Is a Lack of Information Limiting Sanctions Enforcement? *

Avi Dutt, Abhiroop Mukherjee, George Panayotov, Debjit Roy, Xudong Wen

February 26, 2025

Abstract

The enforcement of economic sanctions relies on third-party businesses to identify and avoid violators. Using the case of sanctions-violating oil tankers, we show that (i) such identification is far more challenging than commonly assumed, and (ii) this difficulty weakens sanctions enforcement. To demonstrate the latter, we examine the sudden public disclosure of a list of suspect tankers by a leading maritime AI firm. This revelation triggered sharp market penalties, with affected tankers experiencing a 13% earnings drop and a 37% decline in secondary market liquidity. A dynamic oil shipping model, disciplined by market data, shows that this disclosure alone redirected \$2.3 billion per year away from violators. Surprisingly, our model predicts – and we empirically confirm – that this disclosure also reduced earnings for compliant tankers. Therefore, the benefits of disclosing violators can extend beyond just improving sanctions enforcement – here, it reduced overall shipping costs for legally traded oil.

Keywords: Sanctions enforcement, Information costs, Shipping, Machine learning

*Dutt: Trademo; Mukherjee, Panayotov, and Wen: Hong Kong University of Science and Technology; Roy: Indian Institute of Management Ahmedabad (IIMA). We thank Vishesh Vatsal and Dfy-graviti for their help in analyzing satellite imagery, and Niranjana Unnithan and Benas Zurauskas for excellent research assistance. We also thank Sumit Agarwal, Tania Babina, Utpal Bhattacharya, Briana Chang, Darwin Choi, Shan Ge, John Griffin, Minkang Gu, Harrison Hong, Yan Ji, Navin Kumar, Erica Li, Vikram Mathur, Jagmeet Makkar, S. Lakshmi Naaraayanan, John Nash, Jun Pan, Elias Papaioannou, Krishna Prasad, Robert Ready, Vivek Reshamwala, Arkodipta Sarkar, Rik Sen, Janghoon Shon, Ashok Srinivasan, Yongxiang Wang, Constantine Yannelis, Jianfeng Yu, Xiaoyun Yu, Alminas Zaldokas, Jian Zhang and seminar participants at CKGSB, FMA Asia, Hong Kong University Summer Conference, HKUST, SAIF, The Education University of Hong Kong, and Tsinghua University for various helpful comments and suggestions. We thank the UN Global Platform for access to its AIS database (UN Statistics Division, "Global AIS Data Feed.", 2020).

Sanctioning regimes seek to deter violations not only by threatening legal penalties on violators, but also by requiring third-party businesses and individuals to avoid engaging with them. For example, the U.S. mandates that any entity providing loans, insurance, or transportation to a violator faces potential punishment (Early and Peterson (2022); Eaton and Engers (1992); Pape (1998)). Foreign firms are subject to these restrictions as well, facing the risk of exclusion from the U.S. economy and financial system.[1]

In this paper, we argue that the conventional focus on penalizing third parties for engaging with violators overlooks a key challenge: the difficulty of identifying sanctioned entities and individuals. Even firms committed to compliance struggle to determine whom to avoid, due to limited information on sanction evasion, which weakens the overall effectiveness of sanctions.

While authorities like the U.S. Office of Foreign Assets Control (OFAC) publish lists of sanctioned entities, many violators are still not sanctioned (and so not identified on any government's list). For instance, the OFAC has added fewer than 100 oil tankers to its list from January 2021 until May 2024 (the end of our sample), although the Congressional Research Service estimates that 1,600 tankers have carried sanctioned oil in that period. This leaves businesses to identify around 1,500 violators on their own – a daunting task for many. As Daniel Tadros, Chief Operating Officer of the American Club, a major U.S. shipping insurer, told The New York Times: "It's impossible for us to know on a daily basis exactly what every ship is doing, where it's going, what it's carrying, who its owners are. I would like to think that governments have a lot more capability, manpower, resources to follow that."

This paper examines the role of identification in sanctions enforcement by addressing two key questions: (1) How accurately can businesses identify sanctions violators using publicly available data? (2) If identification were easier, would market forces penalize violators, or would business continue as usual? The latter would suggest that firms engage with violators not due to identification challenges, but because the benefits outweigh the risks (Early (2021)). In short, we ask: Is identifying violators difficult – and, if so, is it consequential for sanctions enforcement?

