

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.

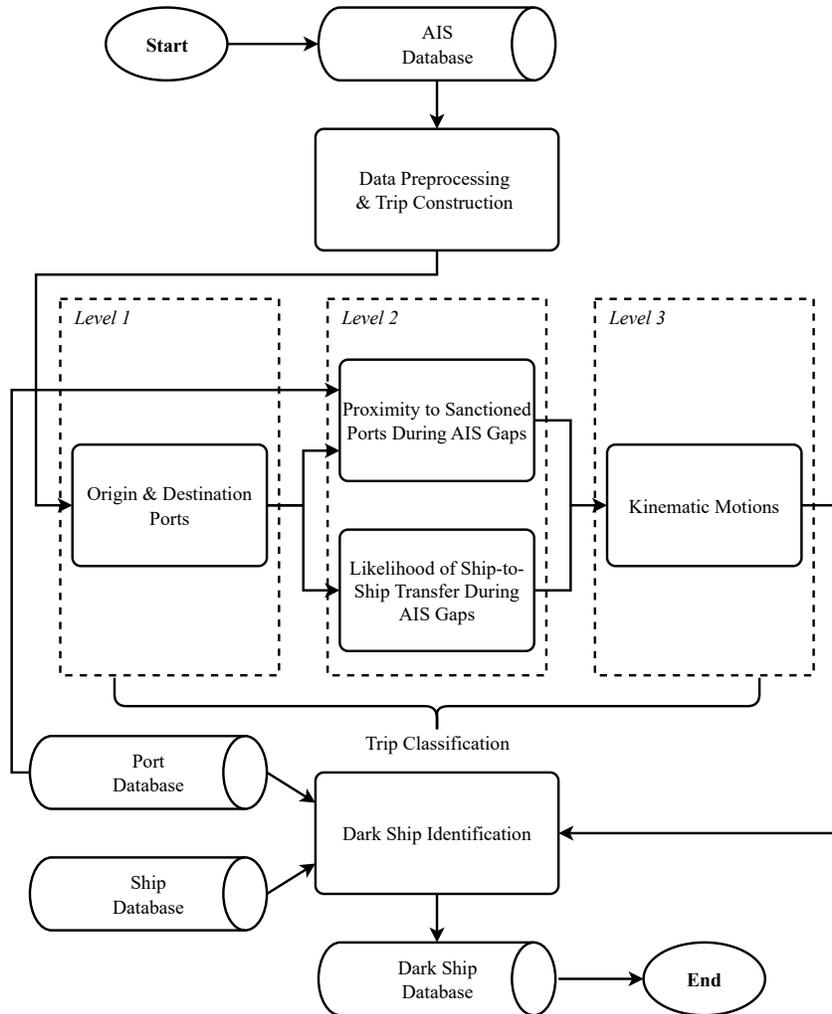


Figure A.1: Methodology Framework of the Multi-Attribute Ship Clustering Model

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 $j \in \mathcal{J}$
 $l_j^{Trip} = \left\{ l_{j,t}^{Trip} \right\}_{t \in t_j}$: trip status labels for ship $j \in \mathcal{J}$
 $l_j^{Gap} = \left\{ l_{j,t}^{Gap} \right\}_{t \in t_j}$: 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:    $(dx_j, dy_j) \leftarrow$  HEADING_TO_VECTOR( $h_{j,t}$ )
9:    $(dx_{j'}, dy_{j'}) \leftarrow$  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.

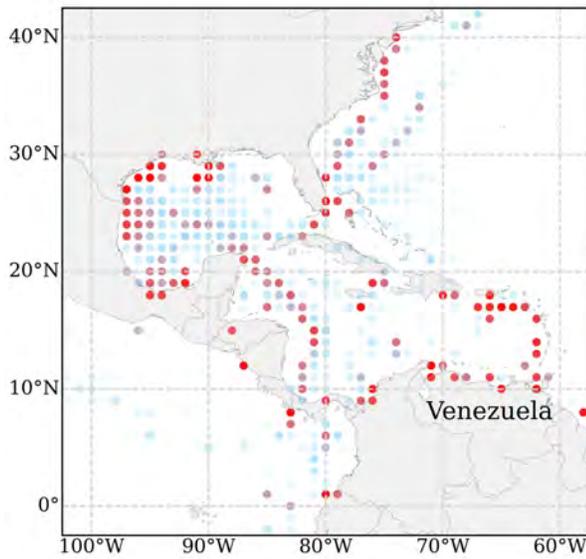
A vessel’s flag state ranking also provides insight into its regulatory environment, as ships registered under higher-risk flags are more likely to engage in non-compliant activities, as discussed in [Section 2.4](#) of the main text.

