levels 2 and 3 indicate suspicious behavior, we assign 1; if only one level does, we assign 0.5; if neither does, we assign 0. We then compute each vessel’s average trip suspicion score by taking the mean of its trip-level scores.

Next, we combine this average trip suspicion score with four additional vessel-level attributes: (i) vessel age (older vessels pose lower financial risks if seized), (ii) operator size (smaller operators face less scrutiny and have limited compliance resources), (iii) the flag state risk ranking from the annually updated [Paris MoU](#) list, and (iv) the idle trip ratio, defined as the fraction of trips where a vessel operates at 1 knot or less for 14 consecutive days.¹²

Finally, we input these vessel-level metrics into a K -means clustering model, which classifies vessels as dark (high suspicion) or white (low suspicion) each year.

3. Identification Results

In this section, we present the identification results of our ship clustering model, starting with an analysis of differences in vessel-specific characteristics, travel patterns, and movement dynamics between dark and white ships. We then examine the fleet size and composition of dark ships and how these evolve over the sample period. Next, we assess the model’s predictive performance by comparing its results to external dark ship classifications and an alternative method based solely on vessel-specific characteristics. Our findings show that incorporating trip-level AIS data

¹¹See [Triebert et al. \(2023\)](#) and [Lloyd’s List Intelligence \(2024a\)](#) for examples of AIS spoofing, where transmitted positions or speeds do not match actual movements.

¹²Larger commercial operators generally have stronger compliance departments and face higher reputational costs if involved in sanction violations, making smaller operators more prone to unauthorized activities. The Paris Memorandum of Understanding (MoU) on Port State Control ranks flag states into White, Gray, and Black lists based on vessel inspections and detentions, provided the flag state has undergone at least 30 inspections in the past three years. The Black list represents the highest risk.

significantly improves the model’s accuracy. Finally, we analyze the geographical concentration of dark shipping activities, demonstrating that our results align with externally documented locations and capture temporal shifts in sanction regimes. Overall, these findings highlight the model’s effectiveness in identifying dark ships.

3.1. Fleet Characteristics and Temporal Dynamics of Dark Ships

Our first check examines whether the dark and white ships identified by our clustering model exhibit distinct vessel-specific characteristics, travel patterns, and movement dynamics from 2017 to 2023.¹³ Across Panels 2a to 2e of Figure 2, we find that, on average, dark ships are older, operated by entities with fewer vessels, flagged by higher-risk states, and have higher trip suspicion scores and idle trip ratios. The differences are statistically significant for most features, though some show no significant variation in specific years.¹⁴ Since these features align with prior studies, our check validates our approach.


Figure 3a presents our estimate of the dark ship fleet size, shown both in absolute numbers and as a percentage of the global crude oil tanker fleet from 2017 to 2023. The fleet grew notably from 2017 to 2018, peaking at 562 vessels in 2018. This rise coincided with the U.S. withdrawal from the JCPOA and the reinstatement of all nuclear-related sanctions on Iran, including those targeting its oil exports (U.S. Department of State, 2018; Laudati and Pesaran, 2023). Amid rising crude prices, these sanctions created a lucrative market for unauthorized Iranian oil trade, prompting many tankers to go dark in pursuit of higher profits.

This elevated level of dark shipping persisted into 2019 as the U.S. imposed further sanctions on Venezuela, specifically targeting its state-owned oil company, *Petróleos de Venezuela, S.A.* (PDVSA), to pressure the Venezuelan government by cutting off its primary revenue source: crude oil exports (U.S. Department of the Treasury, 2019).


In 2020, the COVID-19 pandemic caused a global economic contraction, reducing oil demand and weakening the oil shipping market (Nagle, 2020; Bai et al., 2024). With lower profitability, fewer oil tankers engaged in dark shipping. However, this decline was short-lived, as the 2021 economic reopening, particularly in China, restored oil demand and spurred a rebound in dark shipping activity.

¹³This comparison also evaluates whether the clusters produced by the *K*-means model in Section 2.4 are well-separated, ensuring within-cluster homogeneity and between-cluster heterogeneity.


¹⁴For instance, the idle trip ratio for dark ships is only marginally higher than that for white ships in 2018.




(a) Years of Usage




(b) Number of Vessels Owned by Operator



(c) Flag State Ranking



(d) Average Trip Suspicion Score



(e) Idle Trip Ratio

Figure 2: White vs. Dark Ships

Notes. The figures show the average values for usage years, the number of vessels owned by the ship’s commercial operator, flag state ranking, trip suspicion score, and idle trip ratio for both white and dark ships, as identified using the K -means clustering model in Section 2.4.




Figure 3: Overview of the Dark Ship Fleet

Notes. Figure 3a shows the number and percentage of dark ships within the global crude oil tanker fleet from 2017 to 2023. Figure 3b presents the frequency with which a tanker is identified as a dark ship. This frequency is calculated by dividing the number of years a ship is classified as dark by the total years it has available AIS data. A frequency of 1 means the ship is identified as dark every year it has operational data, while a frequency of 0 indicates it is never classified as dark during the sample period.

By 2022, the fleet size declined again, driven by a global slowdown caused by contractionary monetary policies aimed at curbing inflation and the lingering effects of China’s zero-COVID policy (Bai et al., 2025). However, at the end of 2022 and throughout 2023, sanctions on Russia following its invasion of Ukraine – including oil price caps and export bans – created strong incentives for dark shipping (U.S. Department of the Treasury, 2023). Buyers from non-sanctioning countries, such as China and India, continued purchasing Russian oil, driving another surge in the dark fleet, which reached 566 vessels by the end of 2023.

Despite fluctuations in the absolute size of the dark ship fleet, its share of the global oil tanker fleet steadily declined from 29.8% in 2017 to 23.8% in 2023. This decline was largely driven by the rapid addition of new white tankers (Kalouptsi, 2014; Greenwood and Hanson, 2015; Dunn and Leibovici, 2023), the slow scrapping of older vessels, and heightened scrutiny of illegal oil shipping, which reduced incentives for ships to turn dark.

Figure 3b further corroborates this trend, showing how frequently a tanker was classified as a dark ship during its AIS operational years. Our findings indicate that 68.7% of tankers in our sample were never classified as dark ships. In contrast, 21.0% were consistently identified as dark ships across all available years, while 10.3% were classified as dark ships only in certain years.

This suggests that the dark ship fleet remained relatively stable, with most dark ships consistently identified over time and few white ships turning dark.

3.2. Evaluation of Model Performance

Next, we assess our model’s predictive performance by comparing its identification of dark ships with classifications from external sources across different years. As shown in Table 1, our model correctly identifies most dark ships during the years they are classified as such, achieving an 83.3% prediction accuracy. This result demonstrates that vessel-specific characteristics, combined with AIS-derived travel patterns and movement dynamics, can reliably identify dark ships without relying on private or costly data, such as insurance status (Lloyd’s List Intelligence, 2023), ownership structures (S&P Global Market Intelligence, 2023), or satellite imagery (Triebert et al., 2023).

We also compare our model’s predictive performance against a naive model that relies solely on vessel-specific characteristics – years of usage, the number of vessels owned by the ship’s commercial operator, and its flag state ranking – to identify dark ships. Comparing the last two columns in Table 1, it is clear that incorporating trip-level AIS data improves identification accuracy by over 28%.

Additionally, we validate our approach using the Technique for Order Preference by Similarity to Ideal Solution (TOPSIS), developed by Hwang and Yoon (1981), to rank oil tankers based on the same features as our clustering model. This method provides a more granular assessment of each vessel’s likelihood of engaging in dark shipping. Rankings are available upon request, but they strongly support our findings.

3.3. Geographical Concentration of Dark Shipping

Next, we examine the geographical concentration of dark shipping activities. This analysis provides valuable insights for academics, policymakers, and industry stakeholders in detecting and mitigating dark shipping. Specifically, we focus on locations where vessels frequently disable their AIS transceivers to transport oil for sanctioned countries or conduct suspicious ship-to-ship transfers.

First, we identify geographical hotspots where dark ships transport oil for sanctioned countries by visiting their ports during data gaps. By collecting the last recorded data points before vessels

Table 1: Evaluation of Model Performance Against External Sources and A Naive Model

IMO	Source	Suspicion Years	Model Predictions	
			Full Model	Vessel-Specific Features Only
9182291	Lloyd’s List Intelligence (2023)	2022, 2023	2022, 2023	2022, 2023
9224271	Lloyd’s List Intelligence (2023)	2022, 2023	2022, 2023	N/A
9176993	Lloyd’s List Intelligence (2023)	2021, 2022, 2023	2021	2021
9131357	Lloyd’s List Intelligence (2023)	2022, 2023	2022, 2023	2022, 2023
9183295	Lloyd’s List Intelligence (2023)	2021, 2022, 2023	2021, 2022, 2023	2021, 2022, 2023
9258521	Lloyd’s List Intelligence (2023)	2022, 2023	2022, 2023	2022, 2023
9230969	Lloyd’s List Intelligence (2023)	2022, 2023	2022, 2023	N/A
9307645	Lloyd’s List Intelligence (2024a)	2023	2023	2023
9242118	Lloyd’s List Intelligence (2024b)	2021, 2023	2021, 2023	N/A
9251274	S&P Global Market Intelligence (2023)	2022, 2023	2022, 2023	2022, 2023
9310147	S&P Global Market Intelligence (2023)	2022, 2023	2022, 2023	2022, 2023
9410870	S&P Global Market Intelligence (2023)	2022, 2023	N/A	N/A
9322827	S&P Global Market Intelligence (2023)	2021, 2022, 2023	2022	N/A
9235892	Triebert et al. (2023)	2023	2023	N/A
9249324	Triebert et al. (2023)	2023	2023	N/A
9259745	Triebert et al. (2023)	2023	2023	2023
9259733	Triebert et al. (2023)	2023	2023	2023
9157765	Meade (2023)	2021, 2022, 2023	2021, 2022, 2023	2021, 2022, 2023
9194127	Wijaya (2022)	2022	N/A	N/A
9304667	Wijaya (2022)	2022	2022	2022
9179701	Wijaya (2022)	2022	2022	2022
9255880	Wijaya (2022)	2022	2022	2022
9236353	Wijaya (2022)	2022	2022	N/A
9231767	Wijaya (2022)	2022	2022	N/A
9233208	Wijaya (2022)	2022	2022	N/A
			Summary	
			Total (Ship, Year)	42
			Identified (Ship, Year)	35 (83.3%)
			Unidentified (Ship Year)	7 (16.7%)
				23 (54.8%)
				19 (45.2%)

Notes. The vessel with IMO number 9176993 lacks AIS data for 2022 and 2023, years in which external sources classify it as a dark ship. “N/A” indicates cases where the vessel is not identified as dark in years when external sources classify it as such. The vessel-specific features considered include years of usage, the number of vessels owned by its commercial operator, and its flag state ranking.

disable their AIS transceivers, along with the associated port-based trip suspicion scores, we visualize areas of heightened suspicious activity through mapping.¹⁵

The estimated dark shipping hotspots align with locations frequently reported by external sources, and their variations over time correspond to changes in sanction regimes. For instance,

¹⁵For details on the calculation of the port-based trip suspicion score, refer to Algorithm A.1 in Appendix A.3.

Figure 4 shows dark shipping activity near Venezuela from 2017 to 2019. Both the number of trips with AIS gaps (colored points) and the associated suspicion scores (chromatic intensity) increased sharply in 2019, coinciding with U.S. sanctions on PDVSA, which severely restricted Venezuela’s oil exports. This forced Venezuela to seek alternative buyers, such as China and India, often resorting to dark shipping. This pattern persisted through 2023, as shown in Figure B.1 in Appendix B.

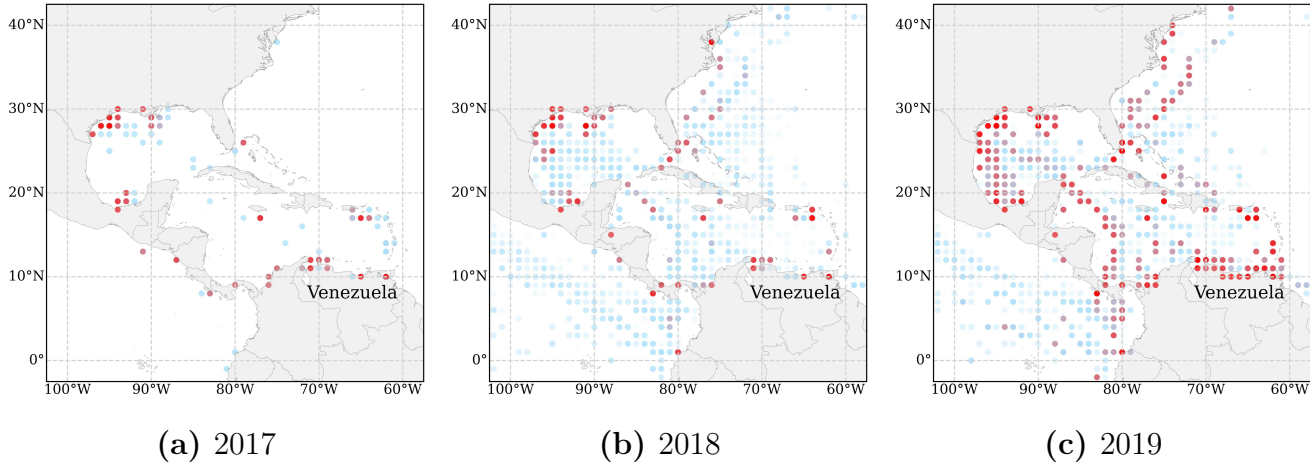


Figure 4: Geographical Concentration of Dark Shipping Activities Near Venezuela: 2017–2019

Notes. Geographical concentration of dark shipping activities near Venezuela from 2017 to 2019. Each colored point represents a geographical square (with 1-degree variations in longitude and latitude) that contains at least one last recorded data point before a vessel disables its AIS transceiver. The chromatic intensity of each point reflects the sum of port-based trip suspicion scores associated with these last recorded data points. The port-based trip suspicion score estimates the likelihood that a vessel traveled to a nearby sanctioned-country port during its AIS gap.

Figure 5 illustrates dark shipping activity near Iran from 2021 to 2023. The number of trips with AIS gaps declined substantially after 2021 (as indicated by the lower density of colored points), though suspicion scores remained high, particularly near the Strait of Hormuz.¹⁶

This decline aligns with intensified Western sanctions on Russia. As noted by [Cook and Sheppard \(2023\)](#) and [S&P Global Market Intelligence \(2023\)](#), tankers previously transporting Iranian oil shifted to Russian oil, as Russia offered premium rates to shield its export revenue from Western measures such as the G7/EU oil price cap. Given that the dark ship fleet size remained relatively stable (Figure 3a) while dark shipping intensified near the Black Sea and Baltic Sea (as discussed next), our findings highlight the geographical shift in dark shipping activity in response to changes in sanctioned regimes.

¹⁶Figure B.2 in Appendix B also presents dark shipping activity near Iran from 2017 to 2020.

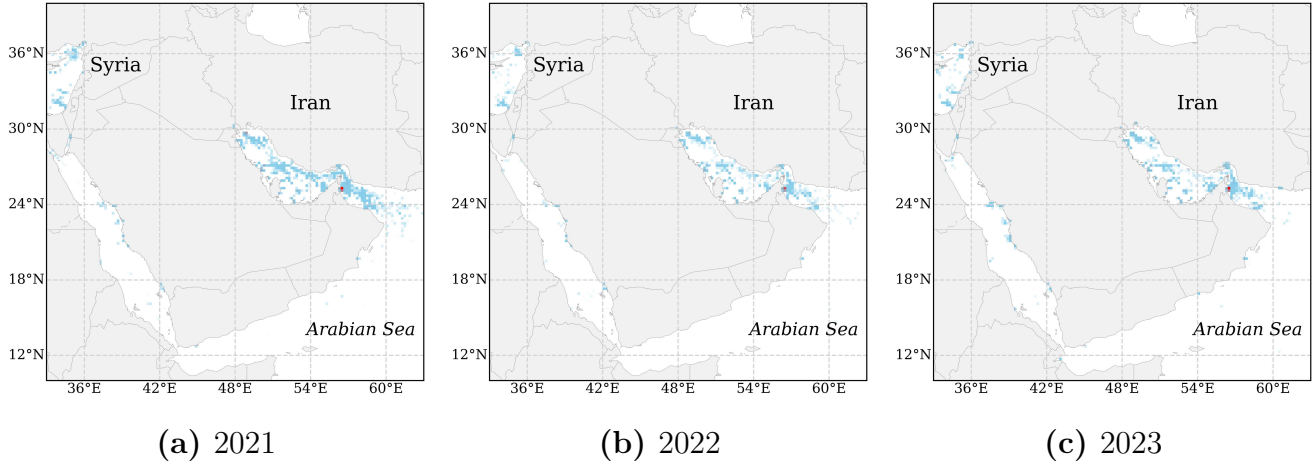


Figure 5: Geographical Concentration of Dark Shipping Activities Near Iran: 2021–2023

Notes. Geographical concentration of dark shipping activities near Iran from 2021 to 2023. Each colored point represents a geographical square (0.2-degree variations in longitude and latitude) containing at least one last recorded data point before a vessel disables its AIS transceiver. The chromatic intensity indicates the cumulative port-based trip suspicion scores associated with these last recorded data points within each square.

Next, we analyze dark shipping related to Russian oil by examining activity intensity near the Black Sea and Baltic Sea. Instead of aggregating data points into small geographical squares, we categorize trips by AIS gap length and plot linear connections between the data points recorded before and after each AIS gap. This approach offers a clearer visualization of dark ship movements, allowing us to trace specific routes and better understand evasive maneuvers – insights that would be lost with simple data point aggregation.

Figure 6 illustrates the concentration of direct vessel visits to suspicious Russian ports near the Black Sea from 2021 to 2023.¹⁷ The top row represents short AIS gaps (less than 10 hours), the middle row medium gaps (10–120 hours), and the bottom row long gaps (over 120 hours). Chromatic intensity indicates the port-based trip suspicion score for each AIS gap.

Following Western sanctions on Russian oil in late 2022, both the scale and intensity of dark shipping – measured by line density and chromatic intensity – increased substantially near the Russian ports of Novorossiysk, Taman, and Tuapse, especially for short and medium AIS gaps. For long AIS gaps, intensity remained stable, but activity expanded, as reflected in the higher line density in later years compared to 2021.¹⁸

¹⁷Figures B.3 and B.4 in Appendix B show visits to Russian ports near the Black Sea and Baltic Sea from 2017 to 2020, respectively.