We answer these questions by focusing on oil tankers, which play a crucial role in the global economy (Brancaccio, Kalouptsidi and Papageorgiou (2020, 2023); Hamilton (1983); Ready (2018)). Oil exports also contribute significantly to the revenues for sanctioned countries, like Russia, Iran, and Venezuela, making the success of sanctions closely tied to re-

[1] A notable example is the secondary sanctions imposed on the Bank of Dandong, Dalian Global Unity Shipping Co., and two individuals (Sun Wei and Li Hong Ri) as a result of their dealings with North Korea-linked entities. The bank was barred from U.S. financial transactions, the company from U.S. trade, and the individuals had their assets frozen and were prohibited from dealing with U.S. entities.

stricting these countries' capabilities to move their oil (Babina et al. (2023); Brown (2020); Farzanegan (2011)).

To address the first question, we analyze how a third party can use publicly available data to identify sanctions violators and assess the accuracy of these predictions. While one might assume that tracking large oil tankers is straightforward with modern satellite and ship-tracking technology, we demonstrate that this assumption is incorrect.

Our analysis here is made possible by a unique "ground truth" dataset provided by an anonymous port agent from a Middle Eastern country. This dataset identifies all oil tankers that violated Iranian sanctions in January 2021, including 33 foreign-flagged (non-Iranian) tankers operating in the Persian Gulf. While limited in size, this dataset is, to our knowledge, the only one of its kind. It provides crucial ground truth for evaluating models that predict sanctions evasion in shipping – a key missing component that has thus far constrained research (e.g., Wolsing et al. (2022)).

Using this dataset, we assess the predictive accuracy of various machine learning (ML) models designed to identify sanctions violators. These models incorporate factors highlighted by industry experts, such as tankers temporarily disappearing from tracking systems or exhibiting unusual travel patterns indicative of route falsification (see Figure 1). Our analysis reveals that even a sophisticated third party using these methods would correctly identify only 39% violators at best (with 95% confidence). This detection rate is low and likely overestimates the overall ease of detection, as our ground truth data is limited to the Persian Gulf – a geographically smaller region compared to the vast and complex waterways around, e.g., Russia.

Given the difficulty of identifying suspect tankers, we next address our second question – whether revealing the identities of suspected sanctions violators harms these suspects, and more importantly, whether and to what extent it hurts true sanctions violators and sanctioned oil exporters. An ideal experiment to answer this question would involve abruptly releasing a list of suspect tankers to the public and then studying its effect on these tankers compared to "counterfactual" tankers – those not on the suspect list, but ex-ante similar to the listed vessels. We get close to this ideal by leveraging Windward.AI's (LON:WNWD) disclosure of a global suspect tanker list on August 16, 2023, via the London Stock Exchange Group's Refinitiv Eikon platform. While Windward and similar firms had previously sold proprietary lists to select clients, these datasets were prohibitively expensive for all but a few major players. For example, in their analysis of maritime sanctions advisories, IHS Markit and the Association of Certified Sanctions Specialists (IHS-ACSS, 2022) reported that "For smaller banks, the platforms and specific shipping tracking software and datasets, via third-party vendors, are considered to be very expensive in terms of their cost outlay.

Every transaction per dollar for managing ship screening checks would probably not make sense in terms of cost from the bank's side." Following the disclosure, Windward's suspect tanker module became available for just $356 per month, dramatically lowering the cost of accessing sanctions-risk information.

We first evaluate the accuracy of Windward's sanctions-risk predictions using our ground truth list of actual violators. We find that 27 out of 33 tankers on our Iranian sanctions-violator list were classified as high or moderate risk by Windward. In contrast, our machine learning models could correctly identify at most 13 of these tankers at a 95% confidence level. Another publicly available list – from the advocacy group United Against Nuclear Iran (UANI) – included only 10 of these tankers as of August 2023. This comparison suggests that Windward's list is significantly more accurate than prior alternatives.[2] Further, this overlap with Windward also reassures us about the authenticity of the Iranian list we received.