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.

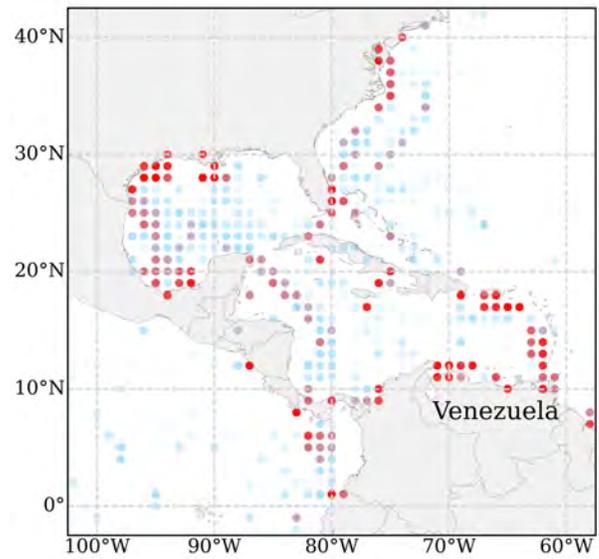
¹¹The idle trip ratio functions similarly to average vessel speed in detecting anomalies in kinematic movements. Dark ships often remain idle for extended periods to facilitate unauthorized ship-to-ship transfers, evade detection, or operate beyond regulatory oversight.

B. Additional Dark Ship Identification Results

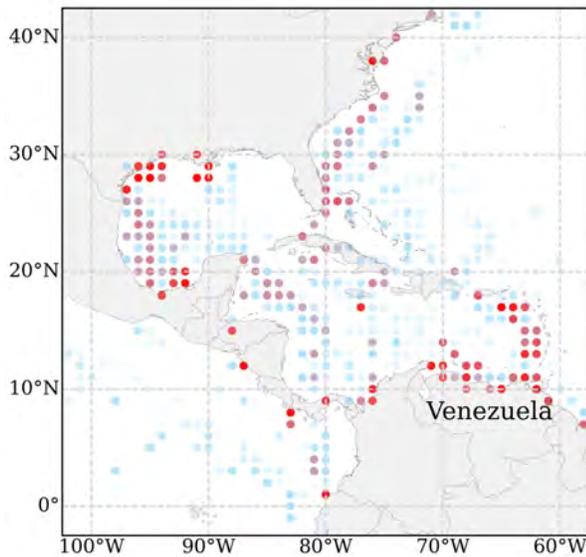
In this appendix, we present Figures B.1 to B.4, which illustrate the concentration of direct vessel visits to suspicious ports in Venezuela, Iran, and Russia near the Black Sea and Baltic Sea during periods not covered in the main text.