¹⁸By construction, trips with longer AIS gaps tend to have higher port-based trip suspicion scores, leading to greater chromatic intensity. Given a fixed geographical distance, a longer AIS gap reduces the required average speed to reach a suspicious port. Since the port-based trip suspicion score in Algorithm A.1 of Appendix A is

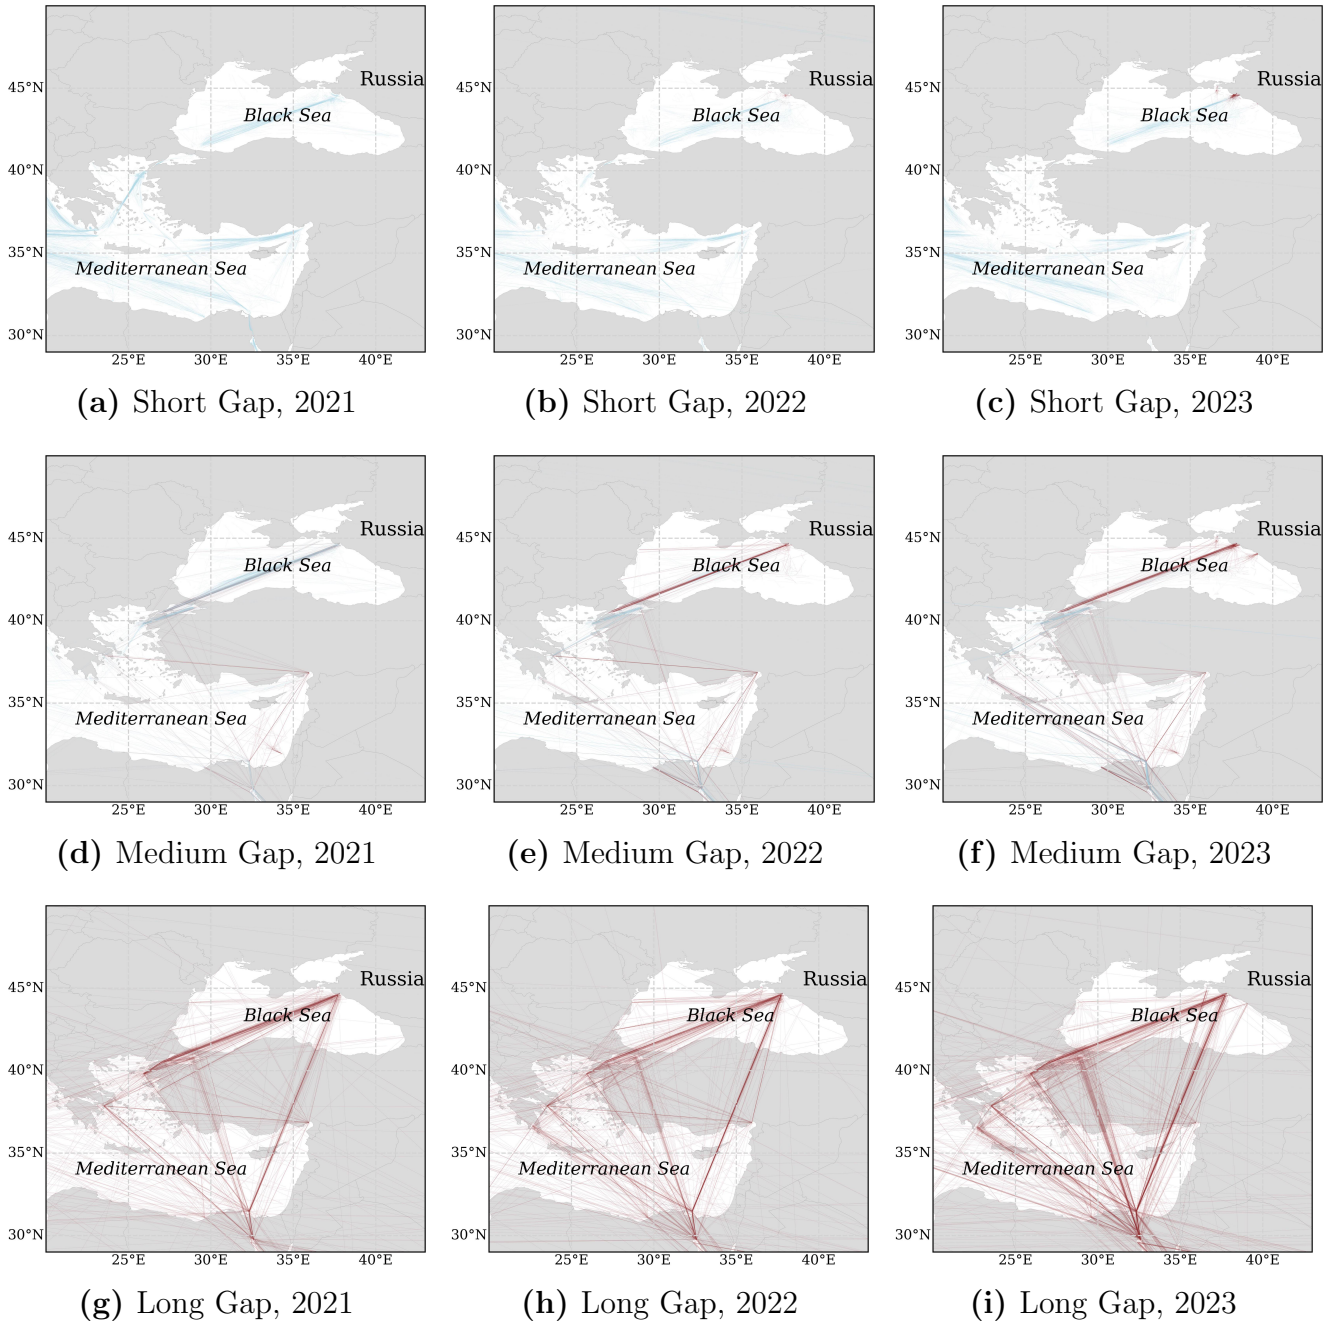


Figure 6: Geographical Concentration of Dark Shipping Activities in the Black Sea: 2021–2023

Notes. Geographical concentration of dark shipping activities in the Black Sea near Russia from 2021 to 2023. Each line represents a linear connection between the two data points recorded before and after an AIS gap during a trip, with chromatic intensity indicating the port-based trip suspicion score. To highlight differences in vessel travel patterns based on AIS gap length, we classify gaps into short, medium, and long categories, using 10 and 120 hours as cutoff points.

Figure 7 illustrates the concentration of tanker visits to Russian ports in the Baltic Sea from 2021 to 2023. Similar to the Black Sea, both the scale and intensity of dark shipping activity have inversely related to speed, a lower speed results in a higher suspicion score.

increased significantly near the Russian ports of Primorsk, Ust-Luga, Vysotsk, and St. Petersburg, particularly for medium and long AIS gaps.




Figure 7: Geographical Concentration of Dark Shipping Activities in the Baltic Sea: 2021–2023

Notes. Geographical concentration of dark shipping activities in the Baltic Sea near Russia from 2021 to 2023. Colored lines and chromatic intensity are represented in the same manner as in Figure 6.

The heightened concentration of dark shipping activities near Russian ports in the Black Sea and Baltic Sea – when AIS transceivers were turned off – illustrates the unintended consequences

of Western sanctions on Russian oil. Originally intended to curb Russia’s oil revenue, these measures have instead incentivized vessels to evade detection through longer AIS gaps and more frequent direct visits to sanctioned ports. Consequently, dark shipping has not only persisted but expanded.




Figure 8: Geographical Concentration of Dark Shipping Activities Near the Strait of Gibraltar and the Gulf of Guinea: 2021–2023

Notes. Geographical concentration of dark shipping activities near the Strait of Gibraltar and the Gulf of Guinea from 2021 to 2023. Each colored point represents a 1-degree geographical square containing at least one last recorded data point before a vessel disables its AIS transceiver. Chromatic intensity indicates the cumulative ship-to-ship transfer-based trip suspicion scores associated with these last recorded data points. The ship-to-ship transfer-based trip suspicion score estimates the likelihood of a transfer occurring when two vessels are geographically close and have overlapping AIS gaps.

Finally, we identify geographical hotspots for suspicious ship-to-ship transfers by aggregating the last recorded data points before tankers disable their AIS transceivers into small geographical squares. We then sum the corresponding ship-to-ship transfer-based suspicion scores, with chromatic intensity representing the overall suspicion level in each square.¹⁹

Figure 8 illustrates suspicious ship-to-ship transfers near the Strait of Gibraltar and the Gulf of Guinea from 2021 to 2023.²⁰ Our findings align closely with external reports, including [Lloyd’s List Intelligence \(2023\)](#) and [S&P Global Market Intelligence \(2023\)](#), which highlight Ceuta near Gibraltar and West Africa as high-risk areas for ship-to-ship transfers involving U.S.-sanctioned

¹⁹For details on the calculation of the ship-to-ship transfer-based trip suspicion score, see Appendix A.3.

²⁰As shown in Figure B.5 in Appendix B, these hotspots have existed since 2017. Their persistence suggests that despite increasing sanctions and enforcement efforts, these regions remain critical for unauthorized transfers due to their strategic importance and regulatory challenges.

Venezuelan, Iranian, or Russian oil.²¹

4. Sanctions and Dark Shipping on the Oil Market

In this section, we quantify the volume of dark-shipped oil and assess the impact of oil sanctions and dark shipping on the global oil market, including recorded world seaborne oil exports in official data and oil prices. Additionally, we identify China as the largest importer of dark-shipped oil, positioning it as a key player in shaping the distinct aggregate effects of oil sanctions on the U.S. and EU economies.

4.1. Oil Exports and Prices

Quantifying dark-shipped oil exports. Using our identification of dark ships, we estimate the monthly crude oil exports transported by these vessels from Iran, Syria, Venezuela, and Russia between 2017 and 2023. This approach quantifies the scale of dark shipping of sanctioned oil.

The estimation process unfolds in several steps. First, we compile a comprehensive set of trips undertaken by dark ships for each sample year. For trips without significant AIS gaps (i.e., those with gaps not exceeding the 99th percentile of time differences between consecutive AIS signals per vessel), we count those originating from ports in each of the four sanctioned countries.

For trips with at least one prolonged AIS gap, we first identify the geographically closest port to each long gap. We then determine the gap with the highest port-based trip suspicion score and attribute the trip to the corresponding sanctioned country.²² The total number of trips is allocated across months based on each trip’s start date.

Finally, we estimate the monthly crude oil exports from each sanctioned country carried by dark ships as:

$$X_{c,t} = \sum_{i \in \mathcal{N}_{c,t}^{EXP}} DWT_i, \quad (1)$$

where $X_{c,t}$ is the estimated crude oil exports from $c \in \{Iran, Syria, Venezuela, Russia\}$ transported by dark ships in month t , $\mathcal{N}_{c,t}^{EXP}$ is the set of qualifying trips for country c in month t ,

²¹Additional hotspots include the Strait of Malacca, the South China Sea, the Yellow Sea (off northern China), and the Arabian Sea (off Oman and the U.A.E.). Results are available upon request.

²²A dark ship may transport oil for multiple sanctioned economies within a single trip. Thus, attributing oil exports based solely on the highest port-based trip suspicion score may underestimate the true scale of sanctioned oil trade.

and DWT_i denotes the deadweight tonnage of the dark ship in trip i . This calculation assumes that each dark ship is fully loaded to its deadweight tonnage capacity per trip, providing an approximation of the maximum potential crude oil export volume.

Oil sanctions and the dynamics of recorded and dark-shipped oil. Figure 9 presents the monthly time series of crude oil exports transported by dark ships from the four sanctioned countries (stacked bars in varying shades of red). This is shown alongside recorded global seaborne crude oil exports (solid blue line) and a monthly index of oil-related sanctions imposed by Western authorities (dashed black line with diamond markers) from January 2017 to December 2023. All time series remain unadjusted for seasonal variation.

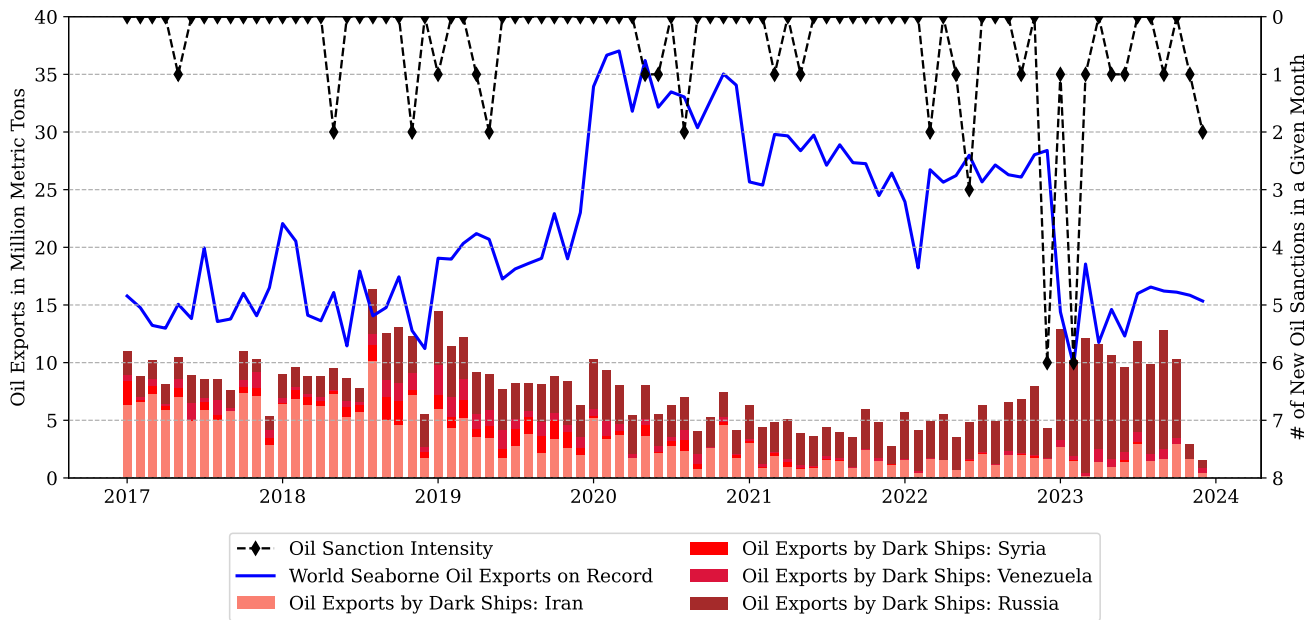


Figure 9: Oil Exports of Sanctioned Countries by Dark Ships, World Seaborne Oil Exports on Record, and Oil Sanction Intensity: 2017M1–2023M12

Notes. The stacked bars depict estimated monthly crude oil exports transported by dark ships from Iran, Syria, Venezuela, and Russia. The solid blue line represents global seaborne crude oil exports, sourced from the UN Comtrade database (HS code 2709: petroleum oils and oils from bituminous minerals, crude). Both series are measured in million metric tons and plotted on the left vertical axis. The dashed black line indicates the monthly oil sanction intensity index, constructed from the number of new oil-related sanctions imposed by Western countries each month. This index is plotted on the right vertical axis. All time series remain unadjusted for seasonal variation.

Data on global seaborne crude oil exports come from the UN Comtrade database, identified using HS code 2709 for petroleum oils and oils from bituminous minerals, crude. The sanction intensity index is derived from a manually curated dataset of oil-related sanctions imposed on Iran, Syria, Venezuela, and Russia by Western countries, including the U.S. and the EU, between

2017 and 2023.²³ For each month, we use the number of new sanctions as a proxy for oil sanction intensity. Theoretically, a higher number of sanctions exerts greater downward pressure on crude oil exports from sanctioned economies, increasing incentives for these countries and shipowners to engage in dark shipping.

Figure 9 highlights three key observations regarding crude oil exports from sanctioned countries transported by dark ships, global seaborne crude oil exports, and oil sanction intensity from January 2017 to December 2023. First, crude oil exports via dark ships from sanctioned countries constitute a substantial share of global seaborne crude oil exports. On average, these exports reach 7.8 million metric tons per month, approximately 43% of total recorded global seaborne crude oil exports in UN Comtrade data.²⁴ This underscores the critical role of dark shipping in the global oil trade, particularly among sanctioned nations.

Second, Iran and Russia are the primary contributors to crude oil exports via dark ships, with their shares fluctuating over time. From 2017 to 2019, Iran dominated sanctioned oil exports through dark shipping, coinciding with the period when Western nuclear-related sanctions were imposed (Laudati and Pesaran, 2023). In 2019, sanctions on Venezuela’s state-owned oil company PDVSA, led to a temporary modest increase in dark shipping. However, following the oil embargoes and price caps imposed by Western countries in response to Russia’s invasion of Ukraine in late 2022, Russia surpassed Iran as the largest exporter of sanctioned oil via dark shipping.

Third, periods of heightened oil sanction intensity – particularly in 2022–2023 – show that such sanctions reduced recorded global seaborne oil exports. However, they also triggered a surge in crude oil exports via dark shipping, as trading sanctioned oil became more profitable. For example, after the December 2022 implementation of the \$60 per barrel price cap on Russian crude (Babina et al., 2023; Kilian et al., 2024b), recorded global seaborne oil exports fell nearly 50% in January 2023. Simultaneously, sanctioned oil exports through dark shipping more than doubled, offsetting a large share of the direct reduction of supply caused by the price cap.

²³Table C.1 in Appendix C provides details on these sanctions, including dates, involved parties, and descriptions. We did not use the GSDB database (version 3; see Syropoulos et al., 2024) because (i) it lacks information on specific trade types (e.g., oil, minerals, or agriculture) and (ii) it records only the sanction year, whereas our analysis requires data with monthly frequency.

²⁴A significant portion of crude oil exports transported by dark ships – especially during prolonged AIS data gaps – may not appear in official records due to the covert nature of such trade, particularly post-sanctions. Consequently, the actual global oil trade volume could exceed UN Comtrade estimates.

Dark-shipped oil and world oil prices. Figure 10 plots de-trended WTI and Brent spot oil prices (dashed green and dash-dotted gray lines, respectively) alongside the share of world seaborne oil exports transported by dark ships (solid red line with diamond markers). When sanctions on Russia’s oil exports began in early 2022, both WTI and Brent prices surged by approximately 45%. However, starting in January 2023, dark shipping responded, with the share of oil exports transported by dark ships from sanctioned countries rising from 13% to 47%.




Figure 10: Oil Prices and Share of Global Seaborne Oil Exports From Sanctioned Countries Transported by Dark Ships: 2017M1–2023M12

Notes. The dashed green line represents the de-trended WTI spot price, and the dash-dotted gray line shows the de-trended Brent spot price, both in U.S. dollars per barrel, plotted on the left vertical axis. The solid red line represents the share of global oil exports transported by dark ships from Iran, Syria, Venezuela, and Russia. This share is calculated as total dark-shipped oil exports from these sanctioned countries divided by the sum of recorded global seaborne oil exports and dark-shipped oil exports from these nations. Expressed as a percentage, it is plotted on the right vertical axis. All time series remain unadjusted for seasonal variation.

The rising share of discounted dark-shipped oil contributed to declines in both WTI and Brent prices, as evidenced by the negative co-movement between the solid red line and the dashed green (WTI) and dash-dotted gray (Brent) lines since early 2022, despite the overall reduction in oil supply due to sanctions.²⁵ This finding highlights the importance of accounting for dark-shipped oil exports alongside direct sanctions when analyzing global oil price dynamics.

²⁵From March 2022 (when the first oil-related sanction on Russia was imposed) to December 2023, the correlation between the share of world oil exports transported by dark ships and the de-trended WTI price is approximately -0.54, while for the de-trended Brent price, it is -0.55.

4.2. China Leads Imports of Dark-Shipped Oil

Our trip-level identification of dark shipping enables us to pinpoint the major importers of discounted sanctioned oil. Countries importing such oil benefit from an increased energy supply and lower energy prices, which enhance domestic production and reduce costs for energy-intensive industries (Bornstein et al., 2023). Additionally, nations that trade extensively with these importers may gain indirectly through global production networks, experiencing higher demand for their goods or access to cheaper intermediate inputs. This dynamic extends the impact of discounted sanctioned oil beyond the primary importers.

To estimate monthly crude oil imports from Iran, Syria, Venezuela, and Russia transported by dark ships, we first analyze trips undertaken by dark ships without significant AIS gaps for each sample year.²⁶ For these trips originating from the four sanctioned countries, we trace their destinations and count the number of trips terminating in each importing country. We then estimate monthly crude oil imports transported by dark ships for each importer as:

$$M_{c,t} = \sum_{i \in \mathcal{N}_{c,t}^{IMP}} DWT_i, \quad (2)$$

where $M_{c,t}$ represents the estimated crude oil imports transported by dark ships to country c in month t , $\mathcal{N}_{c,t}^{IMP}$ denotes the set of trips by dark ships arriving in country c in month t , and DWT_i represents the deadweight tonnage of the dark ship involved in trip i .

Figure 11 presents the largest importers of crude oil exported from the four sanctioned countries via dark ships from 2017 to 2023. Each country’s total crude oil imports transported by dark ships were calculated using the methodology described above. Based on annual import volumes, we identified the top five importers for each year and aggregated their imports over the full period.

China is the largest importer of crude oil from sanctioned countries via dark shipping, with over 97 million metric tons imported during the sample period – about 15% of total crude oil exports from the four sanctioned nations. A significant share of this dark-shipped oil was imported after 2022, following sanctions on Russian oil exports in response to its invasion of Ukraine.

Beyond China, significant crude oil imports via dark shipping were recorded by South Korea, Egypt, India, and Russia itself. In terms of temporal trends, China, South Korea, and India were

²⁶Our analysis focuses on direct crude oil imports transported by dark ships with traceable AIS data. Imports involving prolonged AIS data gaps cannot be linked to specific destination countries, meaning our estimates may understate the true scale of sanctioned oil imports.

the primary importers of crude oil from sanctioned countries in 2017–2018. From 2019 onward, Egypt, Russia, the Netherlands, and Turkey significantly increased their imports of crude oil transported by dark ships.²⁷




Figure 11: Major Importers of Crude Oil from Sanctioned Countries Transported by Dark Ships: 2017–2023

Notes. Largest importers of crude oil exported from Iran, Syria, Venezuela, and Russia via dark ships from 2017 to 2023, along with corresponding import volumes in million metric tons. Oil imports across years are represented by different shades of red. Import volumes are estimated using Equation (2), summing the deadweight tonnage of dark ships arriving in each country.