Next, we examine the extent to which Windward's disclosure improves third parties' ability to detect soon-to-be sanctioned tankers. Specifically, we measure the incremental accuracy gained by incorporating Windward's list, focusing on reductions in the false positive rate while maintaining a 5% risk of unknowingly engaging with such tankers. We show that augmenting a shipping signals-based machine learning model with Windward's list reduces the false positive rate from 51% to 20%. This substantial improvement indicates that Windward's disclosure likely constituted a major information shock for identifying sanctions-violating tankers.

We then analyze the real effects of Windward's disclosure on suspect tankers, using a Propensity Score Matched Difference-in-differences setting (PSM-DiD), as well as Abadie (2005)'s semiparametric DiD approach. We compare high/moderate risk tankers disclosed by Windward with counterfactual tankers, before and after the disclosure. In particular, we focus on three key outcome variables. First, we find that earnings for high-risk tankers dropped by 13% after the disclosure, relative to counterfactuals. This could reflect the market's hesitation in hiring these tankers due to increased scrutiny – and possible impounding if found guilty of earlier violations – which can significantly disrupt shipment plans. Second, high-risk tankers are 17% less likely, compared to other similar tankers, to navigate within 12 nautical miles of sanctioned oil-producing countries like Russia, Iran, and Venezuela. Such behavior suggests that these vessels' fear of heightened scrutiny likely outweighs profit opportunities from continuing to serve these routes. Third, proprietary data from Drewry indicate that the chance of selling a tanker drops by 37% within a year once it is flagged as risky by Windward, likely due to sanctions regimes prohibiting ownership changes for violating tankers.

[2]We discuss potential reasons for Windward's accuracy in Section 4.1.

This analysis, however, is limited in its focus on *suspected* – rather than *actual* – sanctions violators, due to the lack of data on the latter. We address this limitation by developing a dynamic structural model of the oil shipping market. The model explicitly specifies variables on which we lack data, and estimates them using market data and our findings from the Windward event study.

Our results suggest that Windward's disclosure led to sanctioned oil exporters having to pay 2% higher freight rates, and to reduce their overall hiring of tankers by 11.7%. Further, the value of tankers that actually violated sanctions declined by 10%, translating to a loss of approximately $3 million per large 20-year-old tanker – a vessel type commonly used for transporting sanctioned oil. Given that an estimated 1,500 tankers that have not been formally sanctioned yet may be engaged in such violations, this implies a substantial aggregate financial loss for violators.

Finally, the model also produces a counter-intuitive implication: information shocks like Windward's can also reduce the value of "Good" tankers – those with no history of carrying sanctioned oil. This occurs because rogue exporters, fearing detection, avoid using previously implicated tankers, pushing these vessels to seek employment with clean (i.e, not-sanctioned) exporters. This increased tanker supply for clean exporters lowers earnings for Good tankers in the legitimate market. Although rogue exporters are now willing to pay more to attract these Good tankers, this increase does not offset the decline in earnings such tankers face in the much larger legitimate market.

In Section 5.4, we empirically test this prediction using a shift-share approach, leveraging the stickiness of tanker-charterer relationships and differential exposure of tanker types to the Windward shock. Our findings align with model predictions. This suggests that besides redirecting profits away from violators, disclosing violator identities can yield unintended benefits – such as lowering shipping costs for legally traded oil.

Our findings highlight the need to rethink how sanctions violators should be identified. While currently businesses bear most of the burden, a more transparent system – where authoritative sources disclose suspected violators – could ease this burden while strengthening sanctions enforcement by raising the costs of violations.

The rest of the paper is organized as follows: Section 1 discusses related literature, Section 2 provides a brief contextual background, Section 3 presents data, methods, and results from models detecting sanctions violators, Section 4 examines the effect of Windward's sanctions-risk disclosure, Section 5 presents a model and estimates key quantities of interest within its scope, and Section 6 concludes.