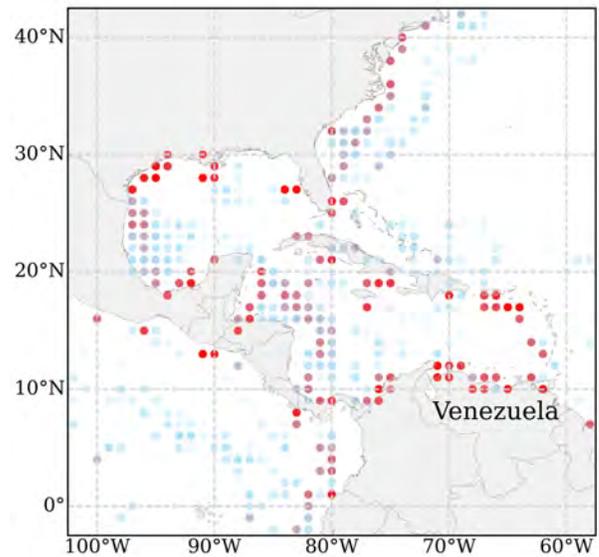
(a) 2020



(b) 2021



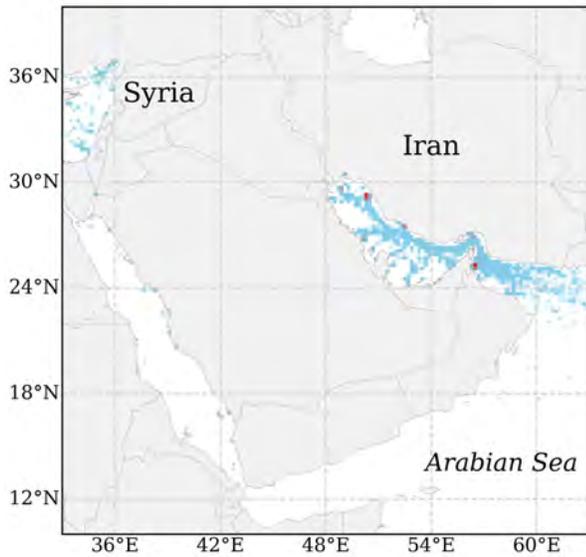
(c) 2022



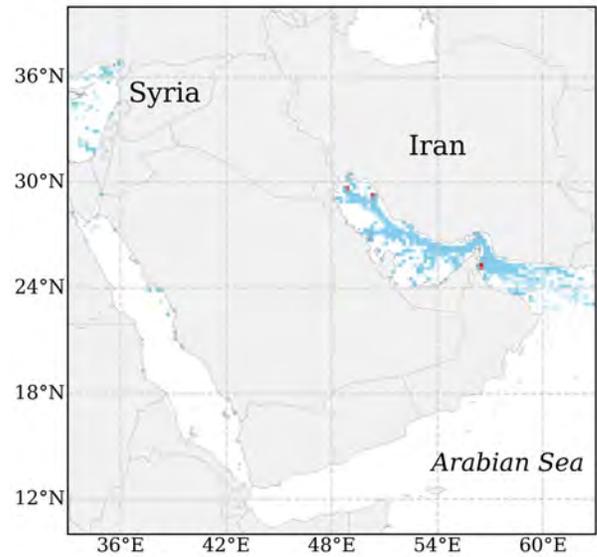
(d) 2023

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

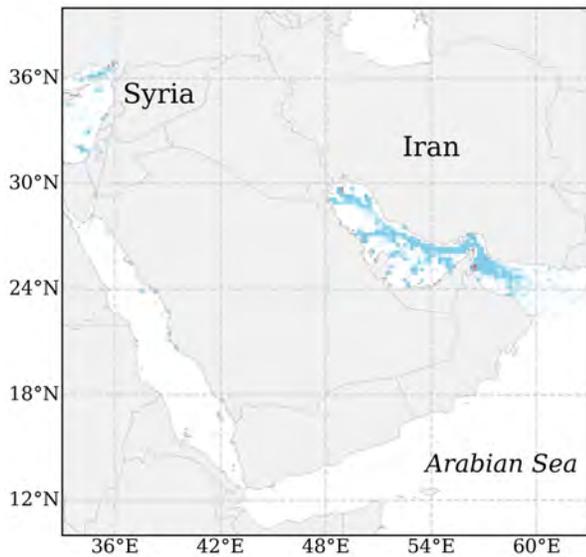
Notes. The figures illustrate the geographical concentration of dark shipping activities near Venezuela from 2020 to 2023. Colored points and chromatic intensity follow the same format as in Figure 4 of the main text.



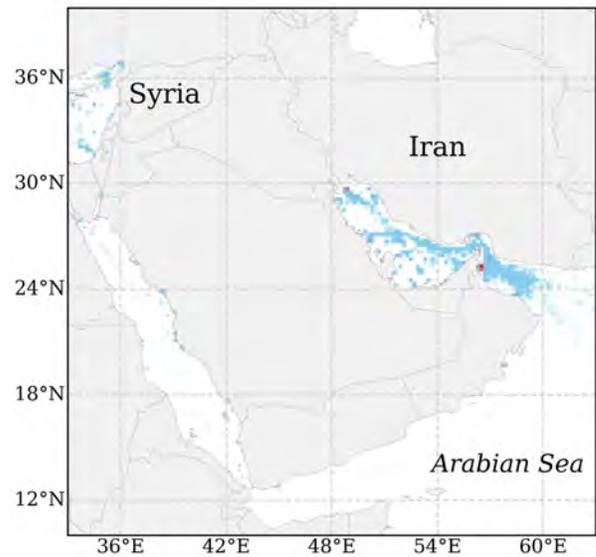
(a) 2017



(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.