5. Macroeconomic Consequences

In this section, we use local projections (LPs) to analyze the dynamic causal effects of unexpected oil sanction intensity shocks – both direct and through dark shipping – on inflation and production in the U.S. and EU.²⁸ We identify the propagation of oil sanction shocks along the global supply chain as the key mechanism driving aggregate dynamics in these economies. China, as the largest importer of discounted dark-shipped oil, plays a central role in this transmission.

²⁷Unlike existing studies (e.g., Hilgenstock et al. 2023; Kilian et al. 2024b), which identify India as a major importer of Russian oil, our results do not find it to be the primary destination for Russian crude. This discrepancy likely arises because India primarily relies on the “shadow fleet” for oil (Kilian et al., 2024b), whereas our algorithm is specifically designed to identify dark ships importing sanctioned Russian oil.

²⁸LPs provide a flexible model specification and are robust against the curse of dimensionality, making them well-suited for capturing dynamic responses over time (Fernández-Villaverde et al., 2025; Li et al., 2024a).

5.1. Baseline Results

Our causality assessment involves identifying an unexpected oil sanction intensity shock using the following LP specification:

$$y_{i,t+k} = \beta_{i,k,0} \cdot SANCTION_t + \sum_{l=1}^L \beta'_{i,k,l} \mathbf{y}_{t-l} + \alpha_k + u_{i,k,t}, \quad (3)$$

where $1 \leq i \leq n$, $0 \leq k \leq K$, \mathbf{y}_t is an $n \times 1$ vector of the endogenous variables, and $SANCTION_t$ represents the sanction intensity index. \mathbf{y}_{t+k} is the $n \times 1$ vector of k -step-ahead forecasts of the endogenous variables, with $y_{i,t+k}$ denoting the value of the i -th variable in \mathbf{y}_{t+k} . The vector \mathbf{y}_{t-l} includes the lagged endogenous variables up to order l , L is the lag length, α_k is a constant, and $u_{i,k,t}$ is the prediction error for the i -th variable. The coefficients of interest are $\beta_{i,k,0}$ for $k = 0, \dots, K$, which capture the dynamic responses of the i -th variable up to K months following a one-standard-deviation unexpected oil sanction intensity shock at time t .

For the baseline estimation, we employ 6-variable LP models to assess the aggregate effects on the U.S. and EU. The U.S. model includes monthly time series for: (i) the sanction intensity index, (ii) world seaborne crude oil exports recorded in UN Comtrade data, (iii) crude oil exports from Iran, Syria, Venezuela, and Russia via dark ships, (iv) the WTI spot price, (v) the U.S. producer price index (PPI), and (vi) U.S. industrial production (IP), covering January 2017 to December 2023. For the EU model, we replace the WTI spot price with Brent and substitute U.S. PPI and IP with their EU equivalents.²⁹ All series, except for the sanction intensity index, are seasonally adjusted. The sanction intensity index is standardized and enters the LP model in standard deviations, while all other endogenous variables enter in log percent. Appendix D provides more details.

The LP model (3) can be estimated using the least squares method (Jordà, 2005), with Newey-West standard errors (Newey and West, 1987) applied to quantify the uncertainty around the implied impulse responses for $\beta_{i,k,0}$. However, as the forecast horizon k increases, the estimates of $\beta_{i,k,0}$ may exhibit higher variance, particularly given the small sample size.

To address this concern, we apply the smooth local projections (SLP) method in Barnichon and Brownlees (2019), approximating $\beta_{i,k,0}$ using a linear B-spline basis function expansion at

²⁹WTI represents the U.S. domestic market, while Brent serves as the global benchmark and is more relevant for EU oil prices.

horizon k :

$$\beta_{i,k,0} \approx \sum_{m=1}^M b_{i,m} B_{i,m}(k), \quad (4)$$

where $B_{i,m}(k) : \mathbb{R} \rightarrow \mathbb{R}$ for $m = 1, \dots, M$ represents a set of B-spline basis functions, and $b_{i,m}$ are scalar parameters. These parameters are estimated using a penalized estimator, which shrinks Equation (4) toward a quadratic polynomial at horizon k .

Following [Barnichon and Brownlees \(2019\)](#), we identify an oil sanction intensity shock by controlling only for the lagged endogenous variables \mathbf{y}_{t-l} (i.e., no contemporaneous values of other endogenous variables are included as controls) and by estimating the LP model (3) starting from horizon $k = 0$. This approach is equivalent to placing the sanction intensity index first in a recursive ordering, allowing all other endogenous variables to respond contemporaneously to the oil sanction intensity shock. We include three lags of the endogenous variables in the LP model (i.e., $L = 3$) based on the Hannan-Quinn Information Criterion (HQIC), and our results remain robust when considering longer lags.

Figures 12 and 13 show the impulse response functions (IRFs) following an unexpected oil sanction intensity shock for the U.S. and EU economies, respectively. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands. Let us discuss the two key findings from these figures.

Sanctions reduce recorded exports, while dark shipping mitigates oil price hikes. An unexpected shock increases the sanction intensity index by over 0.4 standard deviations, equivalent to approximately 0.5 additional oil-related sanctions (Panels 12a and 13a). This shock triggers an immediate and sustained decline in recorded global seaborne oil exports, lasting up to a year post-impact (Panels 12b and 13b). While this contraction exerts upward pressure on global oil prices, its effect is largely offset in the first six months by a surge in dark-shipped oil exports (Panels 12c and 13c), which peak around the fourth month before falling sharply.

Sanctions drive significant increases in WTI and Brent spot prices only after dark-shipped oil exports decline (Panels 12d and 13d). In the months following the oil sanction intensity shock, the high volume of discounted dark-shipped oil suppresses global oil price mark-ups, preventing a Brent price increase in the EU and causing a WTI price decline in the U.S.




Figure 12: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: Baseline Results

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using the SLP method as described in [Barnichon and Brownlees \(2019\)](#). The shock is identified by controlling only for the lagged endogenous variables specified in Equation (3) and by estimating the model from horizon $k = 0$. This approach is equivalent to a recursive identification scheme, with the sanction intensity index positioned first in the ordering. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

Distinct inflation and production responses in the U.S. and EU. Reflecting the divergent post-impact responses of WTI and Brent spot prices, the U.S. and EU economies exhibit distinct patterns in both PPI and IP, differing in both shape and magnitude.

In the U.S., a net oil exporter, the initial WTI price decline reduces revenues for oil producers as importers shift to discounted dark-shipped oil, contributing to a drop in the PPI during the first six months (Panel 12e). However, the U.S. PPI declines more sharply than WTI, and when WTI rebounds after the sixth month, the PPI merely returns to pre-impact levels. This suggests additional supply-side effects through which escalating oil sanctions and dark shipping suppress the PPI. This pattern also appears in U.S. IP, which experiences a mild but positive impact initially.




Figure 13: IRFs to an Oil Sanction Intensity Shock for the EU Economy: Baseline Results

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the EU economy are computed using the SLP method as described in [Barnichon and Brownlees \(2019\)](#). Except for using a set of endogenous variables specific to the EU, the estimation specification is identical to that for the U.S. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

While lower oil prices might typically reduce demand in both U.S. energy and non-energy sectors, potentially leading to a demand-driven decline in industrial production, the observed IP increase suggests offsetting factors (Panel 12f).

In contrast, as a net oil importer, the EU sees only a marginally negative PPI response in the first six months after the oil sanction intensity shock, with confidence bands suggesting limited statistical significance (Panel 13e). This muted response aligns with the subdued Brent spot price, as dark-shipped oil floods the market during this period. However, as Brent prices rise sharply following the step drop in dark-shipped oil exports after the sixth month, the EU PPI begins to inflate, peaking at around 1% near the one-year mark.

Interestingly, the surge in oil prices and energy costs in the EU does not trigger a supply-driven production contraction in later months. Instead, EU IP remains elevated throughout the

forecast horizon, peaking beyond the six-month mark (Panel 13f). Compared to the U.S. IP response (Panel 12f), the EU response is not only statistically significant – evidenced by the 90% confidence band’s lower bound touching the zero-response line at its peak – but also larger (approximately 0.8% for the EU vs. just under 0.3% for the U.S.). This suggests additional demand-side factors driving the increase in EU production.

To assess robustness, we modify the baseline specification in three ways: (i) extending the lag length ($L = 4$), (ii) adding a linear trend, and (iii) replacing WTI and Brent spot prices with U.S. and EU crude oil import prices.³⁰ As shown in Figures E.1 to E.6 in Appendix E, our results remain robust. We further test robustness using a Cholesky-ordered structural vector autoregression (SVAR) model for both the U.S. and the EU. Figures E.7 and E.8 show quantitatively similar SVAR results.

Lastly, Figure E.9 analyzes the response of the Organization of the Petroleum Exporting Countries (OPEC) to an oil sanction intensity shock. The IRFs indicate declines in OPEC’s oil export volume, price, and net revenue, though to varying extents. Meanwhile, spare oil production capacity increases, signaling a looser oil market.³¹ As dark-shipped oil exports decline after the sixth month and OPEC’s exports continue to fall, OPEC’s spot crude oil price rises sharply, boosting net oil export revenue. Spare capacity remains elevated, reflecting OPEC’s strategy of using supply cuts to restore revenue margins.

5.2. Propagation Channels

The missing propagation channels needed to explain the distinct PPI inflation and production responses in the U.S. and EU prompt us to look beyond these economies and examine the transmission of oil sanction intensity shocks and dark shipping along the global supply chain. As shown in Figure 11, China absorbs the largest share of discounted dark-shipped oil and is a key trading partner for both the U.S. and the EU. This positions China as a central hub in propagating these shocks.³²

³⁰Although the U.S. is a net oil exporter, it still imports oil because domestic refineries often require heavier crude that is not produced domestically.

³¹Spare capacity, a measure of market tightness, refers to the volume OPEC can activate within 30 days and sustain for at least 90 days. See <https://www.eia.gov/finance/markets/crudeoil/supply-pec.php> (accessed January 18, 2025) for details.

³²Figure E.10 analyzes China’s response to an unexpected oil sanction intensity shock, focusing on the producer price index (PPI) and the value-added index (VAI) for mining, quarrying, and manufacturing. As the largest importer of discounted dark-shipped oil, China benefits from lower costs and increased production in these

To explore this idea, we expand the U.S. LP model by adding two endogenous variables: China’s crude oil import price and the U.S. import price for Chinese manufactured goods, averaged across NAICS sectors 31–33. For the EU LP model, we include China’s total IP and its total import value from the EU. Additionally, we distinguish between energy and non-energy sectors by replacing the total PPI and IP for the U.S. and EU in the LP model (3) with sector-specific PPI and IP. All additional series are seasonally adjusted and expressed in log percent. Appendix D provides further details.

The estimation specifications remain consistent with the baseline. We include only lagged endogenous variables and estimate the augmented LP models from horizon $k = 0$, maintaining a lag length of three ($L = 3$). In the following discussion, we highlight key panels illustrating the propagation channels, while additional panels are provided in Figures E.11 and E.12 in Appendix E for brevity.

Decline in oil import price and increased production in China. As the leading importer of discounted dark-shipped oil, China experiences an immediate and sharp decline in its oil import price following an unexpected oil sanction intensity shock (Panel 14e).³³ The drop in oil prices propagates through input-output linkages, reducing production costs and boosting China’s industrial production (Panel 15e).³⁴ However, from the sixth month post-impact, the rebound in oil import prices – driven by the rapid decline in dark-shipped oil exports (Panel 12c) – triggers a contraction in Chinese production.

The role of China in propagating oil sanction shocks. The decline in oil import prices and the resulting rise in China’s industrial production impact the U.S. and EU differently, reflecting their distinct input-output linkages with China in global production networks and varying exposure to energy markets.

For the U.S., as shown in Figure 12d, the decline in the WTI spot price and lower oil export revenues for domestic producers reduce PPI (Panels 14a and 14c) and IP (Panels 14b and 14d) across both sectors. The decline in U.S. oil export volumes during the first six months post-impact

industries.

³³In the case of the 2022 sanctions on Russian oil, as observed by Kilian et al. (2024b), the decline in China’s oil import price may result from its increased market power following the shift in oil supply from the EU to China.

³⁴Figure E.13 in Appendix E confirms the reduction in production costs, showing a significant drop in China’s total PPI. This also leads to a sharp decline in the U.S. import price for Chinese goods, including manufactured products, as seen in Figure E.14 and Figure 14f.

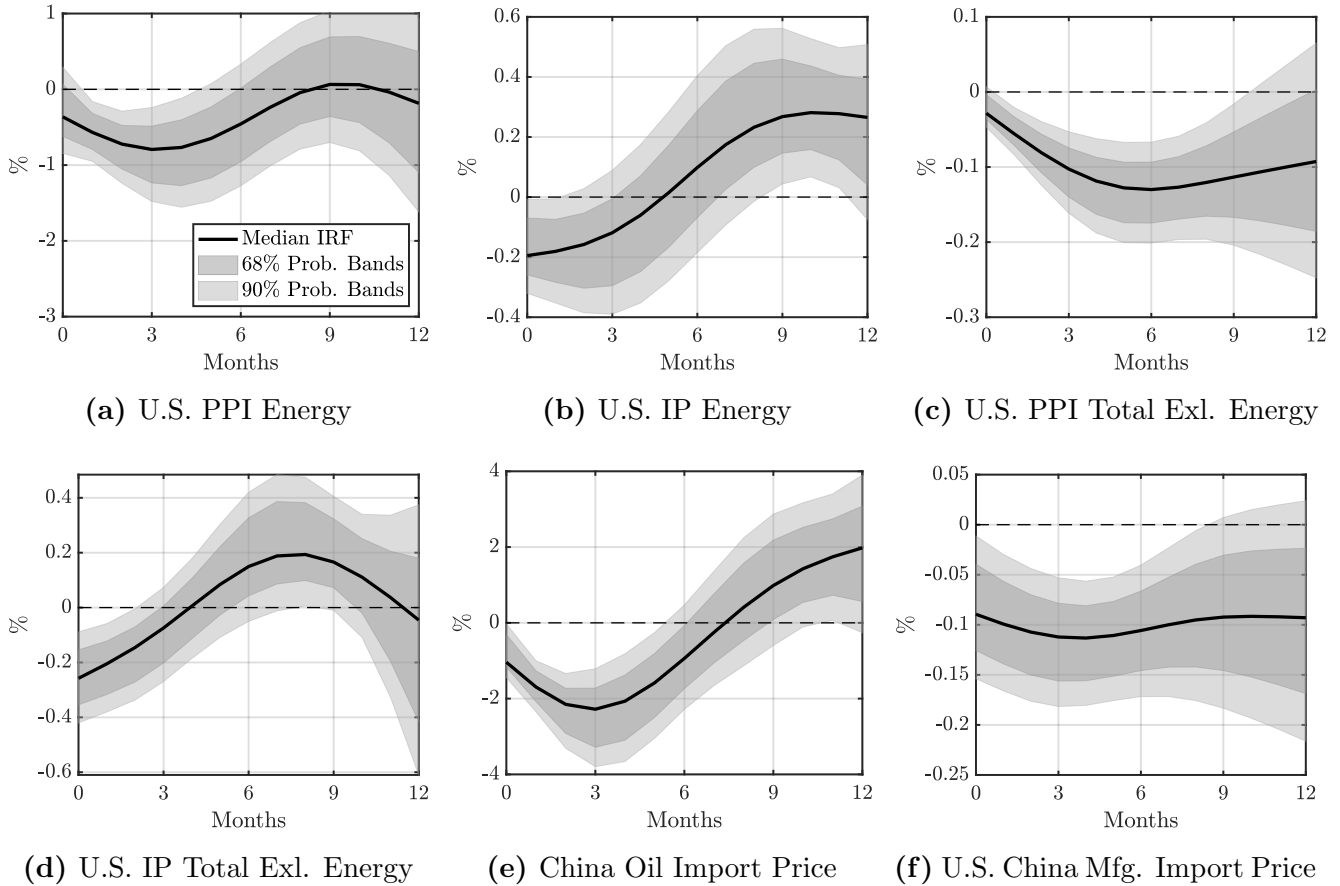


Figure 14: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: Propagation Channels

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using the SLP method described in [Barnichon and Brownlees \(2019\)](#). The estimation specifications follow the baseline outlined in Section 5.1, with two modifications: (i) two additional variables – China’s crude oil import price and the U.S. import price for manufactured goods from China, averaged across NAICS sectors 31 to 33 – are included, and (ii) the total PPI and IP for the U.S. economy are replaced by the PPI and IP for the energy and non-energy sectors. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

reinforces this, as shown in Figure E.15 in Appendix E. Both reduced volumes and lower prices contribute to shrinking oil export revenues.

Additionally, lower production costs in China reduce the U.S. import price of Chinese manufactured goods (Panel 14f), driving cost-saving growth that amplifies PPI deflation while partially offsetting initial production declines. This effect is clearer when excluding China-related variables and re-estimating the LP model. As shown in Figure E.16 of Appendix E, during the first three months post-impact, the median responses of IP for U.S. energy and non-energy sectors are closer to zero-response lines than in Panels 14b and 14d. This suggests a positive omitted variable bias, where lower Chinese import prices support short-term U.S. production.




Figure 15: IRFs to an Oil Sanction Intensity Shock for the EU Economy: Propagation Channels

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the EU economy are computed using the SLP method described in [Barnichon and Brownlees \(2019\)](#). The estimation specifications follow the baseline outlined in Section 5.1, with two modifications: (i) two additional variables – China’s total IP and its total import value from the EU – are included, and (ii) the total PPI and IP for the EU economy are replaced by the PPI and IP for the energy and non-energy sectors. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

In the medium term, as oil prices rebound (Panel 12d), the U.S. energy sector recovers sharply, evident in the rising IP after six months and the mean-reverting PPI. Non-energy sectors also see production gains, though PPI continues to decline due to sustained lower costs from cheaper Chinese inputs (Panels 14c and 14f).

For the EU, rising Chinese demand for its intermediate goods (Panel 15f) – driven by expanding industrial production (Panel 15e) – fuels strong demand-driven growth. This pushes up PPI in non-energy sectors (Panel 15c) and stimulates production (Panel 15d). Growth propagates upstream, increasing energy demand, which in turn raises PPI (Panel 15a) and production (Panel

15b) in the EU energy sector.³⁵

As China’s expansion slows and its imports from the EU decline, production in both EU sectors contracts. However, prices remain elevated, particularly in non-energy sectors (Panel 15c), due to rising oil prices (Panel 13d) and high energy costs. These factors contribute to a stagflationary impact on the EU economy.

Takeaways. Our findings highlight the nuanced impact of oil sanctions and dark shipping on economies within the global supply chain. As a net oil exporter, the U.S. benefits from lower oil prices, which reduce domestic producers’ revenues but also lower import costs from China, driving deflationary pressure and supply-side growth. In contrast, as a net oil importer, the EU faces higher energy costs but gains from rising Chinese demand, leading to inflation and demand-driven growth. Meanwhile, China’s access to discounted dark-shipped oil significantly boosts industrial production, amplifying the effects of oil sanction shocks and shaping aggregate dynamics in both the U.S. and the EU economies.

6. Conclusion

This paper presents a novel multi-attribute ship clustering model to detect dark shipping activities associated with sanctioned oil transport. By integrating vessel-specific traits, evasive behavior during AIS signal gaps, and navigational anomalies, our model identifies covert activity within the global crude oil tanker fleet. Our findings show that dark ships constitute a significant portion of the fleet, facilitating the trade of sanctioned oil. Between 2017 and 2023, an average of 7.8 million metric tons of crude oil per month – approximately 43% of officially recorded global seaborne exports – was transported via dark ships.

Building on these findings, we derive three key empirical insights. First, the increased exports of discounted dark-shipped oil partially offset the reduction in recorded global oil exports caused by oil sanctions, thereby lowering global oil prices. Second, China is the largest importer of dark-shipped oil from sanctioned countries. Third, LP analysis reveals that intensifying sanctions drive sanctioned oil exports via dark shipping, producing asymmetric effects. As a net oil exporter, the U.S. benefits from lower Chinese import prices, which induce deflationary pressure and supply-driven growth, despite reduced revenues for domestic producers. In contrast, the EU, a net oil

³⁵Figure E.12 in Appendix E shows that rising energy demand also lifts the Brent spot price.

importer, faces higher energy costs but benefits from rising Chinese demand, leading to inflation and demand-driven growth. Meanwhile, China’s access to discounted oil boosts industrial production, propagating the economic effects of sanctions and shaping economic dynamics in both the U.S. and the EU.