1 Related Literature

Our paper is related to the literature studying the difficulty in detecting evasion of economic embargoes and sanctions. Earlier studies on the topic include Hsieh and Moretti (2006), who use the price gap between Iraqi oil and its close substitutes to show that paper trails, even including the CIA's documentation, can underestimate the extent to which Iraq used bribes to circumvent the oil embargo. Dellavigna and La Ferrara (2010) examine how the illegal arms trade undercuts the effectiveness of the embargoes to a far greater degree than would be apparent from case-by-case evidence in investigative reports by groups such as Amnesty International. They propose a method to detect potential arms embargo violations based on reactions by individual stocks. More recently, Fisman et al. (2024) document a different type of sanctions circumvention – that of entrepot trading – using the context of Ukraine's economic blockade of the anthracite-rich breakaway Donbas region. Beyond the detection of sanctioned activity, our study relates to the literature on forensic economics, aimed at detecting illegal activities of various forms and fashion (e.g., Fisman and Wang (2017); Fisman and Wei (2004, 2009), Dimmock and Gerken (2012); Duggan and Levitt (2002); Griffin and Kruger (2023); Griffin and Maturana (2016); Griffin and Shams (2018); Jacob and Levitt (2003); Marion and Muehlegger (2008)).

We relate more generally to the literature on the economics of sanctions (e.g., Early and Preble (2019); Eaton and Engers (1992, 1999); Kaplow (1990); Lacy and Niou (2004); Mulder (2022); Siddiquee and van Bergeijk (2012); Clayton et al. (2023)). We also relate specifically to the macroeconomic, political, and firm-level effects of trade restrictions with Iran (Dizaji and van Bergeijk (2013); Haidar (2017)). More recent research discusses sanctions on Russia (e.g., Itskhoki and Mukhin (2022); Mamonov and Pestova (2023)), including a few studies on how firms, banks, and other entities affected by embargoes or sanctions respond (e.g., Crozet et al. (2021); Efing et al. (2023); Huynh et al. (2023); Lastauskas et al. (2023)). Different from this literature, our paper examines the impact of providing identifying information on suspected sanctions violators, rather than the impact of the imposition of sanctions themselves.

Finally, the disclosure of risk scores aspect of our paper is related to work on sudden disclosure of information through credit registries, as studied by Liberti et al. (2016). More generally, we relate to a large literature on the diffusion of information in financial markets (e.g., Hong and Stein (1999); Hong et al. (2007); Hong and Yogo (2012); Peress (2014)).

2 Identifying Suspected Sanctions Violators: Do Businesses Care?

Before we delve into the difficulty of identifying sanctions violators, we examine the extent to which third party businesses are investing in sanctions compliance, if at all. This is an important question for us because identification difficulties can be a binding constraint for sanctions enforcement only if third party businesses are actually making efforts to identify violators.

On one hand, regulators motivate them to do so by making sanctions-related penalties severe. For example, global financial institutions have paid billions in settlements for violations (e.g., $8.97 billion by BNP Paribas in 2014). Even transactions with entities that violated sanctions years prior are at risk of costly disruption, as extended statutes of limitations imply that such entities are still at risk of sanctions. Moreover, both primary sanctions (that target U.S. entities' activities directly) and secondary sanctions (that target non-U.S. entities by threatening their access to the U.S. financial system) operate on a strict liability basis, not requiring proof of intent to enforce penalties (FCA, 2024, OFAC, 2024). For instance, Eagle Shipping's new owner had to settle sanctions violations from 2011-2012 despite no personal involvement (OFAC, 2020).

On the other hand, compliance can be difficult and costly for those who are serious about it. For instance, as highlighted by the U.S. Treasury itself in its 2021 Sanctions Review, small businesses "may lack the resources to bear the costs of sanctions compliance". Even for larger entities, sanctions compliance costs can be substantial – e.g., they can constitute up to 15% of a bank's annual expenditure (Lloyd's List, 3 May 2022).

To shed light on whether third party businesses are indeed investing in sanctions compliance, we analyze annual (10-K) reports of major US financial institutions. These businesses are often particularly exposed to the risk of dealing with counterparties (e.g., borrowers) who might be violating sanctions without their knowledge. Figure 2, based on machine-reading these reports using an LLM, shows that over 75% of the 50 largest US financial firms expressed concerns about sanctions in 2023. In addition, ultimately the costs of avoiding sanctions risks and resultant disruptions hinge on the ability of such businesses to discern violators from legitimate entities. Accordingly, nearly all of them (98%) pointed out that identifying sanctions violators is a major challenge for their business. This suggests that the issue of identifying sanctions violators is indeed an important one, at least for the financial sector.