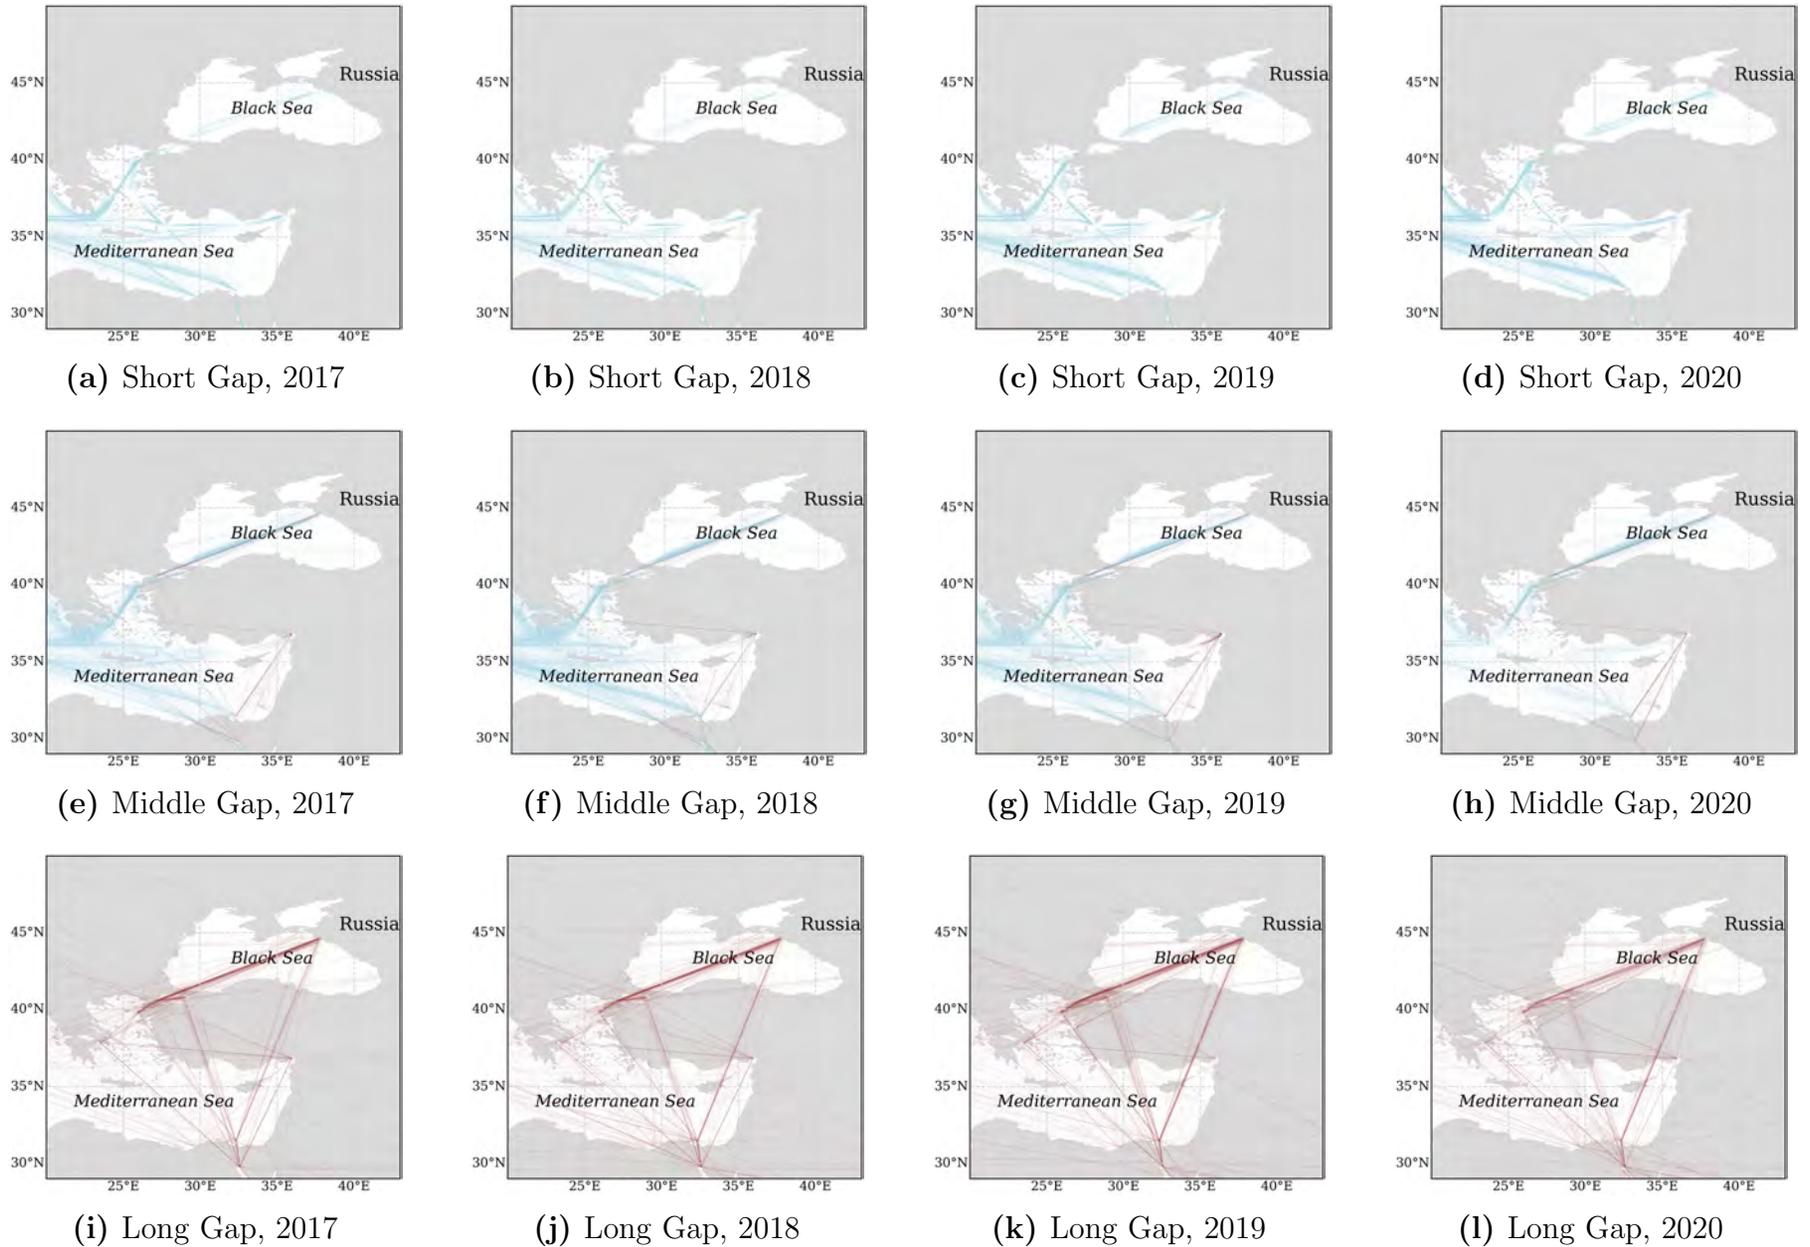


Figure B.3: Geographical Concentration of Dark Shipping Activities in the Black Sea: 2017–2020

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.