Several promising avenues for future research remain. First, modeling the global economy with input-output linkages and a micro-founded international oil market could provide deeper insights into the distinct responses of the U.S. and the EU to oil sanction shocks. Extending the analysis to examine production networks (as in [Ghassibe, 2024](#) or [Qiu et al., 2025](#)) could help determine whether variations in the geographical impact of sanctions are primarily driven by spatial proximity ([Redding and Sturm, 2008](#)), increased trading costs ([Xu et al., 2025](#)), or production synergies ([Fernández-Villaverde et al., 2021, 2024, 2025](#)), which play an important role in the formation and persistence of trade relationships. Second, policy counterfactuals could evaluate alternative sanction regimes and explore measures to mitigate unintended consequences, such as the EU’s inflationary pressure resulting from higher energy costs and rising Chinese demand. Finally, integrating satellite imagery could enhance validation, particularly in detecting ship-to-ship transfers, thereby improving model accuracy. We plan to pursue some of these extensions in future work.

References

- Androjna, A., Pavić, I., Gucma, L., Vidmar, P., and Perkovič, M. (2024). AIS Data Manipulation in the Illicit Global Oil Trade. *Journal of Marine Science and Engineering*, 12(1).
- Asker, J., Collard-Wexler, A., and De Loecker, J. (2019). (Mis)Allocation, Market Power, and Global Oil Extraction. *American Economic Review*, 109(4):1568–1615.
- Babina, T., Hilgenstock, B., Itskhoki, O., Mironov, M., and Ribakova, E. (2023). Assessing the Impact of International Sanctions on Russian Oil Exports. VoxEU.
- Bai, X., Fernández-Villaverde, J., Li, Y., and Zanetti, F. (2024). The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory. Working Paper 32098, National Bureau of Economic Research.
- Bai, X., Fernández-Villaverde, J., Li, Y., and Zanetti, F. (2025). State Dependence of Monetary Policy During Global Supply Chain Disruptions. Working paper.
- Bai, X., Xu, M., Han, T., and Yang, D. (2022). Quantifying the Impact of Pandemic Lockdown Policies on Global Port Calls. *Transportation Research Part A: Policy and Practice*, 164:224–241.
- Ballinger, O. (2024). Automatic Detection of Dark Ship-To-Ship Transfers Using Deep Learning and Satellite Imagery. *arXiv preprint*.
- Barnichon, R. and Brownlees, C. (2019). Impulse Response Estimation by Smooth Local Projections. *Review of Economics and Statistics*, 101(3):522–530.
- Baumeister, C. and Hamilton, J. D. (2019). Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks. *American Economic Review*, 109(5):1873–1910.
- Baumeister, C. and Kilian, L. (2016). Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us. *Journal of Economic Perspectives*, 30(1):139–60.
- Bernabé, P., Gotlieb, A., Legeard, B., Marijan, D., Sem-Jacobsen, F. O., and Spieker, H. (2024). Detecting Intentional AIS Shutdown in Open Sea Maritime Surveillance Using Self-Supervised Deep Learning. *IEEE Transactions on Intelligent Transportation Systems*, 25(2):1166–1177.
- Bjørnland, H. C., Nordvik, F. M., and Rohrer, M. (2021). Supply Flexibility in the Shale Patch: Evidence From North Dakota. *Journal of Applied Econometrics*, 36(3):273–292.
- Bodenstein, M., Erceg, C. J., and Guerrieri, L. (2011). Oil Shocks and External Adjustment. *Journal of International Economics*, 83(2):168–184.
- Bornstein, G., Krusell, P., and Rebelo, S. (2023). A World Equilibrium Model of the Oil Market. *The Review of Economic Studies*, 90(1):132–164.
- Brancaccio, G., Kalouptsi, M., and Papageorgiou, T. (2020). Geography, Transportation, and Endogenous Trade Costs. *Econometrica*, 88(2):657–691.

- Brancaccio, G., Kalouptsi, M., and Papageorgiou, T. (2024). Investment in Infrastructure and Trade: The Case of Ports. Working Paper 32503, National Bureau of Economic Research.
- Cardoso, D. S., Salant, S. W., and Daubanes, J. (2025). The Dynamics of Evasion: The Price Cap on Russian Oil Exports and the Amassing of the Shadow Fleet. Working Paper 11618, CESifo.
- Cerdeiro, D. A., Komaromi, A., Liu, Y., and Saeed, M. (2020). World Seaborne Trade in Real Time: A Proof of Concept for Building AIS-Based Nowcasts From Scratch. Working Paper 2020/57, International Monetary Fund.
- Cook, C. and Sheppard, D. (2023). Iran’s ‘Ghost Fleet’ Switches Into Russian Oil. Financial Times.
- de Souza, G., Hu, N., Li, H., and Mei, Y. (2024). (Trade) War and Peace: How to Impose International Trade Sanctions. *Journal of Monetary Economics*, 146:103572.
- Dunn, J. and Leibovici, F. (2023). Navigating the Waves of Global Shipping: Drivers and Aggregate Implications. Working Paper 2023-002, Federal Reserve Bank of St. Louis.
- European Parliament (2024). Parliament Calls for an EU Crackdown on Russia’s ‘Shadow Fleet’. European Parliament Press Releases.
- Felbermayr, G., Morgan, T. C., Syropoulos, C., and Yotov, Y. V. (2024). Economic Sanctions: Stylized Facts and Quantitative Evidence. *Annual Review of Economics*. forthcoming.
- Fernández-Villaverde, J., Mandelman, F., Yu, Y., and Zanetti, F. (2021). The “Matthew Effect” and Market Concentration: Search Complementarities and Monopsony Power. *Journal of Monetary Economics*, 121:62–90.
- Fernández-Villaverde, J., Mandelman, F., Yu, Y., and Zanetti, F. (2025). Search Complementarities, Aggregate Fluctuations, and Fiscal Policy. *Review of Economic Studies*, page rdae053.
- Fernández-Villaverde, J., Mineyama, T., and Song, D. (2025). Are We Fragmented Yet? Measuring Geopolitical Fragmentation and Its Causal Effect. Working Paper 32638, National Bureau of Economic Research.
- Fernández-Villaverde, J., Yu, Y., and Zanetti, F. (2024). Technological Synergies, Heterogeneous Firms, and Idiosyncratic Volatility. Working Paper 32247, National Bureau of Economic Research.
- Ghassibe, M. (2024). Endogenous Production Networks and Non-Linear Monetary Transmission. Working paper, CREi.
- Ghironi, F., Kim, D., and Ozhan, G. K. (2024). International Economic Sanctions and Third-Country Effects. *IMF Economic Review*, 72(2):611–652.
- Greenwood, R. and Hanson, S. G. (2015). Waves in Ship Prices and Investment. *Quarterly Journal of Economics*, 130(1):55–109.

- Heiland, I., Moxnes, A., Ulltveit-Moe, K. H., and Zi, Y. (2019). Trade From Space: Shipping Networks and the Global Implications of Local Shocks. Discussion Paper 14193, Centre for Economic Policy Research.
- Hilgenstock, B., Ribakova, E., Shapoval, N., Babina, T., Itskhoki, O., and Mironov, M. (2023). Russian Oil Exports Under International Sanctions. Working paper, SSRN.
- Hwang, C.-L. and Yoon, K. (1981). *Multiple Attribute Decision Making: Methods and Applications A State-Of-The-Art Survey*. Springer Berlin Heidelberg.
- Itskhoki, O. and Mukhin, D. (2022). Sanctions and the Exchange Rate. Working Paper 32247, National Bureau of Economic Research.
- Johnson, S., Rachel, L., and Wolfram, C. (2023a). A Theory of Price Caps on Non-Renewable Resources. Working Paper 31347, National Bureau of Economic Research.
- Johnson, S., Rachel, L., and Wolfram, C. (2023b). Design and Implementation of the Price Cap on Russian Oil Exports. *Journal of Comparative Economics*, 51(4):1244–1252.
- Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.
- Kalouptsi, M. (2014). Time to Build and Fluctuations in Bulk Shipping. *American Economic Review*, 104(2):564–608.
- Känzig, D. R. (2021). The Macroeconomic Effects of Oil Supply News: Evidence From OPEC Announcements. *American Economic Review*, 111:1092–1125.
- Kilian, L. (2009). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review*, 99(3):1053–69.
- Kilian, L. (2014). Oil Price Shocks: Causes and Consequences. *Annual Review of Resource Economics*, 6:133–154.
- Kilian, L., Plante, M. D., and Richter, A. W. (2024a). Geopolitical Oil Price Risk and Economic Fluctuations. Working Paper 2403, Federal Reserve Bank of Dallas.
- Kilian, L., Rapson, D., and Schipper, B. (2024b). The Impact of the 2022 Oil Embargo and Price Cap on Russian Oil Prices. Working Paper 2401, Federal Reserve Bank of Dallas.
- Komaromi, A., Cerdeiro, D. A., and Liu, Y. (2022). Supply Chains and Port Congestion Around the World. Working Paper 2022/59, International Monetary Fund.
- Laudati, D. and Pesaran, M. H. (2023). Identifying the Effects of Sanctions on the Iranian Economy Using Newspaper Coverage. *Journal of Applied Econometrics*, 38(3):271–294.
- Levi, I., Katinas, P., Myllyvirta, L., and Hemalatha, K. (2023). Shedding Light on Shadow Tankers: Who Transports Russian Oil 18 Months Into the Invasion? Technical report, Centre for Research on Energy and Clean Air.
- Li, D., Plagborg-Møller, M., and Wolf, C. K. (2024a). Local Projections vs. VARs: Lessons From Thousands of DGPs. *Journal of Econometrics*, page 105722.

- Li, H., Li, Z., Park, Z., Wang, Y., and Wu, J. (2024b). To Comply or Not to Comply: Understanding Neutral Country Supply Chain Responses to Russian Sanctions. Working paper, SSRN.
- Li, W., Bai, X., Yang, D., and Hou, Y. (2023). Maritime Connectivity, Transport Infrastructure Expansion and Economic Growth: A Global Perspective. *Transportation Research Part A: Policy and Practice*, 170:103609.
- Li, Y., Bai, X., Wang, Q., and Ma, Z. (2022). A Big Data Approach to Cargo Type Prediction and Its Implications for Oil Trade Estimation. *Transportation Research Part E: Logistics and Transportation Review*, 165:102831.
- Lippi, F. and Nobili, A. (2012). Oil and the Macroeconomy: A Quantitative Structural Analysis. *Journal of the European Economic Association*, 10(5):1059–1083.
- Lloyd, S. (1982). Least Squares Quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137.
- Lloyd’s List Intelligence (2023). Shifty Shades of Grey: The Different Risk Profiles of the Dark Fleet Explained. Technical report, Lloyd’s List Intelligence.
- Lloyd’s List Intelligence (2024a). Understanding AIS Spoofing: The Case of the Blazers. Seasearcher Advanced Risk & Compliance.
- Lloyd’s List Intelligence (2024b). Understanding AIS Spoofing: The Case of the Shanaye Queen. Seasearcher Advanced Risk & Compliance.
- MacQueen, J. B. (1967). Some Methods for Classification and Analysis of Multivariate Observations. In Cam, L. M. L. and Neyman, J., editors, *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297. University of California Press.
- Meade, R. (2023). Euronav Investigating Claims It Unwittingly Stored Sanctioned Iranian Oil. Lloyd’s List Intelligence.
- Melek, N., Plante, M., and Yücel, M. K. (2017). The U.S. Shale Oil Boom, the Oil Export Ban, and the Economy: A General Equilibrium Analysis. Working Paper 23818, National Bureau of Economic Research.
- Nagle, P. (2020). The Oil Market Outlook: Lasting Scars From the Pandemic. World Bank Data Blog.
- Newey, W. K. and West, K. D. (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–708.
- Notteboom, T., Haralambides, H., and Cullinane, K. (2024). The Red Sea Crisis: Ramifications for Vessel Operations, Shipping Networks, and Maritime Supply Chains. *Maritime Economics & Logistics*, 26(1):1–20.
- Paris MoU (2018). Explanatory Note – “White,” “Grey” and “Black List”. Technical report, Paris Memorandum of Understanding on Port State Control.

- Prochazka, V., Adland, R., and Wolff, F.-C. (2019). Contracting Decisions in the Crude Oil Transportation Market: Evidence From Fixtures Matched With AIS Data. *Transportation Research Part A: Policy and Practice*, 130:37–53.
- Qiu, Z., Wang, Y., Xu, L., and Zanetti, F. (2025). Monetary Policy in Open Economies With Production Networks. Discussion Paper 19828, Centre for Economic Policy Research.
- Redding, S. J. and Sturm, D. M. (2008). The Costs of Remoteness: Evidence From German Division and Reunification. *American Economic Review*, 98(5):1766–1797.
- Rodger, M. and Guida, R. (2022). Mapping Dark Shipping Zones Using Multi-Temporal SAR and AIS Data for Maritime Domain Awareness. In *2022 IEEE International Geoscience and Remote Sensing Symposium*, pages 3672–3675, Kuala Lumpur, Malaysia.
- Rodríguez, F. R. (2025). Sanctioning Venezuela. In *The Collapse of Venezuela: Scorched Earth Politics and Economic Decline, 2012-2020*, chapter 9. University of Notre Dame Press, Notre Dame, 1 edition. forthcoming.
- Sarker, I. H. (2021). Machine Learning: Algorithms, Real-World Applications and Research Directions. *SN Computer Science*, 2(3):160.
- S&P Global Market Intelligence (2023). Russia’s Shadow Fleet – Understanding Its Size, Activity and Relationships. Technical report, S&P Global Market Intelligence.
- Sturm, J. (2023). How Should Sanctions Account for Bystander Countries? *AEA Papers and Proceedings*, 113:39–42.
- Syropoulos, C., Felbermayr, G., Kirilakha, A., Yalcin, E., and Yotov, Y. V. (2024). The Global Sanctions Data Base–Release 3: COVID-19, Russia, and Multilateral Sanctions. *Review of International Economics*, 32(1):12–48.
- Triebert, C., Migliozi, B., Bedi, N., and Cardia, A. (2024). The \$2.8 Billion Hole in U.S. Sanctions on Iran. *The New York Times*.
- Triebert, C., Migliozi, B., Cardia, A., Xiao, M., and Botti, D. (2023). Fake Signals and American Insurance: How a Dark Fleet Moves Russian Oil. *The New York Times*.
- UNCTAD (2022). Review of Maritime Transport 2022. Technical report, United Nations Conference on Trade and Development.
- U.S. Department of State (2018). Background Briefing on U.S. Withdrawal From the Joint Comprehensive Plan of Action (JCPOA). Foreign Press Centers Briefing.
- U.S. Department of State (2020). Sanctions advisory for the maritime industry, energy and metals sectors, and related communities. Technical report, U.S. Department of State, U.S. Department of the Treasury’s Office of Foreign Assets Control, and U.S. Coast Guard.
- U.S. Department of the Treasury (2019). Treasury Sanctions Venezuela’s State-Owned Oil Company Petroleos de Venezuela, S.A. Office of Foreign Assets Control.

- U.S. Department of the Treasury (2023). With Over 300 Sanctions, U.S. Targets Russia's Circumvention and Evasion, Military-Industrial Supply Chains, and Future Energy Revenues. Office of Foreign Assets Control.
- U.S. Energy Information Administration (2024). Country Analysis Brief: World Oil Transit Chokepoints. Technical report, U.S. Energy Information Administration.
- Verduzco-Bustos, G. and Zanetti, F. (2025). The Effects of Geopolitical Oil Price Shocks. Working paper.
- Wijaya, M. (2022). Palau Cancels Sanctioned VLCC Young Yong's Insurance Certificate. Lloyd's List Intelligence.
- Wu, J. (2012). *Advances in K-Means Clustering: A Data Mining Thinking*. Springer Berlin Heidelberg.
- Xu, D. and Tian, Y. (2015). A Comprehensive Survey of Clustering Algorithms. *Annals of Data Science*, 2(2):165–193.
- Xu, L., Yu, Y., and Zanetti, F. (2025). The Adoption and Termination of Suppliers Over the Business Cycle. *Journal of Monetary Economics*, forthcoming.
- Zhang, L., Yang, D., Bai, X., and Lai, K.-h. (2023). How Liner Shipping Heals Schedule Disruption: A Data-Driven Framework to Uncover the Strategic Behavior of Port-Skipping. *Transportation Research Part E: Logistics and Transportation Review*, 176:103229.

Online Appendices

The (Un)Intended Consequences of Oil Sanctions Through the Dark Shipping of Sanctioned Oil

Xiwen Bai,[†] *Jesús Fernández-Villaverde*,[‡] *Yiliang Li*,[§] *Le Xu*,[◇] *Francesco Zanetti*[¶]

Contents

A	A Multi-Attribute Ship Clustering Model for Dark Ship Identification	A-2
A.1	Additional Data Sources	A-3
A.2	Data Preprocessing and Trip Construction	A-4
A.3	Feature Extraction and Trip Classification	A-5
A.4	Dark Ship Identification	A-15
B	Additional Dark Ship Identification Results	A-16
C	Details of Oil Sanctions on Iran, Syria, Venezuela, and Russia in 2017–2023	A-21
D	Data Sources for LP Estimation	A-22
E	Robustness of LP Results	A-23
E.1	Baseline Results	A-23
E.2	Propagation Channels	A-31
	References for Appendices	A-37

[†] Bai: Tsinghua University, China. xiwenbai@mail.tsinghua.edu.cn. [‡] Fernández-Villaverde: University of Pennsylvania, U.S. jesusfv@econ.upenn.edu. [§] Li: University of International Business and Economics, China. yiliang.li@uibe.edu.cn. [◇] Xu: Shanghai Jiao Tong University, China. lexu1@sjtu.edu.cn. [¶] Zanetti: University of Oxford, U.K. francesco.zanetti@economics.ox.ac.uk.

A. A Multi-Attribute Ship Clustering Model for Dark Ship Identification

Before detailing our ship clustering model, we present a high-level overview of its main components and their interconnections, as illustrated in Figure A.1.

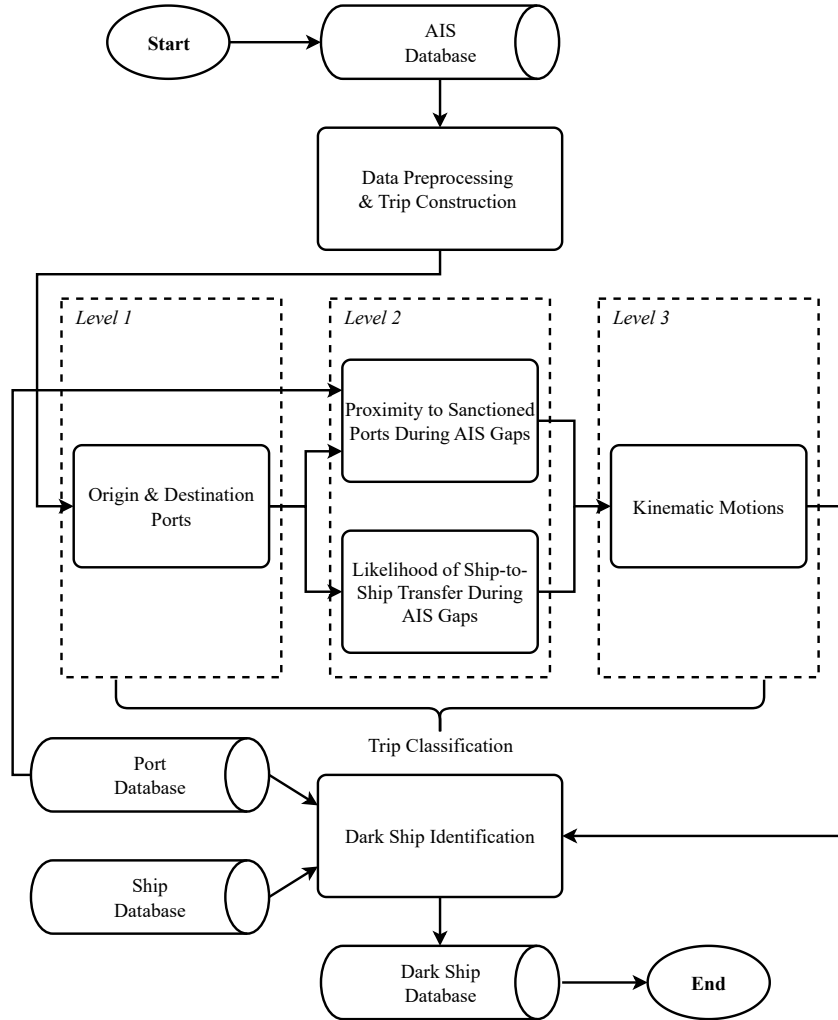


Figure A.1: Methodology Framework of the Multi-Attribute Ship Clustering Model

First, we preprocess raw AIS data from the global crude oil tanker fleet for the 2017–2023 sample period. This process involves removing abnormal data points and segmenting AIS data into trip-based units by analyzing vessel position, speed, and anomalies related to port entry and exit. It identifies trip start and end points and categorizes idle trips.