3 How difficult is it to detect sanctions violators?

As mentioned earlier, in this paper we focus on oil shipping and examine the burden faced by third parties in identifying potential violators in this sector. Sanctions compliance in this sector is far from trivial, since, per industry estimates, about 10% of oil tankers may have violated sanctions.[3]

To quantify the challenge of detecting sanctions evaders, we take the view of an entity that uses advanced analytical tools, such as machine learning, combined with radio-signal-based ship-tracking data from the Automatic Identification System (AIS) and satellite imagery. First we assess detection accuracy using only AIS data, and subsequently incorporate satellite data to determine how much it enhances accuracy.

3.1 Data on Oil Tankers

We utilize three datasets: AIS data on high-frequency shipping signals, Sentinel-1 satellite data, and a proprietary Iran dataset, to evaluate detection performance.

Testing detection efficacy requires a known, ”ground truth” subset of the data. We use such a subset from Iran, provided by an anonymous source with extensive Middle-Eastern shipping experience, and use it to evaluate AIS-based detection of sanctions-violating oil tankers. This proprietary set lists all non-Iranian-flagged tankers seen in Iran's Persian Gulf waters in January 2021 (total of 33 tankers). We cross-validated these data, which have been collected via aggregated text messages from ports, brokers, and charterers, against Windward's list of high-risk tankers.

Second, the AIS data we use was originally designed for collision avoidance. Vessels over 300 gross tonnage must carry AIS equipment (IMO, 2000). This data has been used in academic research before (e.g., Brancaccio et al. (2020)), and comprises of fixed information (IMO number, MMSI number, vessel type and size), dynamic information (position, speed), and voyage details (draught, destination, ETA), all transmitted at short intervals. Dynamic data is auto-updated, while voyage data is manually entered. AIS is a widely recognized tool for sanctions compliance (U.N. Doc. S/2019/171, Kilpatrick (2022)), and even flag state registries have used it to de-flag ships (Lloyd's List, Feb. 2020, Oct. 2020).

Finally, we employ data from the Sentinel-1 mission, which has two satellites, providing near-real-time, high-resolution images at six-day intervals of most locations on Earth via synthetic aperture radar (`https://skywatch.com/free-sources-of-satellite-data/`). The utility of commercial satellite imagery for identification is acknowledged by regulators

[3] "A mysterious fleet is helping Russia ship oil around the world. And it's growing", CNN, Mar. 2023; "Dark fleet of tankers now comprises 10% of seaborne oil transport", Lloyd's List, Mar 2023.

(OFAC, 2023, Economist, 16 Apr. 2022). We analyze images of the Persian Gulf from January 2021 to match our ground truth dataset. Figure 3 displays three of these images as an example, with enlargements showing individual ships at sea. Dfy Graviti – a company specializing in aerospace and maritime AI – helped us process these images to identify and locate ships.

3.2 Patterns of Suspicious Activity

Regulators enforcing sanctions advise the shipping industry (OFAC, 2023) to monitor suspicious AIS transmission gaps (”dark activity”), and geolocation ”spoofing,” where AIS data is falsified to display a ship in a false location, akin to using a VPN (Windward, 2020). Such spoofing, reported by the UN in 2019 (U.N. Doc. S/2019/171), has rapidly proliferated since (New York Times, 3 Sep. 2022, Economist, 16 Apr. 2022). Figure 1 illustrates spoofing and dark activity for two tankers in our Iranian sanctions violators dataset. In the top panel, the tanker's AIS signals falsely indicated its presence in the northern Persian Gulf while it was actually in Iran according to our ground truth data. The tanker in the bottom panel ceased AIS transmissions for four days during which it was also observed in Iran.

3.3 Detection: Predictor Variables

3.3.1 Predictors using AIS Data

In line with previous research, regulatory guidance (OFAC, 2020), and industry experts (e.g., Wolsing et al. (2022)), we construct six categories of sanctions evasion predictors, as follows:

- Identity Change: equals one if a tanker has changed its MMSI number or flag in the past six months.
- Risky Flag: equals one if a tanker is registered under a flag frequently linked to past violations.
- Ship-to-ship Transfer (STS): A deceptive practice where a tanker loads oil from another vessel away from ports to evade scrutiny.
- Irregular Trajectory: equals one if a tanker's AIS trajectory displays abnormal patterns indicative of spoofing. The top left panel of Figure 5 shows an example from the Persian Gulf in January 2021, confirmed as spoofing based on our proprietary data.
- DBSCAN Outlier: equals one if a tanker sends signals from anomalous locations, identified using the data-driven clustering method DBSCAN.