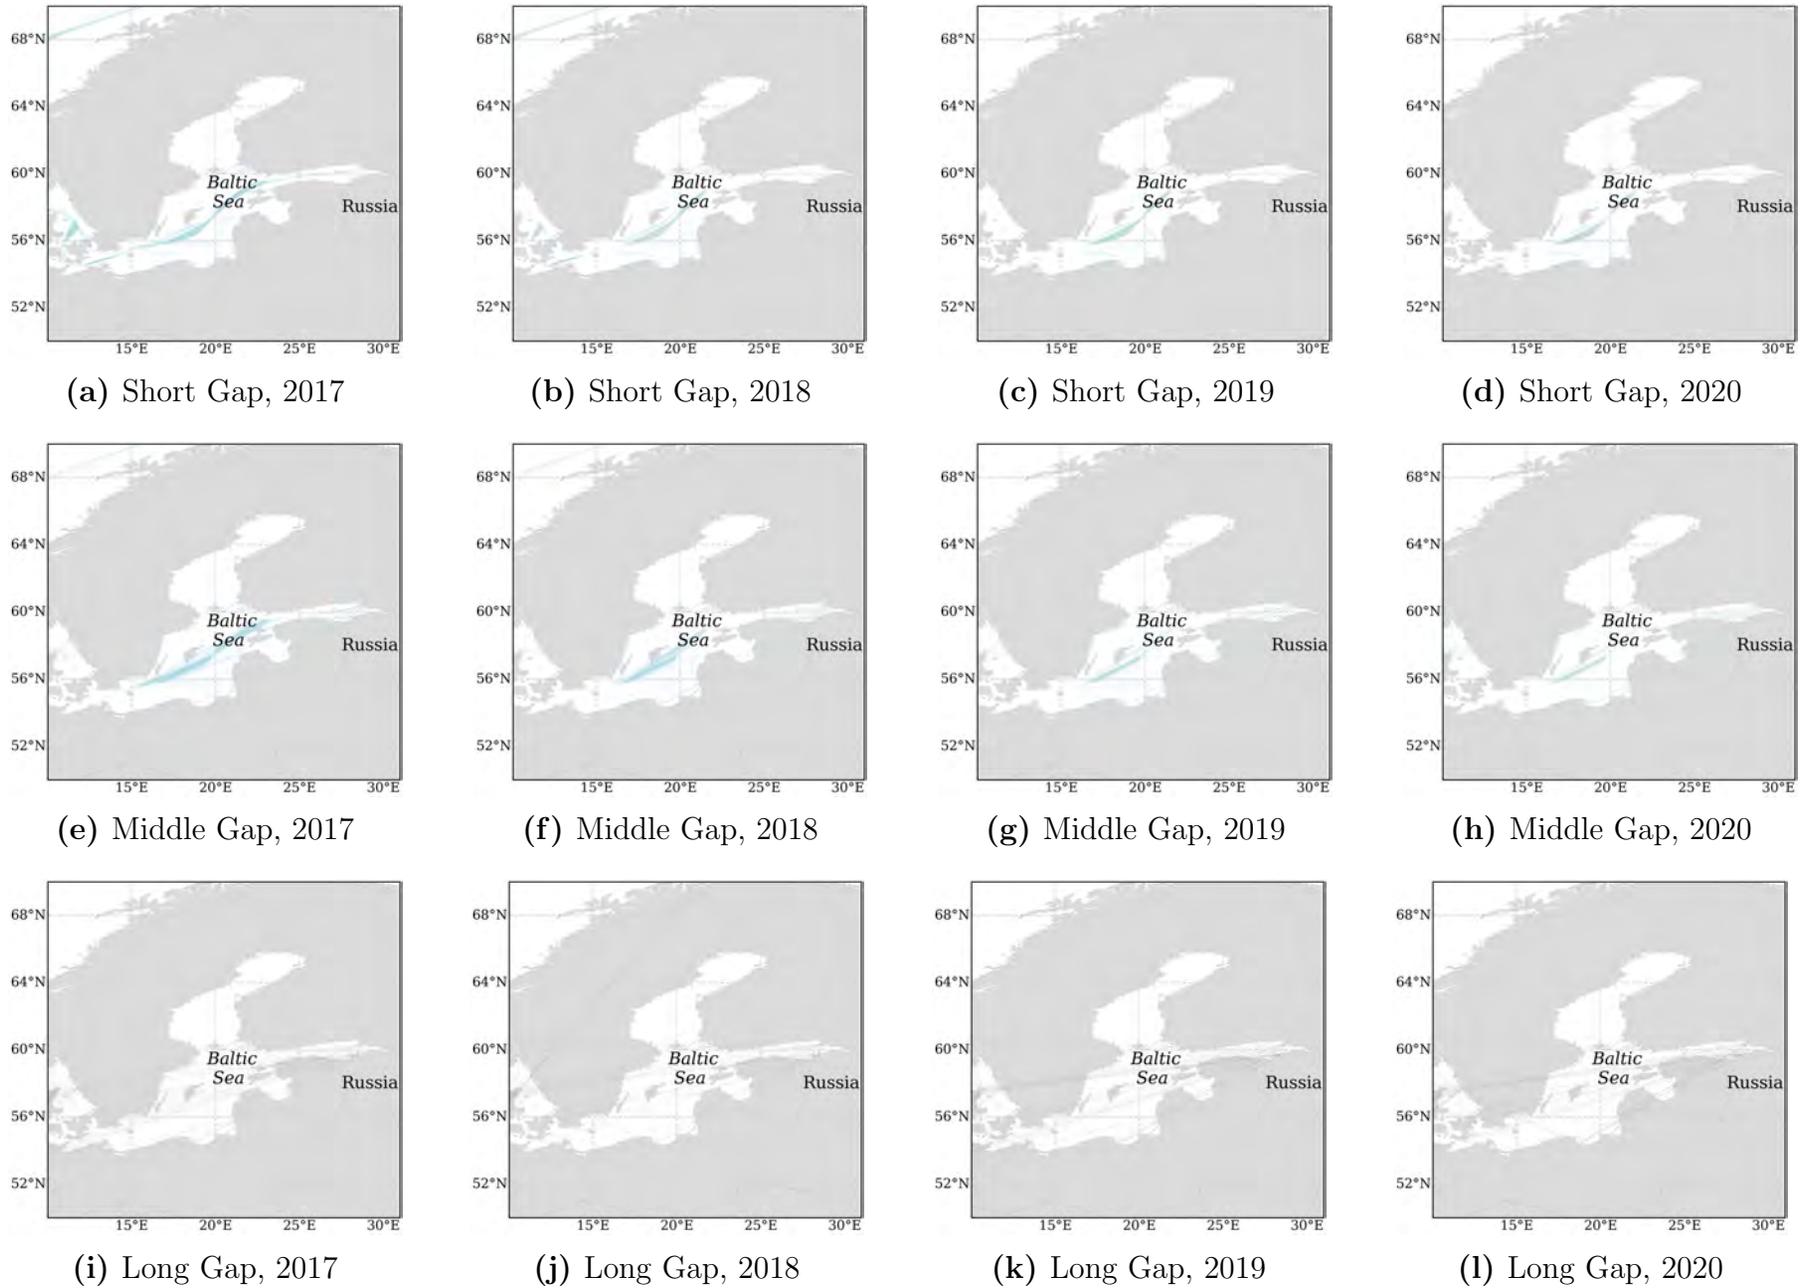


Figure B.4: Geographical Concentration of Dark Shipping Activities in the Baltic Sea: 2017–2020

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.