Second, we implement a three-level trip classification process to assess: (i) whether the trip’s origin and/or destination port is in a sanctioned country, (ii) whether the vessel is near a sanctioned-country port during a prolonged AIS data gap or near another ship with overlapping data gaps, potentially enabling an unauthorized ship-to-ship transfer, and (iii) whether the vessel’s kinematic movements exhibit abnormal patterns.

Third, using this information, we assign categorical suspicion scores to each trip and compute the average trip suspicion score per vessel, along with the idle trip ratio. We use these metrics, combined with three additional features extracted from ship and port databases – years of vessel usage, the number of vessels owned by its commercial operator, and its flag state ranking – to identify dark oil tankers.

Finally, we estimate a k -means clustering model to classify ships as either dark or white for each year in the sample period (MacQueen, 1967; Lloyd, 1982). The following sections provide further details.¹

A.1. Additional Data Sources

Our algorithm extracts and analyzes four key indicators: AIS signal gaps, vessel kinematics, intrinsic vessel characteristics, and geographical information. The goal is to identify and leverage indicators with strong predictive power for dark shipping. AIS signal gaps and vessel kinematics are derived from AIS data, our primary source, as detailed in Section 2.2 of the main text. To enhance robustness, we also incorporate vessel-specific and geographical data.

The vessel data used in the clustering model come primarily from the Seasearcher platform by Lloyd’s List Intelligence. We focus on two key attributes: vessel age and the number of vessels owned by its commercial operator. Additionally, we incorporate flag state rankings from the publicly available, annually updated Paris MoU List.² Port data, mainly sourced from Lloyd’s List, are essential for segmenting continuous AIS data into distinct trips for granular analysis. Specifically, we use three port attributes: name, geographic coordinates (longitude and latitude), and country.

¹The complete pseudocode and detailed technical specifications for our multi-attribute ship clustering model are available upon request.

²See <https://parismou.org/Statistics%26Current-Lists/white-grey-and-black-list> (accessed July 3, 2024).

A.2. Data Preprocessing and Trip Construction

We begin by preprocessing AIS data and segmenting them into consecutive trips for each tanker, which serves as the fundamental unit of analysis. AIS data provide continuous vessel tracking but contain noise from abnormal situations, such as idling in port, repairs, and maintenance. By dividing the AIS records into trip-based units and extracting relevant features, we can better capture vessel behavior and movement patterns (Yan et al., 2020; Li et al., 2022). This approach helps distinguish normal from abnormal trips, providing more granular input features for our clustering model.

Second, we filter out anomalous data points by removing entries where the vessel’s draft is less than zero or its speed exceeds 20 knots. This speed threshold is chosen because most tankers typically operate well below this limit under normal conditions (Adland et al., 2020). Speeds exceeding 20 knots likely indicate AIS errors or anomalous events that deviate from standard vessel operations.

Third, we determine port entry and exit states by analyzing speed changes. A vessel is flagged as entering a port if its speed drops from at least 1 knot to below 1 knot, and as leaving if its speed rises from below 1 knot to 1 knot or more. However, not all detected port entries and exits are valid. To retain only meaningful movements, we filter out data points near ports where speeds fluctuate due to congestion, loading, or unloading (e.g., a speed drop below 1 knot followed by a return to 1 knot or more). Additionally, if the time between consecutive port entries and exits is under three hours, we exclude these cases, as they likely reflect speed fluctuations rather than actual port activity.

To identify ports of entry and exit, our algorithm first determines the set of ports geographically close to a vessel’s entry point, as defined by speed variations. It does so by drawing a small geographical square centered on the entry point coordinates, with longitude and latitude variations of 2 degrees (approximately 200 kilometers, depending on latitude). Only ports within this square are considered valid entry and exit candidates. The algorithm then identifies the nearest port by computing the geographical distance (in kilometers) between the vessel’s entry point and each

candidate port using the Haversine formula:

$$\begin{aligned}
 D & \left[\left(lon_{j,t}^{Ship}, lat_{j,t}^{Ship} \right), \left(lon_p^{Port}, lat_p^{Port} \right) \right] \\
 & = 2R \arcsin \left[\sqrt{ \sin^2 \left(\frac{lat_{j,t}^{Ship} - lat_p^{Port}}{2} \right) + \cos \left(lat_{j,t}^{Ship} \right) \cos \left(lat_p^{Port} \right) \sin^2 \left(\frac{lon_{j,t}^{Ship} - lon_p^{Port}}{2} \right) } \right], \quad (A.1)
 \end{aligned}$$

where $\left(lon_{j,t}^{Ship}, lat_{j,t}^{Ship} \right)$ and $\left(lon_p^{Port}, lat_p^{Port} \right)$ are the coordinates of ship j at time t and port p in radians, respectively, and $R = 6,371$ kilometers is Earth’s radius.³ The algorithm also excludes cases where a vessel enters the same port consecutively, recording only the first entry as valid. This removes redundancies and ensures the focus remains on significant vessel movements.

Finally, the algorithm identifies and marks idle trips, defined as instances where a vessel’s average speed remains 1 knot or less for 14 consecutive days.⁴ These trips are flagged for further analysis.

A.3. Feature Extraction and Trip Classification

Our next task is to distinguish between normal and abnormal trips first and then categorize the abnormal ones accordingly. However, there is no generally accepted standard for defining normal and abnormal trips, as different studies use various methods, such as detecting intentional AIS shutdowns (Bernabé et al., 2024) or identifying deviations in vessel movement patterns (Rong et al., 2024), to determine abnormal ship behaviors.

Thus, we develop a three-level trip classification process that extends existing features in the literature – such as direct involvement in shipping sanctioned oil via loading or unloading in sanctioned-country ports and anomalies in kinematic movements – while also leveraging the duration of AIS data gaps and the vessel’s location when disabling its transceiver. Specifically, we introduce a novel method for detecting suspicious ship-to-ship transfers by analyzing the geographical coordinates and timestamps of two vessels before and after their respective data gaps, particularly when these gaps overlap. This approach relies solely on AIS data, eliminating the computational burden of methods that require combining AIS data with satellite imagery to

³If the distance between the nearest port and the vessel’s entry point is 50 kilometers or more, the entry is deemed invalid.

⁴According to Clarksons’ Shipping Intelligence Network, a vessel is classified as idle if it has not recorded an average speed above 1 knot for at least 14 days, is not listed as laid-up, under repair, or in storage, and has not undertaken voyage fixture activity in selected sectors. Additionally, it must not subsequently record an average speed above 1 knot for two or more consecutive days. We use the first criterion to identify idle trips, as the others require additional data. Our results remain robust when using a 7-day cutoff instead.

identify unauthorized transfers (Rodger and Guida, 2022; Androjna et al., 2024; Ballinger, 2024).

With this trip classification process, we assign categorical suspicion scores to each trip and compute an average suspicion score for each vessel. The following sections examine each level of the trip classification process in detail.⁵

Level 1: Origin and Destination Ports

The first level of the trip classification process identifies trips directly linked to sanctioned countries by analyzing both origin and destination ports. Our analysis focuses on five oil-sanctioned countries: Iran, Syria, North Korea, Venezuela, and Russia. These nations faced significant international sanctions restricting their crude oil and petroleum trade during the 2017–2023 sample period. While four of these countries are oil exporters, we also include North Korea in our clustering model due to Western sanctions on its oil imports. Table A.1 summarizes the sanctions on the five countries.

Table A.1: Overview of Oil Export/Import Sanctions on Key Countries

Country	Sanctioning Parties	Timing	Details
Iran	U.S., European Union, Canada, etc.	Since 1995, intensified in 2012 and 2018	Sanctions targeting oil exports due to nuclear program
Syria	U.S., European Union, Canada, etc.	Since 2011	Prohibition on oil exports due to the Syrian civil war
North Korea	U.N. Security Council, U.S., European Union, etc.	Since 2017	Caps on oil and petroleum product imports due to nuclear program
Venezuela	U.S., European Union, Canada, etc.	Since 2019, intensified in 2020	Sanctions targeting PDVSA and oil embargo to pressure the government
Russia	U.S., European Union, U.K., Canada, etc.	Since 2022	Restrictions on oil exports due to the invasion of Ukraine

Notes. We exclude sanctions on Russia prior to 2022, as they primarily targeted specific companies and individuals, restricting access to Western technology, financing, and investment for oil exploration and production (Chupilkin et al., 2024). In contrast, post-2022 sanctions directly targeted Russian oil exports through import bans and price caps designed to reduce oil revenue. Additionally, Table C.1 in Appendix C provides a detailed month-by-month breakdown of sanctions on Iran, Syria, Venezuela, and Russia.

⁵We avoid inputting all features from the three levels into a single clustering model to prevent overfitting and maintain interpretability. Including too many diverse features at once could obscure the relationship between specific behaviors and their associated risk factors. By classifying trips at different levels, we gain a better understanding of the contribution of each feature set, enabling a more targeted analysis of abnormal ship behaviors.

Our algorithm first identifies trips departing from ports in sanctioned countries. If a trip originates at such a port, it assigns a suspicion score (S^{Trip}) of 1 for the entire duration. This flags high-risk trips early. Next, the algorithm examines trips arriving at sanctioned-country ports. If a trip ends at such a port and has not already been flagged based on its origin, it also receives a suspicion score of 1. By analyzing both origin and destination ports, the algorithm ensures that all trips linked to sanctioned countries are captured, regardless of where they begin or end.

We assess the suspicion level of each trip only after sanctions have been imposed on the respective country, as there is no reason to suspect dark shipping of sanctioned oil before that point. For example, since Western sanctions on Venezuela began in 2019, trips to or from Venezuelan ports before 2019 are not considered suspicious.⁶

While the first level of the trip classification process identifies trips linked to sanctioned countries, such trips are not necessarily conducted by dark ships. A vessel classified as white can transport oil for sanctioned countries, provided it complies with sanction regulations. For example, the tanker *Sino Star* (IMO 9263693) called at St. Petersburg in early 2023 but adhered to the oil price cap regulations set by the Group of Seven (G7) nations. Despite its connections to vessels engaged in dark shipping, *Sino Star* was not classified as part of the dark fleet, as no evidence suggested deceptive or evasive practices in its operations ([Lloyd's List Intelligence, 2023](#)). This example highlights the need for additional criteria – such as AIS data gaps and anomalies in kinematic movements – to assess trip suspicion more accurately.

Level 2: AIS Data Gaps

The definition of dark ships varies across the literature, but two deceptive shipping practices outlined by the [U.S. Department of State \(2020\)](#) are widely recognized: (i) disabling a ship's AIS transceiver or manipulating its data to obscure movements (“AIS spoofing”) and (ii) engaging in ship-to-ship transfers, particularly in high-risk areas for sanctions violation or illegal activities. Although AIS spoofing is difficult to detect – a topic we explore later – our algorithm introduces a novel approach by assessing a vessel's proximity to a suspicious port during an AIS data gap and evaluating the likelihood of an unauthorized ship-to-ship transfer when two vessels are geographically close with overlapping data gaps. By integrating these key indicators, our method

⁶This consideration also applies when calculating the port-based trip suspicion score in the second level of the trip classification process. A vessel's proximity to a port in a country not yet sanctioned is not factored into the suspicion assessment, even if its AIS transceiver is turned off.

enhances the identification of potential evasion of sanctions.

Direct visit to a suspicious port. We first analyze a vessel’s proximity to ports in the five sanctioned countries during periods when its AIS transceiver is off. Such behavior often indicates potential involvement in transporting oil for sanctioned nations. Our algorithm detects AIS data gaps by computing time differences between consecutive AIS observations for each vessel. To focus on suspicious activity rather than routine operations (e.g., idling in port, repairs, or maintenance), it filters out data gaps occurring during port calls, isolating gaps more likely to signal unauthorized behavior.

Next, the algorithm identifies particularly long data gaps, defined as those exceeding the 99th percentile of time differences between consecutive AIS signals for each vessel. These gaps are marked by their start and end times, allowing the algorithm to focus on unaccounted periods that may involve the evasion of sanctions. After detecting prolonged data gaps, the algorithm calculates the geographical distance between the vessel and the nearest port in a sanctioned country listed in Table A.1. It applies the Haversine formula in Equation (A.1) for accuracy (see Figure 1a in the main text and Algorithm A.1 for illustrations).

The algorithm then calculates the required average speed for the vessel to reach and return from the suspicious port within the duration of the gap. The suspicion score for each long data gap is determined as one minus the percentile (divided by 100) of the required average speed within the historical speed distribution for all vessels in the sample year. More specifically, a lower required speed increases the plausibility that the vessel visited a suspicious port, resulting in a higher suspicion score.

Finally, the algorithm assigns a suspicion score (S^{Port}) to each trip, based on the highest suspicion score of any long data gap within that trip. This score estimates the likelihood that the vessel transported oil for a sanctioned country while its AIS transceiver was off.

Ship-to-ship transfer. Identifying potential unauthorized ship-to-ship transfers during AIS data gaps involves several steps. The algorithm first converts each ship’s heading, measured in degrees, into a vector.⁷ This conversion allows us to project the ship’s travel path based on its geographical coordinates and heading.

⁷In AIS data, the heading represents the direction of a vessel’s bow relative to true north (0°).

Algorithm A.1 Calculation of Port-Based Trip Suspicion Scores

Inputs: t_j : set of timestamps for ship $j \in \mathcal{J}$
 $\left\{ \left(lon_{j,t}^{Ship}, lat_{j,t}^{Ship} \right) \right\}_{t \in t_j}$: geographical coordinates of ship $j \in \mathcal{J}$
 $s_j = \{s_{j,t}\}_{t \in t_j}$: speeds of ship $j \in \mathcal{J}$
 $l_j^{Trip} = \{l_{j,t}^{Trip}\}_{t \in t_j}$: trip status labels for ship $j \in \mathcal{J}$
 $l_j^{Gap} = \{l_{j,t}^{Gap}\}_{t \in t_j}$: long data gap status labels for ship $j \in \mathcal{J}$
 $\left\{ \left(lon_p^{Port}, lat_p^{Port} \right) \right\}_{p \in \mathcal{P}}$: geographical coordinates of each port $p \in \mathcal{P}$
 $\mathcal{P}^{Sanction} \subset \mathcal{P}$: set of ports in sanctioned countries

Outputs: S_j^{Port} : set of port-based trip suspicion scores for each ship $j \in \mathcal{J}$

- 1: **function** PERCENTILE_OF_VALUE(s, \bar{s})
- 2: Sort s in ascending order
- 3: $n \leftarrow$ length of s
- 4: $count \leftarrow$ number of elements in s where $s \leq \bar{s}$
- 5: **return** $(count/n) \times 100$
- 6: **end function**
- 7: **function** CALCULATE_PORT_BASED_TRIP_SUSPICION_SCORES
- 8: **for all** $j \in \mathcal{J}$ **do**
- 9: **for all** long data gaps with start time t and end time t^\diamond in l_j^{Gap} **do**
- 10: Initialize $Dist \leftarrow \{\}$
- 11: **for all** $p \in \mathcal{P}^{Sanction}$ **do**
- 12: $D_{j,t,p} \leftarrow D \left[\left(lon_{j,t}^{Ship}, lat_{j,t}^{Ship} \right), \left(lon_p^{Port}, lat_p^{Port} \right) \right]$
- 13: $D_{j,t^\diamond,p} \leftarrow D \left[\left(lon_{j,t^\diamond}^{Ship}, lat_{j,t^\diamond}^{Ship} \right), \left(lon_p^{Port}, lat_p^{Port} \right) \right]$
- 14: $Dist_{j,p} \leftarrow D_{j,t,p} + D_{j,t^\diamond,p}$
- 15: $Dist \leftarrow Dist \cup \{Dist_{j,p}\}$
- 16: **end for**
- 17: $p^* \leftarrow \arg \min_{p \in \mathcal{P}^{Sanction}} Dist$
- 18: $\bar{s}_{j,t} \leftarrow Dist_{j,p^*} / \text{HOURS_DIFFERENCE}(t, t^\diamond)$
- 19: $S_{j,t} \leftarrow 1 - \text{PERCENTILE_OF_VALUE}(s_j, \bar{s}_{j,t}) / 100$
- 20: Store $S_{j,t}$
- 21: **end for**
- 22: **for all** trips with start time t and end time t^\S in l_j^{Trip} **do**
- 23: $S_{j,t^*} \leftarrow \max_{t' \in [t, t^\S]} S_{j,t'}$
- 24: **for all** $t' \in [t, t^\S]$ **do**
- 25: $S_{j,t'}^{Port} \leftarrow S_{j,t^*}$
- 26: **end for**
- 27: **end for**
- 28: **end for**
- 29: **end function**

Next, the algorithm examines pairs of ships with overlapping AIS gaps to determine whether their paths could have intersected. It calculates the intersection point of their projected paths using the geographical coordinates and headings recorded before their AIS signals disappeared (see Figure 1c in the main text and Algorithms A.2 and A.3 for illustrations). This intersection point is considered valid only if both vessels could have feasibly reached it based on their projected paths and speeds.

Algorithm A.2 Calculation of Ship-To-Ship Transfer-Based Trip Suspicion Scores

Inputs: t_j : set of timestamps for ship $j \in \mathcal{F}$
 $\left\{ \left(lon_{j,t}^{Ship}, lat_{j,t}^{Ship} \right) \right\}_{t \in t_j}$: geographical coordinates of ship $j \in \mathcal{F}$
 $h_j = \{h_{j,t}\}_{t \in t_j}$: headings of ship $j \in \mathcal{F}$
 $s_j = \{s_{j,t}\}_{t \in t_j}$: speeds of ship $j \in \mathcal{F}$
 $l_j^{Trip} = \{l_{j,t}^{Trip}\}_{t \in t_j}$: trip status labels for ship $j \in \mathcal{F}$
 $l_j^{Gap} = \{l_{j,t}^{Gap}\}_{t \in t_j}$: long data gap status labels for ship $j \in \mathcal{F}$
Outputs: S_j^{STS} : set of ship-to-ship transfer-based trip suspicion scores for each ship $j \in \mathcal{F}$

```

1: function CALCULATE_STS_BASED_TRIP_SUSPICION_SCORES
2:   for all  $j \in \mathcal{F}$  do
3:     for all long data gaps with start time  $t$  and end time  $t^\circ$  in  $l_j^{Gap}$  do
4:       Initialize  $\bar{s} \leftarrow \{\}$ 
5:       for all  $j' \in \mathcal{F} \setminus \{j\}$  do
6:         for all long data gaps with start time  $t^\dagger$  and end time  $t^\ddagger$  in  $l_{j'}^{Gap}$  do
7:           if  $t < t^\ddagger$  and  $t^\circ > t^\dagger$  then
8:             Compute intersection point  $(lon^{Intersect}, lat^{Intersect})$  of paths using positions at times  $t$ 
              and  $t^\dagger$ , and headings  $h_{j,t}$  and  $h_{j',t^\dagger}$   $\triangleright$  See Algorithm A.3
9:             if intersection point exists then
10:               $t^{mid} \leftarrow (\max\{t, t^\dagger\} + \min\{t^\circ, t^\ddagger\}) / 2$ 
11:              Compute distances:
12:               $D_{j,t,Intersect} \leftarrow D \left[ \left( lon_{j,t}^{Ship}, lat_{j,t}^{Ship} \right), \left( lon^{Intersect}, lat^{Intersect} \right) \right]$ 
13:               $D_{j,t^\circ,Intersect} \leftarrow D \left[ \left( lon_{j,t^\circ}^{Ship}, lat_{j,t^\circ}^{Ship} \right), \left( lon^{Intersect}, lat^{Intersect} \right) \right]$ 
14:              Compute required average speed:
15:               $\bar{s}_{j,j'} \leftarrow \max \{ D_{j,t,Intersect} / (t^{mid} - t), D_{j,t^\circ,Intersect} / (t^\circ - t^{mid}) \}$ 
16:               $\bar{s} \leftarrow \bar{s} \cup \{ \bar{s}_{j,j'} \}$ 
17:            end if
18:          end if
19:        end for
20:      end for
21:      if  $\bar{s} \neq \emptyset$  then
22:         $\bar{s}_{min} \leftarrow \min \bar{s}$ 
23:         $S_{j,t} \leftarrow 1 - \text{PERCENTILE\_OF\_VALUE}(s_j, \bar{s}_{min}) / 100$ 
24:        Store  $S_{j,t}$ 
25:      end if
26:    end for
27:    for all trips with start time  $t$  and end time  $t^{\S}$  in  $l_j^{Trip}$  do
28:       $S_{j,t^*} \leftarrow \max_{t' \in [t, t^{\S}]} S_{j,t'}$ 
29:      for all  $t' \in [t, t^{\S}]$  do
30:         $S_{j,t'}^{STS} \leftarrow S_{j,t^*}$ 
31:      end for
32:    end for
33:  end for
34: end function