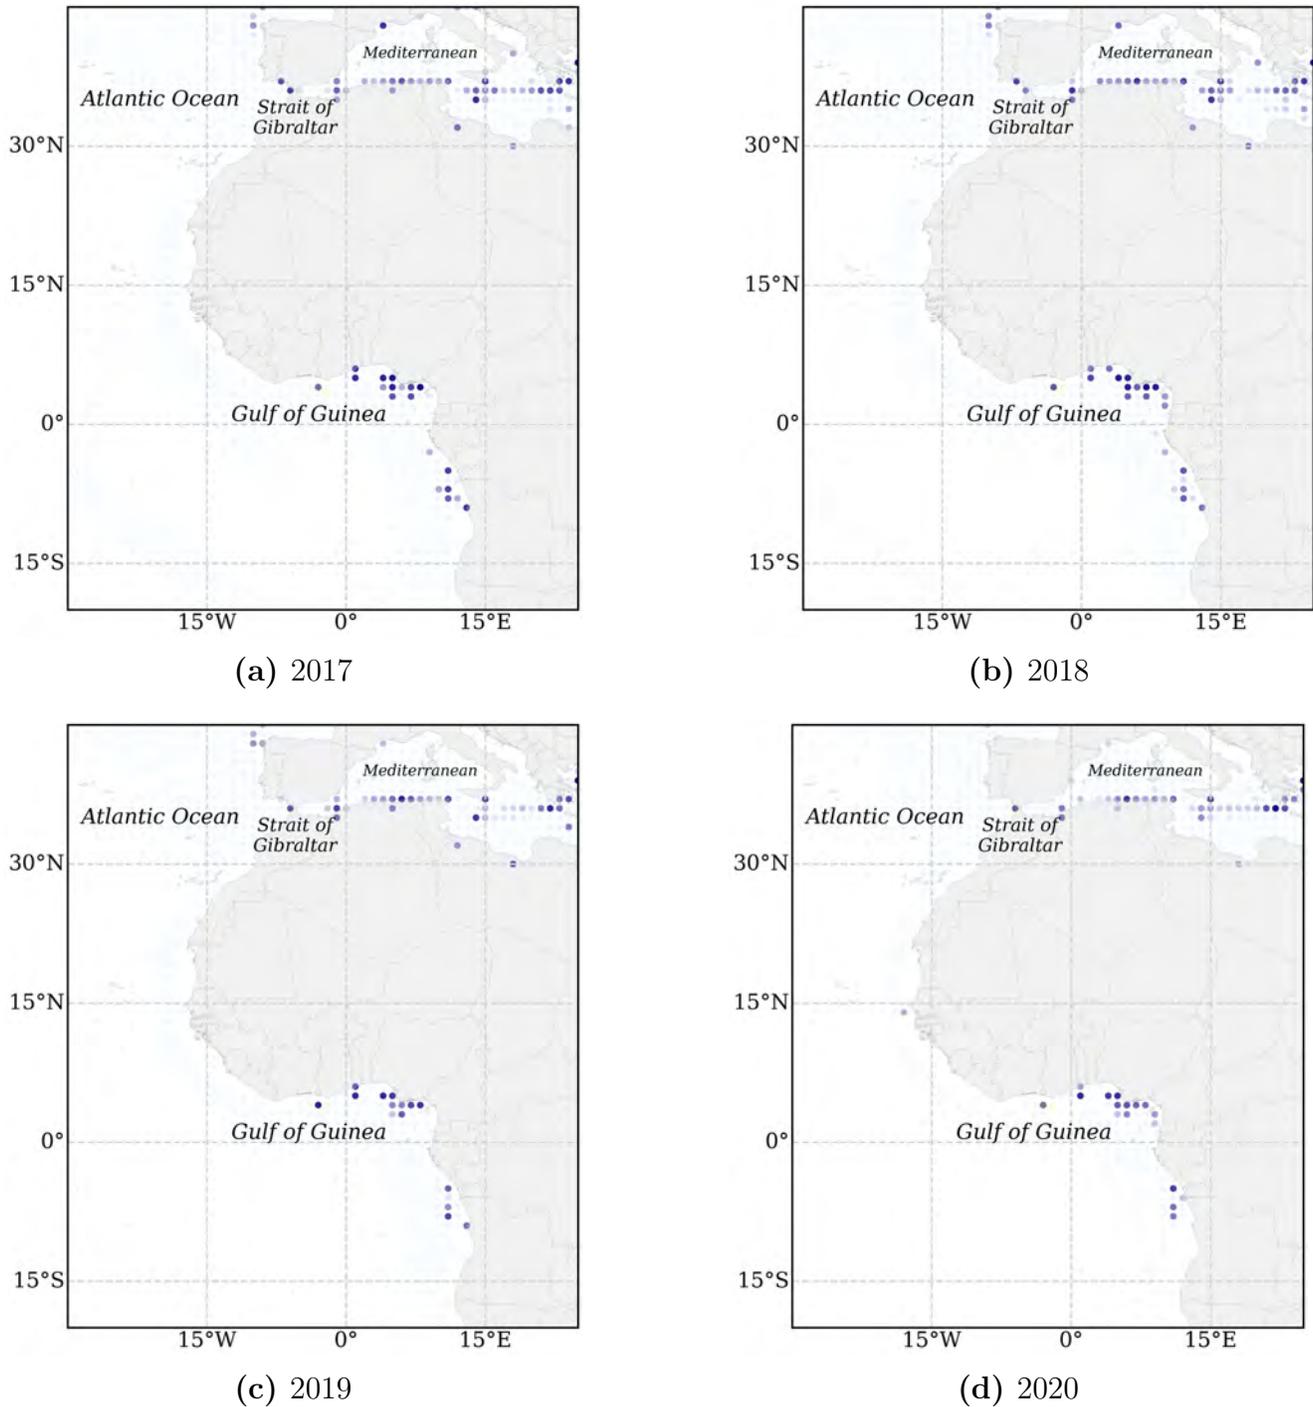


Figure B.5: Geographical Concentration of Dark Shipping Activities Near the Strait of Gibraltar and the Gulf of Guinea: 2017–2020

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.