```

To refine the analysis, the algorithm calculates the midpoint in time between overlapping AIS gaps for each ship pair. This midpoint helps determine the required average speed for each vessel to reach the intersection point from the location where its AIS signal disappeared, potentially

Algorithm A.3 Calculation of Intersection Point of Paths

Inputs: t_j : set of timestamps for ship $j \in \mathcal{J}$ $\left\{ \left(lon_{j,t}^{Ship}, lat_{j,t}^{Ship} \right) \right\}_{t \in t_j}$: geographical coordinates of ship $j \in \mathcal{J}$ $h_j = \{h_{j,t}\}_{t \in t_j}$: headings of ship $j \in \mathcal{J}$ **Outputs:** $\left(lon_{j,j'}^{Intersect}, lat_{j,j'}^{Intersect} \right)$: intersection point of paths for ships $j, j' \in \mathcal{J}$

```
1: function HEADING_TO_VECTOR( $h_{j,t}$ )
2:    $\theta \leftarrow h_{j,t} \cdot \pi / 180$ 
3:    $dx \leftarrow \sin(\theta)$ 
4:    $dy \leftarrow \cos(\theta)$ 
5:   return ( $dx, dy$ )
6: end function
7: function FIND_INTERSECTION( $\left( lon_{j,t}^{Ship}, lat_{j,t}^{Ship} \right), \left( lon_{j',t'}^{Ship}, lat_{j',t'}^{Ship} \right), h_{j,t}, h_{j',t'}$ )
8:    $\left( dx_j, dy_j \right) \leftarrow$  HEADING_TO_VECTOR( $h_{j,t}$ )
9:    $\left( dx_{j'}, dy_{j'} \right) \leftarrow$  HEADING_TO_VECTOR( $h_{j',t'}$ )
10:   $\Delta t \leftarrow \left[ \left( lat_{j',t'}^{Ship} - lat_{j,t}^{Ship} \right) \cdot dx_{j'} - \left( lon_{j',t'}^{Ship} - lon_{j,t}^{Ship} \right) \cdot dy_{j'} \right] / \left( dx_{j'} \cdot dy_j - dy_{j'} \cdot dx_j \right)$ 
11:   $\Delta t' \leftarrow \left[ \left( lat_{j',t'}^{Ship} - lat_{j,t}^{Ship} \right) \cdot dx_j - \left( lon_{j',t'}^{Ship} - lon_{j,t}^{Ship} \right) \cdot dy_j \right] / \left( dx_{j'} \cdot dy_j - dy_{j'} \cdot dx_j \right)$ 
12:  if  $\Delta t \geq 0$  and  $\Delta t' \geq 0$  then ▷ Check if the intersection point is within a valid range
13:     $lon_{j,j'}^{Intersect} \leftarrow lon_{j,t}^{Ship} + \Delta t \cdot dx_j$ 
14:     $lat_{j,j'}^{Intersect} \leftarrow lat_{j,t}^{Ship} + \Delta t \cdot dy_j$ 
15:    return  $\left( lon_{j,j'}^{Intersect}, lat_{j,j'}^{Intersect} \right)$ 
16:  end if
17: end function
```

conduct an unauthorized ship-to-ship transfer involving sanctioned oil, and return to the point where its AIS signal reappears.⁸

The suspicion score is calculated as one minus the percentile (divided by 100) of the required average speed within the historical speed distribution for all vessels in the sample year. This score assesses the plausibility of the transfer: the lower the required speed, the more likely the vessels completed an unauthorized ship-to-ship transfer, resulting in a higher suspicion score.

Finally, the algorithm assigns a suspicion score (S^{STS}) to each trip, based on the highest score among overlapping AIS gaps within that trip. This score reflects the likelihood that the vessel engaged in a suspicious ship-to-ship transfer of sanctioned oil while its AIS transceiver was off.

K -means clustering of trips. After normalizing the trip suspicion scores for both port-based and ship-to-ship transfer-based assessments, our algorithm applies k -means clustering to categorize trips. As outlined in Algorithm A.4, the process begins with randomly initializing centroids for two

⁸The algorithm assumes ships travel in straight lines and at constant speeds during AIS data gaps, which may not always reflect actual movements due to currents and operational adjustments. While it effectively identifies potential intersection points, it relies on the accuracy of pre-gap data and does not account for real-time deviations. Incorporating models for path deviations or additional data sources (e.g., satellite imagery) could improve detection accuracy. We leave these enhancements for future research.

clusters. Each trip is then assigned to the cluster with the closest centroid, based on the squared Euclidean distance between the trip’s normalized suspicion scores and the centroids. Once all trips are assigned, the algorithm recalculates the centroids based on current cluster memberships. This iterative process continues until the centroids stabilize, indicating that further iterations do not change the cluster assignments.

Algorithm A.4 K -Means Clustering of Trips

Inputs: t_j : set of timestamps for ship $j \in \mathcal{J}$

$l_j^{Trip} = \{l_{j,t}^{Trip}\}_{t \in t_j}$: trip status labels for ship $j \in \mathcal{J}$

$S_j^{Port} = \{S_{j,t}^{Port}\}_{t \in t_j}$: port-based trip suspicion scores for ship $j \in \mathcal{J}$

$S_j^{STS} = \{S_{j,t}^{STS}\}_{t \in t_j}$: ship-to-ship transfer-based trip suspicion scores for ship $j \in \mathcal{J}$

Outputs: $\{\mu_1^{Gap}, \mu_2^{Gap}\}$: cluster centroids

$\{c_{j,t}^{Gap} \mid j \in \mathcal{J}, t \in t_j, l_{j,t}^{Trip} = \text{Trip start}\}$: cluster assignments

```

1: function NORMALIZE( $\mathcal{X}$ )
2:    $\underline{x} \leftarrow \min \mathcal{X}$ 
3:    $\bar{x} \leftarrow \max \mathcal{X}$ 
4:   for  $x_n \in \mathcal{X}$  do
5:      $x_n^{norm} \leftarrow (x_n - \underline{x}) / (\bar{x} - \underline{x})$ 
6:   end for
7:   return  $\{x_n^{norm}\}_{n=1}^{|\mathcal{X}|}$ 
8: end function
9: function K-MEANS( $\mathcal{X}, K, I$ )
10:  Randomly initialize centroids  $\mu_1, \dots, \mu_K$ 
11:  for  $i = 1$  to  $I$  do
12:    for  $k = 1$  to  $K$  do
13:       $\mathcal{C}_k \leftarrow \{\}$ 
14:    end for
15:    for all  $x_n \in \mathcal{X}$  do
16:       $k^* \leftarrow \arg \min_k \|x_n - \mu_k\|^2$ 
17:       $\mathcal{C}_{k^*} \leftarrow \mathcal{C}_{k^*} \cup \{x_n\}$ 
18:       $c_n \leftarrow k^*$ 
19:    end for
20:     $\mu_{old} \leftarrow (\mu_1, \dots, \mu_K)$ 
21:    for  $k = 1$  to  $K$  do
22:      if  $|\mathcal{C}_k| > 0$  then
23:         $\mu_k \leftarrow \sum_{x_n \in \mathcal{C}_k} x_n / |\mathcal{C}_k|$ 
24:      end if
25:    end for
26:    if  $\mu_k = \mu_{old}$  for all  $k$  then
27:      break
28:    end if
29:  end for
30:  return  $\{\mu_k\}_{k=1}^K, \{c_n\}_{n=1}^{|\mathcal{X}|}$ 
31: end function
32:  $S^{Port} \leftarrow \text{NORMALIZE} \left( \left\{ S_{j,t}^{Port} \mid j \in \mathcal{J}, t \in t_j, l_{j,t}^{Trip} = \text{Trip start} \right\} \right)$ 
33:  $S^{STS} \leftarrow \text{NORMALIZE} \left( \left\{ S_{j,t}^{STS} \mid j \in \mathcal{J}, t \in t_j, l_{j,t}^{Trip} = \text{Trip start} \right\} \right)$ 
34:  $\{\mu_1^{Gap}, \mu_2^{Gap}\}, \{c_{j,t}^{Gap} \mid j \in \mathcal{J}, t \in t_j, l_{j,t}^{Trip} = \text{Trip start}\} \leftarrow \text{K-MEANS}(\{S^{Port}, S^{STS}\}, 2, 300)$ 

```

After clustering, the final assignments determine whether a trip is classified as suspicious based on its port-based and ship-to-ship transfer-based suspicion scores (S^{Port} and S^{STS}). Trips with higher values for both scores are grouped together and assigned a clustering label (c^{Gap}) of 1, indicating a greater likelihood of suspicious activity when the vessel turns off its AIS transceiver.

Level 3: Kinematic Movements

In the third level of the trip classification process, our algorithm analyzes anomalies in a ship’s kinematic movements to identify suspicious behavior. The process begins by calculating three key metrics for each trip: the average speed, the standard deviation of the speed, and the detour factor. These metrics are widely used in the literature for detecting and classifying abnormal ship behaviors (Ristic et al., 2008; Rong et al., 2024). They provide insights into the vessel’s movement patterns, allowing for the detection of irregularities that may indicate evasive maneuvers or other suspicious activities.⁹

First, the algorithm calculates the average speed of the vessel during each trip by summing the vessel’s recorded speeds and dividing by the number of speed measurements. Lower average speeds may indicate dark shipping activities because vessels involved in covert operations often reduce their speed to avoid detection or may spend prolonged periods idling, either to wait for favorable conditions or to engage in unauthorized ship-to-ship transfers.

Next, the algorithm calculates the standard deviation of the ship’s speed, which measures the variability in speed. A high standard deviation may suggest erratic behavior, potentially indicating attempts to avoid detection or engage in unauthorized activities. In addition to these speed metrics, the algorithm computes the detour factor (DF^{Trip}), which compares the actual path taken by the vessel to the direct path between the trip’s start and end points, as defined in Equation (A.2):

$$DF_j^{Trip} = \frac{\sum_{n=1}^{N-1} D \left[\left(lon_{j,n}^{Ship}, lat_{j,n}^{Ship} \right), \left(lon_{j,n+1}^{Ship}, lat_{j,n+1}^{Ship} \right) \right]}{D \left[\left(lon_{j,1}^{Ship}, lat_{j,1}^{Ship} \right), \left(lon_{j,N}^{Ship}, lat_{j,N}^{Ship} \right) \right]}, \quad (\text{A.2})$$

⁹Incorporating these features can also be useful in detecting AIS spoofing because spoofed AIS signals often show inconsistencies in a vessel’s movement patterns (Triebert et al., 2023; Lloyd’s List Intelligence, 2024a,b). For instance, abnormal speed variations or unusual detours that deviate from typical shipping routes can indicate that a vessel’s AIS data are being manipulated. These features help flag suspicious behavior by comparing actual movements with expected norms, enhancing the detection of spoofing attempts.

where $(lon_{j,n}^{Ship}, lat_{j,n}^{Ship})$ are the n -th coordinates of ship j in radians, and the Haversine formula in Equation (A.1) is applied to calculate the geographical distance between each pair of coordinates. A detour factor greater than one indicates that the vessel deviated from a direct route, which could be indicative of deceptive shipping practices (U.S. Department of State, 2020; Rong et al., 2024).

After computing these metrics, the algorithm normalizes the values to ensure they are on a comparable scale. The normalized metrics are then input into a k -means clustering algorithm, which groups the trips into two clusters: those with typical kinematic behavior and those with anomalous kinematic behavior. The algorithm iteratively refines these clusters until the centroids stabilize, ensuring accurate classification of each trip.

Subsequently, the algorithm reassigns the clustering labels ($c^{Kinematic}$) to reflect the level of suspicion associated with each trip based on kinematic movements. Trips characterized by lower average speeds, higher speed variability, and significant detours are grouped together and re-assigned a clustering label of 1, indicating the vessel’s potential involvement in dark shipping activities.

Finally, the clustering results from the second and third levels of the trip classification process enable us to assign different categorical suspicion scores (S^{Trip}) to each trip based on anomalies identified in both significant data gaps and kinematic movements. For each vessel, the algorithm assesses whether any of its trips exhibit anomalies in both prolonged AIS data gaps and kinematic movements.¹⁰ If both levels of clustering indicate abnormalities (i.e., $c^{Gap} = 1$ and $c^{Kinematic} = 1$), the trip is assigned a suspicion score of 1, indicating a high likelihood of dark shipping activities. Conversely, if no anomalies are detected in either level, the trip receives a score of 0, marking it as a normal trip. In cases where only one of the two levels flags an anomaly, a moderate suspicion score of 0.5 is assigned, reflecting partial abnormality.

After assigning suspicion scores to each trip, the algorithm calculates the average trip suspicion score (\bar{S}^{Trip}) for each vessel. This is achieved by summing the suspicion scores across all trips taken by a vessel and dividing the total by the number of trips. The resulting average score provides a comprehensive measure of the vessel’s overall involvement in the dark shipping of sanctioned oil.

¹⁰We only assign trip suspicion scores to previously unassigned trips. Recall that in the first level of the trip classification process, trips that start and/or end at a port located in a sanctioned country have already been assigned a suspicion score (S^{Trip}) of 1.

A.4. Dark Ship Identification

The three-level trip classification process provides a comprehensive analysis of vessel travel patterns and movement dynamics. The resulting average trip suspicion score and idle trip ratio – defined as the fraction of trips classified as idle – serve as key indicators for identifying dark ships.¹¹

Additional vessel-specific indicators must be incorporated to fully capture vessel characteristics. As discussed, we consider three key features: the vessel’s years in service, the number of vessels owned by its commercial operator, and its flag state ranking, as outlined in the Paris MoU List ([Paris MoU, 2018](#)).

Older vessels are more likely to engage in dark shipping, as they are typically near the end of their service life and hold less value for their operators ([Miller, 2023](#)). This lower financial risk makes them attractive for unauthorized activities, as seizure or decommissioning carries fewer economic consequences. Additionally, vessels owned by smaller commercial operators may be more prone to dark shipping due to reduced regulatory scrutiny and limited compliance resources ([Lloyd’s List Intelligence, 2023](#)). In contrast, larger operators face greater reputational and financial risks, discouraging their involvement.

A vessel’s flag state ranking also provides insight into its regulatory environment, as ships registered under higher-risk flags are more likely to engage in non-compliant activities, as discussed in [Section 2.4](#) of the main text.

After normalizing the five selected features, our algorithm applies the k -means clustering method to classify vessels into two groups: those with a higher likelihood of dark shipping and those that are less suspicious. The algorithm then performs a final reassignment of clustering labels, adjusting the weighting of the number of vessels owned by an operator. A smaller fleet size is associated with higher suspicion, ensuring that vessels owned by smaller operators – who may face less scrutiny or have fewer compliance resources – are appropriately flagged as higher risk. Finally, vessels in the high-suspicion cluster are labeled as dark ships, while those in the lower-risk cluster are labeled as white ships.

¹¹The idle trip ratio functions similarly to average vessel speed in detecting anomalies in kinematic movements. Dark ships often remain idle for extended periods to facilitate unauthorized ship-to-ship transfers, evade detection, or operate beyond regulatory oversight.

B. Additional Dark Ship Identification Results

In this appendix, we present Figures B.1 to B.4, which illustrate the concentration of direct vessel visits to suspicious ports in Venezuela, Iran, and Russia near the Black Sea and Baltic Sea during periods not covered in the main text.

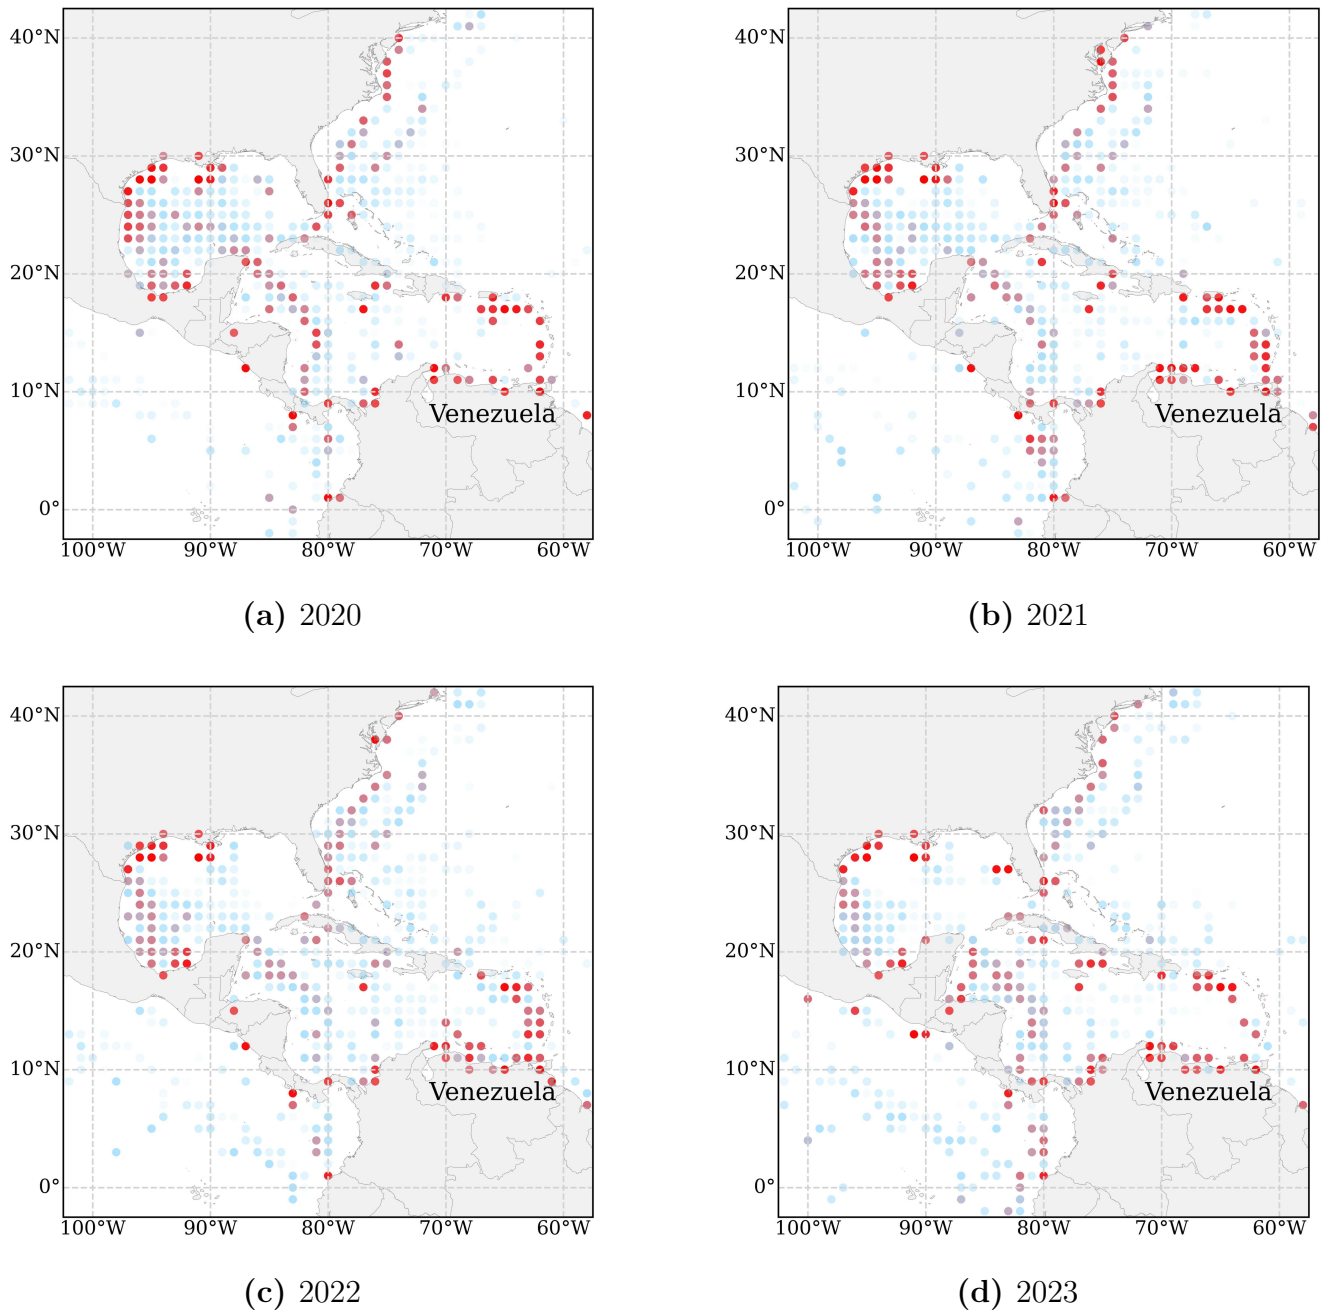




Figure B.1: Geographical Concentration of Dark Shipping Activities Near Venezuela: 2020–2023


Notes. The figures illustrate the geographical concentration of dark shipping activities near Venezuela from 2020 to 2023. Colored points and chromatic intensity follow the same format as in Figure 4 of the main text.



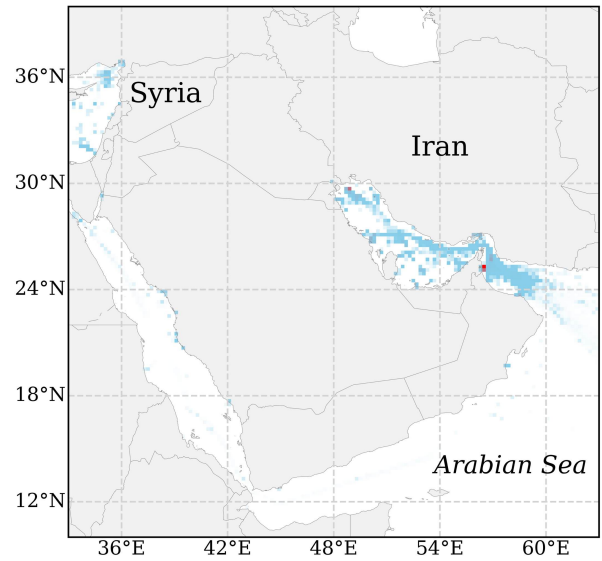
(a) 2017



(b) 2018



(c) 2019



(d) 2020

Figure B.2: Geographical Concentration of Dark Shipping Activities Near Iran: 2017–2020

Notes. The figures show the geographical concentration of dark shipping activities near Iran from 2017 to 2020, following the same format as Figure 5 in the main text.




Figure B.3: Geographical Concentration of Dark Shipping Activities in the Black Sea: 2017–2020

Notes. The figures illustrate the geographical concentration of dark shipping activities in the Black Sea near Russia from 2017 to 2020. The colored lines and chromatic intensity follow the same format as in Figure 6 of the main text.




Figure B.4: Geographical Concentration of Dark Shipping Activities in the Baltic Sea: 2017–2020

Notes. The figures show the geographical concentration of dark shipping activities in the Baltic Sea near Russia from 2017 to 2020. The colored lines and chromatic intensity follow the same format as in Figure 7 of the main text.

Figure B.5 shows the concentration of suspicious ship-to-ship transfers near West Africa and Gibraltar for years not reported in the main text.




Figure B.5: Geographical Concentration of Dark Shipping Activities Near the Strait of Gibraltar and the Gulf of Guinea: 2017–2020

Notes. The figures show the geographical concentration of dark shipping activities near the Strait of Gibraltar and the Gulf of Guinea during the years 2017–2020. The colored points and chromatic intensity follow the same format as in Figure 8 of the main text.

C. Details of Oil Sanctions on Iran, Syria, Venezuela, and Russia in 2017–2023

Table C.1: Details of Oil-Related Sanctions on Iran, Syria, Venezuela, and Russia During 2017–2023

Date	Sanctioning Parties	Sanctioned Party	Brief Details of the Sanction
May-17	European Union	Syria	Extended oil embargo and sanctions against Syria until June 2018 due to ongoing repression.
May-18	European Union	Syria	Extended oil embargo and sanctions against Syria until June 2019.
May-18	United States	Iran	Announced withdrawal from JCPOA; planned to reimpose sanctions on Iran’s oil exports and energy sector.
Nov-18	United States	Iran	Re-imposed sanctions targeting Iran’s oil exports, banking, and shipping; aimed to reduce Iran’s oil revenue.
Nov-18	United States	Syria	Sanctioned entities involved in transporting Iranian oil to Syria.
Jan-19	United States	Venezuela	Imposed sanctions on PDVSA to cut off Venezuela’s main revenue source.
Apr-19	United States	Venezuela	Expanded sanctions to entities transporting Venezuelan oil to Cuba; imposed restrictions on vessels and shipping companies.
May-19	European Union	Syria	Renewed oil embargo and sanctions against Syria until June 2020.
May-19	United States	Iran	Ended SREs allowing import of Iranian oil without sanctions; aimed to reduce Iran’s oil exports to zero.
May-20	European Union	Syria	Extended oil embargo and sanctions against Syria until June 2021.
Jun-20	United States	Syria	Implemented Caesar Act; expanded sanctions on Syrian government and oil sector; targeted foreign entities aiding Syria’s oil industry.
Aug-20	United States	Iran	Sanctioned entities facilitating oil shipments from Iran; aimed at disrupting Iran’s oil exports.
Aug-20	United States	Syria	Sanctioned entities receiving oil shipments from Iran; intended to cut off oil supplies to Syria.
Mar-21	United States	Iran	Imposed sanctions on entities involved in sale and transport of Iranian petrochemical products, including companies in China and UAE.
May-21	European Union	Syria	Extended oil embargo and sanctions against Syria until June 2022.
Mar-22	United States	Russia	Banned imports of Russian crude oil, LNG, and coal in response to invasion of Ukraine.
Mar-22	United Kingdom	Russia	Announced phase-out of Russian oil imports by end of 2022.
May-22	European Union	Syria	Extended oil embargo and sanctions against Syria until June 2023.
Jun-22	United States	Russia	Imposed sanctions on Russian maritime entities to restrict transport of oil.
Jun-22	Canada	Russia	Banned imports of certain petroleum products; prohibited export of goods related to oil exploration and production.
Jun-22	European Union	Russia	Implemented ban on imports of Russian crude oil and refined petroleum products, with limited exceptions.
Oct-22	European Union	Russia	Adopted 8th sanctions package, including price cap on maritime transport of Russian oil.
Dec-22	EU, G7, Australia	Russia	Agreed on price cap of \$60 per barrel on Russian crude oil exports effective from 5 Dec 2022.
Jan-23	European Union	Russia	Prolonged economic sanctions, including ban on import or transfer of seaborne crude oil and certain petroleum products from Russia to EU.
Feb-23	EU, G7, Australia	Russia	Extended price cap mechanism to include Russian refined petroleum products, setting separate price caps.
Mar-23	United States	Iran	Sanctioned companies facilitating sale and shipment of Iranian petrochemical products and petroleum.
May-23	European Union	Syria	Extended oil embargo and sanctions against Syria until June 2024.
Jun-23	European Union	Russia	Adopted 11th sanctions package; prohibited transit of goods via Russia and tightened export restrictions on crude oil and petroleum products.
Sep-23	United States	Russia	Imposed sanctions on additional Russian energy companies and individuals involved in the energy sector.
Nov-23	United States	Iran	Escalated sanctions enforcement on illegal Iranian oil exports; lawmakers introduced ‘SHIP Act’ to strengthen measures.
Dec-23	United States	Russia	Imposed further sanctions on entities and vessels violating oil price cap, especially those engaging in deceptive practices.
Dec-23	European Union	Russia	Adopted 12th sanctions package; enforced oil price cap and prohibited import of liquefied propane from Russia.

Notes. In constructing the sanction intensity index, we treat the European Union (EU) as a single sanctioning entity and count the number of sanctions it imposed accordingly. For the \$60 per barrel price cap on Russian crude oil exports introduced in December 2022 and the extended price cap mechanism in February 2023, the sanctioning parties were the EU, Canada, Japan, the United Kingdom, the United States, and Australia. Thus, for these two sanctions, the sanction intensity is assigned a value of 6.

D. Data Sources for LP Estimation

Table D.1 provides details on the data used in the baseline estimation and propagation channel analysis, including sources and any construction or adjustment processes. Additional data for robustness checks are outlined in the following appendices.

Table D.1: Data Sources and Description

Variable	Mnemonic/Series Key	Source	Notes on Construction/Adjustment
Baseline Estimation			
Oil Sanction Intensity Index	N/A	Various Sources	Constructed from a dataset of oil-related sanctions (Table A.1).
World Seaborne Oil Exports on Record	N/A	UN Comtrade	Extracted using HS code 2709 and manually seasonally adjusted.
Oil Exports by Dark Ships	N/A	AIS	Derived from dark ship identification and manually seasonally adjusted.
WTI Spot Price	WTISPLC	FRED	Manually seasonally adjusted.
U.S. PPI Total	PPIFIS	FRED	Raw series obtained directly from FRED.
U.S. IP Total	INDPRO	FRED	Raw series obtained directly from FRED.
Brent Spot Price	MCOILBRETEU	FRED	Manually seasonally adjusted.
EU PPI Total	STS.M.I9.N.PRON.NS0020.4.000	Eurostat	Manually seasonally adjusted.
EU IP Total	STS.M.I9.Y.PROD.NS0020.4.000	Eurostat	Raw series obtained directly from Eurostat.
Propagation Channels			
U.S. PPI Energy	PPIDES	FRED	Raw series obtained directly from FRED.
U.S. IP Energy	IPB50089S	FRED	Raw series obtained directly from FRED.
U.S. PPI Total Exl. Energy	WPSFD49107	FRED	Raw series obtained directly from FRED.
U.S. IP Total Exl. Energy	IPX5001ES	FRED	Raw series obtained directly from FRED.
China Oil Import Price	N/A	National Bureau of Statistics of China	Manually seasonally adjusted.
U.S. Import Price for China, NAICS 31	COCHNZ31	FRED	Manually seasonally adjusted.
U.S. Import Price for China, NAICS 32	COCHNZ32	FRED	Manually seasonally adjusted.
U.S. Import Price for China, NAICS 33	COCHNZ33	FRED	Manually seasonally adjusted.
U.S. Crude Oil Exports	N/A	U.S. Energy Information Administration	Manually seasonally adjusted.
EU PPI Energy	STS.M.I9.N.PRON.NS0090.4.000	Eurostat	Manually seasonally adjusted.
EU IP Energy	STS.M.I9.Y.PROD.NS0090.4.000	Eurostat	Raw series obtained directly from Eurostat.
EU PPI Total Exl. Energy	STS.M.I9.N.PRON.NS0021.4.000	Eurostat	Manually seasonally adjusted.
EU IP Total Exl. Energy	STS.M.I9.Y.PROD.NS0021.4.000	FRED	Raw series obtained directly from Eurostat.
China IP Total	N/A	National Bureau of Statistics of China	Constructed from the month-on-month IP growth rate.
China Total Import Value From EU	N/A	General Administration of Customs of China	Manually seasonally adjusted.

E. Robustness of LP Results

E.1. Baseline Results

Figures E.1 and E.2 display the IRFs to an unexpected oil sanction intensity shock for the U.S. and EU economies, incorporating four lags of the endogenous variables in the LP model (3). Additionally, Figures E.3 and E.4 present the IRFs from the LP model (3) with a constant and a linear trend. The results remain robust across these alternative specifications.

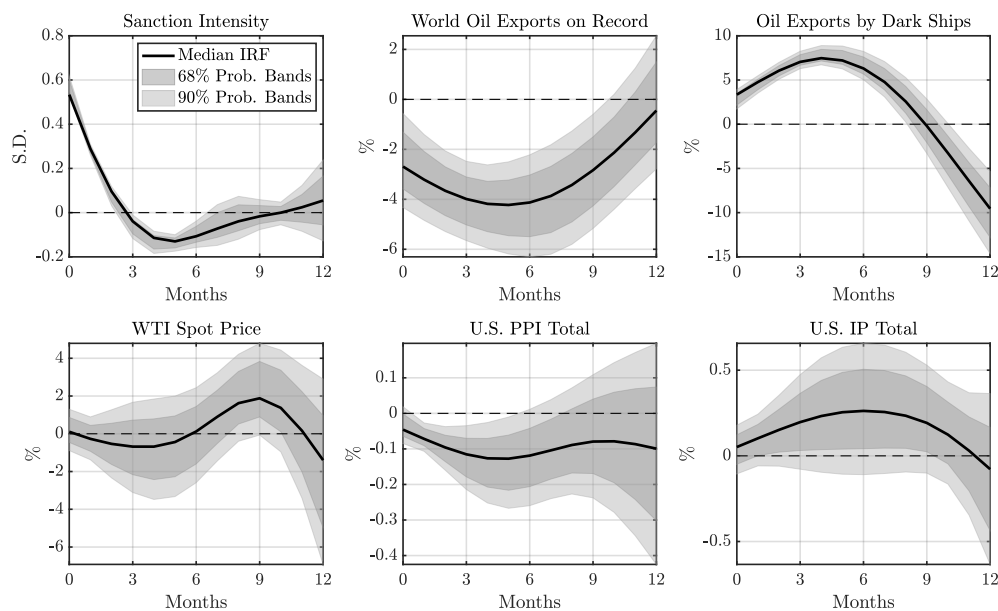


Figure E.1: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: $L = 4$

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using the SLP method (Barnichon and Brownlees, 2019). Apart from incorporating four lags of the endogenous variables, the estimation follows the same specification as in Figure 12. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.




Figure E.2: IRFs to an Oil Sanction Intensity Shock for the EU Economy: $L = 4$

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the EU economy are computed using the SLP method (Barnichon and Brownlees, 2019). Apart from incorporating four lags of the endogenous variables, the estimation follows the same specification as in Figure 13. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

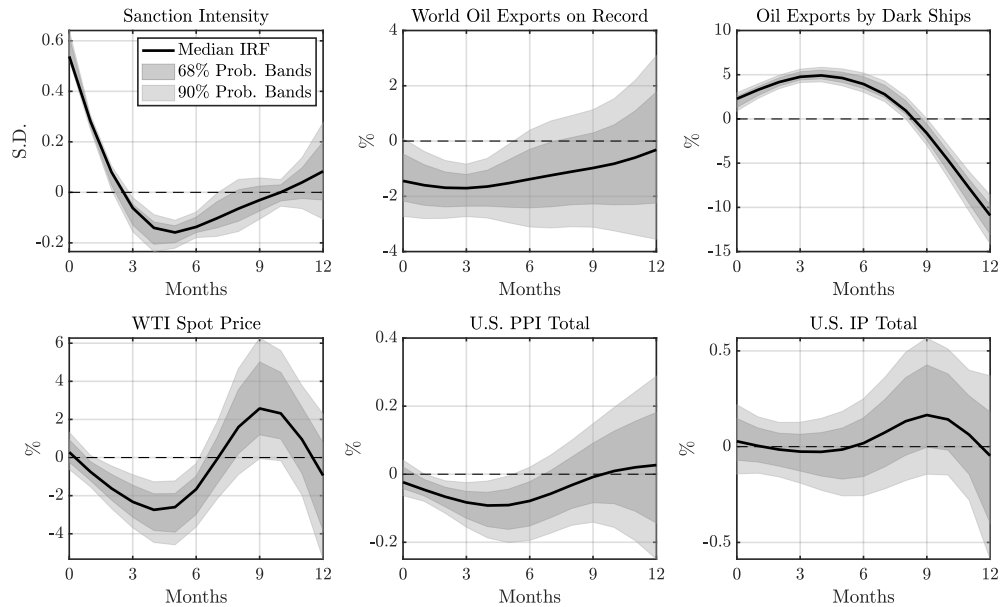


Figure E.3: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: Linear Trend

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using the SLP method (Barnichon and Brownlees, 2019). Apart from incorporating both a constant and a linear trend, the estimation matches the specification used in Figure 12. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.




Figure E.4: IRFs to an Oil Sanction Intensity Shock for the EU Economy: Linear Trend

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the EU economy are computed using the SLP method (Barnichon and Brownlees, 2019). Apart from incorporating both a constant and a linear trend, the estimation matches the specification used in Figure 13. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

Figures E.5 and E.6 present the IRFs to an unexpected oil sanction intensity shock for the U.S. and EU economies, replacing the WTI spot price with the U.S. crude oil import price and the Brent spot price with the EU crude oil import price. The U.S. import price is sourced from FRED (IR10000), and the EU import price from Eurostat (STS.M.I9.N.IMPR.2B0610.4.000), both of which are manually seasonally adjusted. The results remain robust.




Figure E.5: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: Oil Import Price

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using the SLP method (Barnichon and Brownlees, 2019). Apart from replacing the WTI spot price with the U.S. crude oil import price, the estimation specification remains identical to that used for Figure 12. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.




Figure E.6: IRFs to an Oil Sanction Intensity Shock for the EU Economy: Oil Import Price

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the EU economy are computed using the SLP method (Barnichon and Brownlees, 2019). Apart from replacing the Brent spot price with the EU crude oil import price, the estimation specification remains identical to that used for Figure 13. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

Figures E.7 and E.8 present the IRFs to an unexpected oil sanction intensity shock for the U.S. and EU economies, identified using a Cholesky-ordered SVAR model. The endogenous variables are ordered as follows: the oil sanction intensity index, recorded world seaborne crude oil exports, crude oil exports from sanctioned countries transported by dark ships, the WTI spot price, U.S./EU IP, and U.S./EU PPI.

Our timing assumption holds that an oil sanction intensity shock immediately affects all subsequent variables. Conversely, shocks to recorded world seaborne and dark-shipped oil exports influence only later variables, without affecting the sanction intensity index contemporaneously.

Overall, the IRFs are less precisely estimated for both economies. For the U.S., the median responses of the endogenous variables remain largely unchanged from Figure 12. In contrast, for the EU, the median Brent spot price response is positive for the first three months after the shock, while the total PPI response remains consistently negative across forecast horizons. However, none of these differences are statistically significant.

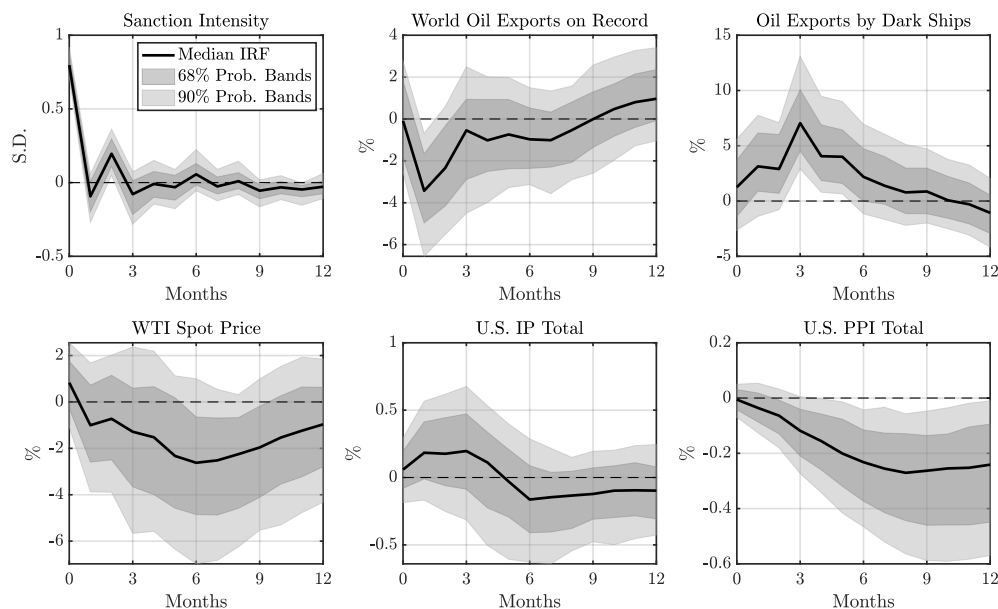


Figure E.7: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: Cholesky-Ordered SVAR

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using a Cholesky-ordered SVAR model with three lags. The ordering of endogenous variables follows Figure 12, except that total PPI and IP are switched. Black solid lines represent the median responses, while the gray-shaded areas indicate the 68% and 90% confidence bands.

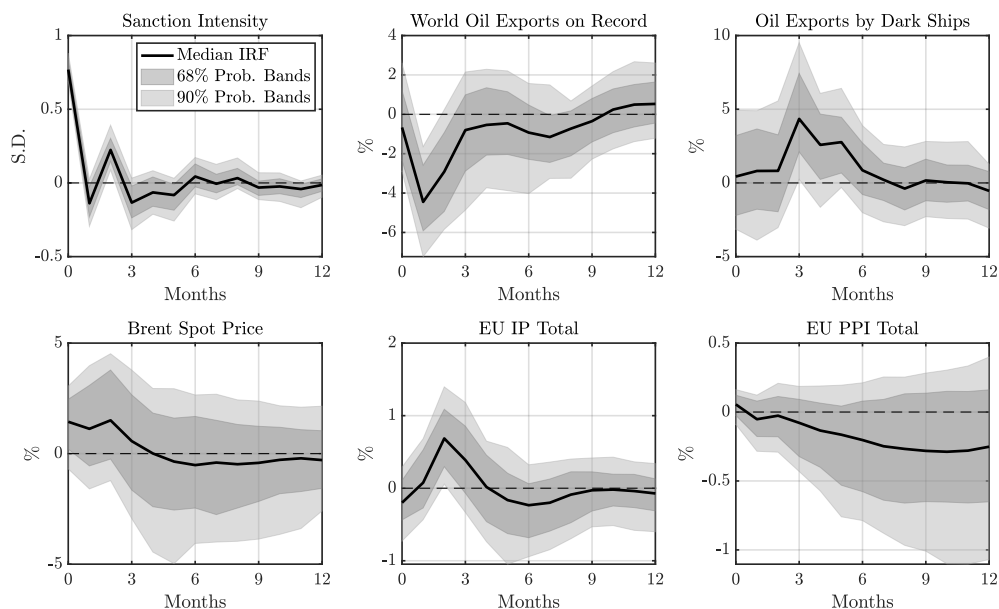


Figure E.8: IRFs to an Oil Sanction Intensity Shock for the EU Economy: Cholesky-Ordered SVAR

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the EU economy are computed using a Cholesky-ordered SVAR model with a lag length of three. The ordering of endogenous variables mirrors that in Figure 13, except that the order of total PPI and IP is switched. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

Figure E.9 presents the IRFs to an unexpected oil sanction intensity shock for OPEC countries. The IRFs are derived using the LP model (3), which includes the following monthly time series: (i) the sanction intensity index, (ii) crude oil exports from sanctioned countries transported by dark ships, (iii) OPEC crude oil exports, (iv) OPEC spot crude oil price, (v) OPEC net oil export revenue, and (vi) OPEC total spare crude oil production capacity. The sample spans January 2017 to December 2023.

Crude oil export data are sourced from the UN Comtrade database (HS code 2709), and spot crude oil prices are sourced from OPEC’s official website. The monthly OPEC net oil export revenue series is interpolated using the Chow-Lin method (Chow and Lin, 1971), based on the annual series from the U.S. Energy Information Administration (EIA) and the monthly OPEC crude oil export value. Spare production capacity data are also obtained from the EIA.

The oil sanction intensity shock is identified by controlling for two lags of the endogenous variables, as determined by the HQIC criterion, and estimating the LP model from horizon $k = 0$. Similar to the estimation for the U.S. and EU economies, the IRFs are approximated using the SLP method described in Barnichon and Brownlees (2019).

Tracing the dynamic responses of OPEC-related variables shows that an unexpected rise in oil sanction intensity reduces OPEC’s oil export volume, price, and net revenue, though to varying degrees. Meanwhile, OPEC’s spare production capacity increases, indicating a less tight market for its oil. As dark-shipped oil exports decline after the sixth month, while OPEC’s oil exports continue falling, OPEC’s spot crude price rises significantly, boosting net oil export revenue. Spare capacity remains elevated, reflecting OPEC’s strategy of using supply cuts to restore revenue margins.




Figure E.9: IRFs to an Oil Sanction Intensity Shock for OPEC Countries

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for OPEC countries are computed using the SLP method (Barnichon and Brownlees, 2019). Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

Lastly, Figure E.10 presents the IRFs to an unexpected oil sanction intensity shock for China. The impulse responses are derived using the LP model in Equation (3), which includes the following monthly time series: (i) the sanction intensity index, (ii) crude oil exports from sanctioned countries transported by dark ships, (iii) China’s PPI for mining and quarrying, (iv) China’s VAI for mining, (v) China’s PPI for manufactured, and (vi) China’s VAI for manufacturing, covering January 2017 to December 2023. All China-related data are sourced from the National Bureau of Statistics of China.

The oil sanction intensity shock is identified by controlling for three lags of endogenous variables, as determined by the HQIC criterion, and estimating the LP model from horizon $k = 0$.

As shown in Figure E.10, China benefits from reduced costs and increased production in both the energy (proxied by mining and quarrying) and manufacturing industries.




Figure E.10: IRFs to an Oil Sanction Intensity Shock for China

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for China are computed using the SLP method (Barnichon and Brownlees, 2019). Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

E.2. Propagation Channels

To complement Figures 14 and 15 in the main text, Figures E.11 and E.12 show the full IRFs to an unexpected oil sanction intensity shock for the U.S. and EU economies, respectively.

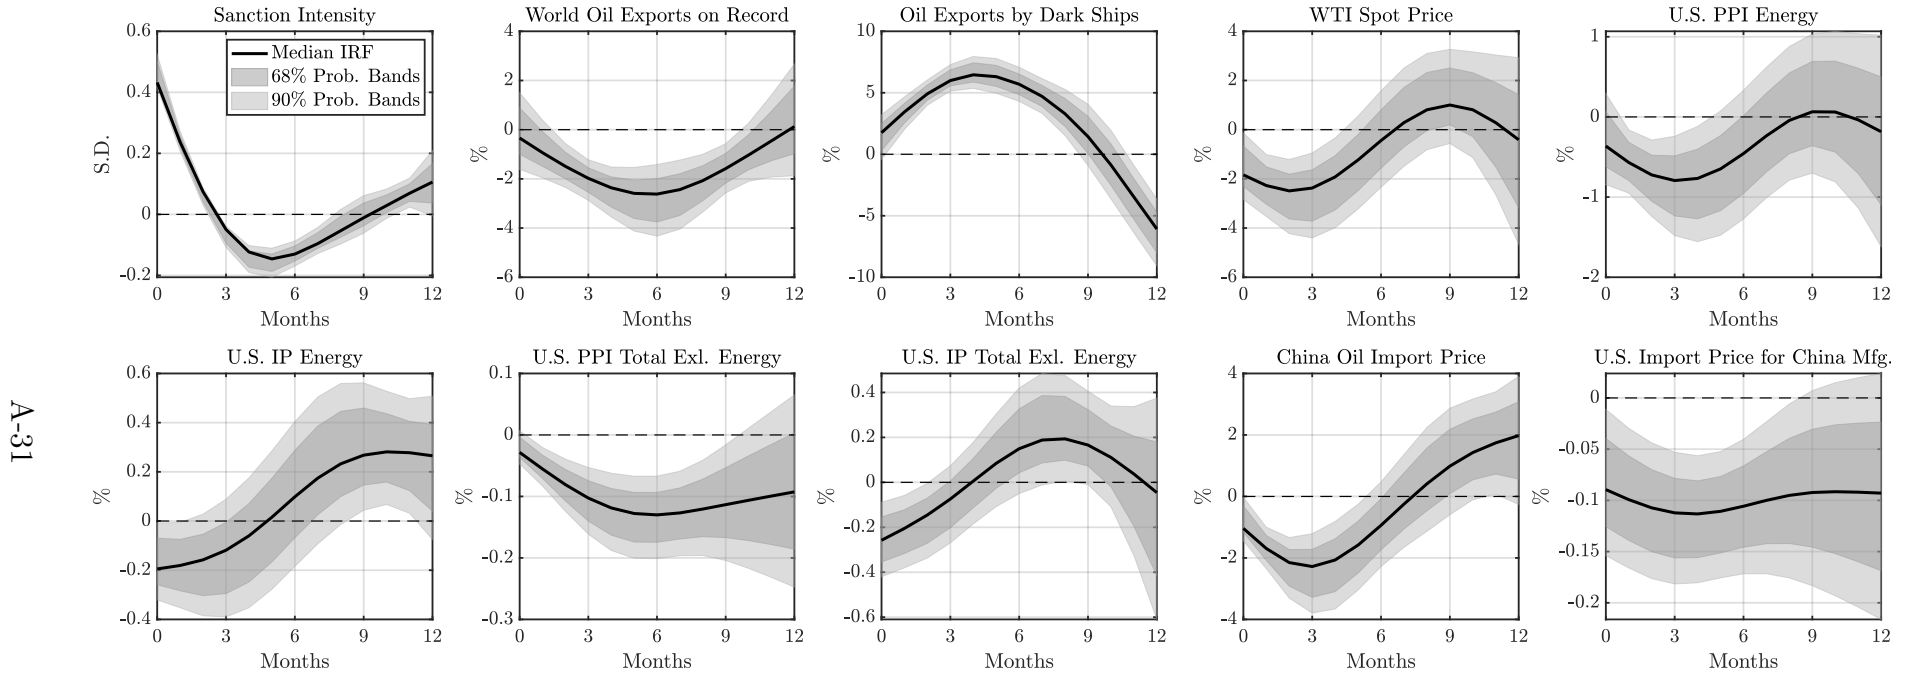


Figure E.11: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: Propagation Channels (Full IRFs)

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using the SLP method described in [Barnichon and Brownlees \(2019\)](#). The estimation specifications follow the baseline outlined in Section 5.1 of the main text, with two modifications: (i) two additional variables – China’s crude oil import price and the U.S. import price for manufacturing goods from China, averaged across NAICS sectors 31 to 33 – are included, and (ii) the total PPI and IP for the U.S. economy are replaced by the PPI and IP for the energy and non-energy sectors. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.




Figure E.12: IRFs to an Oil Sanction Intensity Shock for the EU Economy: Propagation Channels (Full IRFs)

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the EU economy are computed using the SLP method described in [Barnichon and Brownlees \(2019\)](#). The estimation specifications follow the baseline outlined in Section 5.1 of the main text, with two modifications: (i) two additional variables – China’s total IP and its total import value from the EU – are included, and (ii) the total PPI and IP for the EU economy are replaced by the PPI and IP for the energy and non-energy sectors. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

Figure E.13 presents the IRFs to an unexpected oil sanction intensity shock for the U.S. economy, using the same model specification as Figure 14, except that China's oil import price is replaced with its total PPI. The monthly series for China's total PPI is sourced from the National Bureau of Statistics of China and has been seasonally adjusted.

The results indicate that this substitution has little impact, as China's total PPI consistently declines across all forecast horizons.




Figure E.13: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: China's Total PPI

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using the same LP model specification as in Figure 14, with the exception of replacing China's oil import price with its total PPI. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

Figure E.14 presents the IRFs to an unexpected oil sanction intensity shock for the U.S. economy, using the same model specification as Figure 14, except that the U.S. import price for China’s manufactured goods is replaced with the import price for all Chinese products. The monthly series for the U.S. import price of all Chinese products is sourced directly from FRED (CHNTOT) and has been seasonally adjusted. Once again, the results appear largely unaffected by this variable substitution.




Figure E.14: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: U.S. Import Price for China

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using the same LP model specification as in Figure 14, except that the U.S. import price for China’s manufactured goods is replaced with the import price for all Chinese goods. Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

Figure E.15 presents the IRFs to an unexpected oil sanction intensity shock for the U.S. economy, based on an 11-variable LP model (3). This model extends the ten endogenous variables from Figure 14 by adding U.S. crude oil exports, using data from the U.S. EIA.

The ten original variables exhibit responses similar to those in Figure 14. U.S. crude oil exports closely track the WTI spot price, declining sharply in the first six months before stabilizing near the zero-response line.

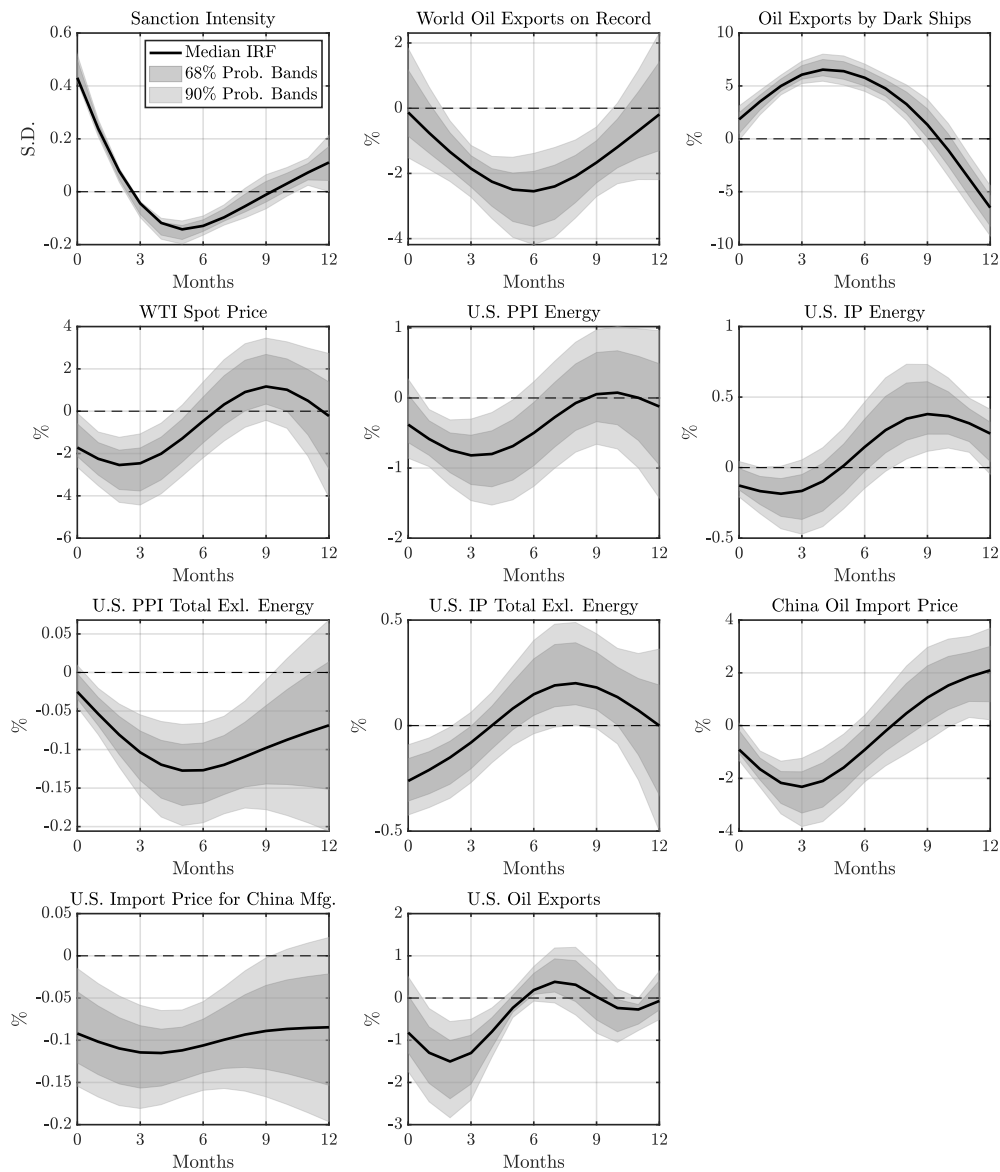


Figure E.15: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: U.S. Oil Exports

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using an 11-variable LP model (3), which includes the same ten endogenous variables as in Figure 14 of the main text, plus U.S. crude oil exports. The model is estimated using the SLP method outlined in Barnichon and Brownlees (2019). Black solid lines represent median responses, while gray-shaded areas indicate the 68% and 90% confidence bands.

Figure E.16 compares the IRFs for the U.S. economy with and without two China-related variables – China’s oil import price and the U.S. import price for Chinese manufactured goods. In the first three months, the median responses of U.S. energy and non-energy IP are closer to the zero-response lines than in Panels 14b and 14d of the main text. This suggests a positive omitted variable bias, where lower Chinese import prices support short-term U.S. production.




Figure E.16: IRFs to an Oil Sanction Intensity Shock for the U.S. Economy: With and Without China-Related Variables

Notes. The IRFs to a one-standard-deviation unexpected oil sanction intensity shock for the U.S. economy are computed using both a 10-variable LP model, as in Figure 14, and a counterfactual 8-variable LP model that excludes the two China-related variables: China’s oil import price and the U.S. import price for China’s manufactured products. The IRFs are approximated using the SLP method (Barnichon and Brownlees, 2019). The black solid lines represent the median responses estimated from the 10-variable LP model, with gray-shaded areas showing the corresponding 68% and 90% confidence bands. The red solid lines show the median responses from the counterfactual 8-variable LP model, while the red dash-dotted and dotted lines indicate the respective 68% and 90% confidence bands.

References for Appendices

- Adland, R., Cariou, P., and Wolff, F.-C. (2020). Optimal Ship Speed and the Cubic Law Revisited: Empirical Evidence From an Oil Tanker Fleet. *Transportation Research Part E: Logistics and Transportation Review*, 140:101972.
- Androjna, A., Pavić, I., Gucma, L., Vidmar, P., and Perkovič, M. (2024). AIS Data Manipulation in the Illicit Global Oil Trade. *Journal of Marine Science and Engineering*, 12(1).
- Ballinger, O. (2024). Automatic Detection of Dark Ship-To-Ship Transfers Using Deep Learning and Satellite Imagery. *arXiv preprint*.
- Barnichon, R. and Brownlees, C. (2019). Impulse Response Estimation by Smooth Local Projections. *Review of Economics and Statistics*, 101(3):522–530.
- Bernabé, P., Gotlieb, A., Legéard, B., Marijan, D., Sem-Jacobsen, F. O., and Spieker, H. (2024). Detecting Intentional AIS Shutdown in Open Sea Maritime Surveillance Using Self-Supervised Deep Learning. *IEEE Transactions on Intelligent Transportation Systems*, 25(2):1166–1177.
- Chow, G. C. and Lin, A. (1971). Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series. *Review of Economics and Statistics*, 53:372–375.
- Chupilkin, M., Javorcik, B., Peeva, A., and Plekhanov, A. (2024). Decision to Leave: Economic Sanctions and Intermediated Trade. Working paper, European Bank for Reconstruction and Development.
- Li, Y., Bai, X., Wang, Q., and Ma, Z. (2022). A Big Data Approach to Cargo Type Prediction and Its Implications for Oil Trade Estimation. *Transportation Research Part E: Logistics and Transportation Review*, 165:102831.
- Lloyd, S. (1982). Least Squares Quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137.
- Lloyd’s List Intelligence (2023). Shifty Shades of Grey: The Different Risk Profiles of the Dark Fleet Explained. Technical report, Lloyd’s List Intelligence.
- Lloyd’s List Intelligence (2024a). Understanding AIS Spoofing: The Case of the Blazers. Seasearcher Advanced Risk & Compliance.
- Lloyd’s List Intelligence (2024b). Understanding AIS Spoofing: The Case of the Shanaye Queen. Seasearcher Advanced Risk & Compliance.
- MacQueen, J. B. (1967). Some Methods for Classification and Analysis of Multivariate Observations. In Cam, L. M. L. and Neyman, J., editors, *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297, Berkeley, CA. University of California Press.
- Miller, G. (2023). Welcome to the Dark Side: The Rise of Tanker Shipping’s ‘Shadow Fleet’. FreightWaves.

- Paris MoU (2018). Explanatory Note – “White,” “Grey” and “Black List”. Technical report, Paris Memorandum of Understanding on Port State Control.
- Ristic, B., La Scala, B., Morelande, M. R., and Gordon, N. (2008). Statistical Analysis of Motion Patterns in AIS Data: Anomaly Detection and Motion Prediction. In *2008 11th International Conference on Information Fusion*, pages 1–7, Cologne, Germany.
- Rodger, M. and Guida, R. (2022). Mapping Dark Shipping Zones Using Multi-Temporal SAR and AIS Data for Maritime Domain Awareness. In *2022 IEEE International Geoscience and Remote Sensing Symposium*, pages 3672–3675, Kuala Lumpur, Malaysia.
- Rong, H., Teixeira, A., and Guedes Soares, C. (2024). A Framework for Ship Abnormal Behaviour Detection and Classification Using AIS Data. *Reliability Engineering & System Safety*, 247:110105.
- Triebert, C., Migliozi, B., Cardia, A., Xiao, M., and Botti, D. (2023). Fake Signals and American Insurance: How a Dark Fleet Moves Russian Oil. *The New York Times*.
- U.S. Department of State (2020). Sanctions advisory for the maritime industry, energy and metals sectors, and related communities. Technical report, U.S. Department of State, U.S. Department of the Treasury’s Office of Foreign Assets Control, and U.S. Coast Guard.
- Yan, Z., Xiao, Y., Cheng, L., Chen, S., Zhou, X., Ruan, X., Li, M., He, R., and Ran, B. (2020). Analysis of Global Marine Oil Trade Based on Automatic Identification System (AIS) Data. *Journal of Transport Geography*, 83:102637.