- Dark Activity: equals one if there is an unusually long interval between two AIS signals, as shown in the bottom panel of Figure 1.

3.3.2 Predictors using Sentinel Satellite Data

Ideally, spoofing could be detected if a satellite image at the time of a tanker's AIS signal shows no ship at the transmitted location. Similarly, dark activity could be detected by tracking, using satellite imagery, a tanker's actual trajectory during a dark period (e.g., if heading towards an Iranian port).[4] However, this ideal scenario would require frequent, if not continuous, satellite observations of the area of interest. However, such observations are quite sparse in reality – the satellites generating our data revisit each area only at six-day intervals.

To address this lack of simultaneous AIS signals and satellite imagery, we interpolate tanker trajectories in AIS. For example, if a Sentinel image captures a region at 10 am and the closest AIS signals from a target tanker are at 7 am and 1 pm, we use available trajectories of other tankers with similar start and end points, and similar travel time between such points, to estimate the tanker's likely location at 10 am. Figure 4 illustrates this interpolation (detailed methodology is presented in the Internet Appendix B.2). If no tankers with similar trajectories are found, spoofing is suspected. For dark activity detection, we examine tankers that go dark and then resume AIS signals at about the same place, defining search areas around these spots. If a satellite image taken during the dark period shows no tanker, wrongdoing is likely suspected.

We add the following satellite-based predictors to our AIS-based list:

- Satellite Detection: (a) For spoofing, equals one for day t if a tanker is not seen in its search area in $[t-r, t+r]$ where $r = 0$, 1, or 2. (b) For dark activity, equals one if a tanker is not seen in its search area during its dark period.

The flexibility of machine learning allows for various versions and non-linear combinations of these variables. Detailed descriptions of variable construction and summary statistics are presented in Internet Appendix B.

3.4 Detection: Methods

For spoofing detection, we create a tanker-day sample where the dependent variable is one if a tanker is observed in Iran, per our proprietary dataset, but its AIS signals indicate another

[4] Alternative detection methods, such as identifying ships at port are infeasible, given that satellite images of standard resolutions cannot distinguish between multiple vessels of similar size and shape. Furthermore, sanctions violators sometimes have altered their vessels' appearance (IHS-ACSS, 2022).

location; it is zero otherwise. For detecting dark activity, we construct a sample where the dependent variable is one if a tanker is observed in Iran during a dark period exceeding 24 hours, with zero otherwise.

To optimally combine predictors, potentially in a non-linear manner, we employ decision trees and neural networks, using 10-fold cross-validation to fine-tune model parameters. The sample is randomly divided into 10 sub-samples; models are trained with nine sub-samples, leaving one for validation. This procedure is repeated 10 times to select the best hyper-parameter combination based on McFadden's pseudo-R^2.

Detection performance is evaluated using the pseudo-R^2 from the best cross-validated model as well as the implied false positive rate. The false positive rate represents the proportion of business a firm must forgo to keep the likelihood of engaging with a violator below a set threshold. Further details are in Internet Appendix C.

3.5 Results

Results from our machine learning models are detailed in Table 1. Panel A reveals that the best AIS-based detection models achieve a pseudo-R^2 of about 10% for detecting spoofing and dark activity, with high implied false positive rates of approximately 50% to avoid 95% of true violators. This means firms must forgo over half their potential counterparties to maintain a 5% or lower risk of engaging with violators. Panel B shows that incorporating satellite data enhances detection models only modestly. The false positive rates decrease slightly; but to maintain a 5% or lower risk of engaging with violators, firms must still forgo 35.8-38% of business from compliant tankers.

Panel C evaluates the incremental contribution of each predictor category. While overall pseudo-R^2 values remain low, the Dark Activity and Ship-to-ship Transfer predictors are somewhat significant, with incremental false positive rates of 24.5% and 14.8%, respectively.

We conclude that, despite combining in a sophisticated manner numerous predictors mentioned in industry and regulatory sources, detection of sanctions evaders is highly inaccurate, which can adversely impact compliant firms. Our analysis, limited to Iran and the Persian Gulf, suggests that broader applications, such as in larger regions like Russia, could be even more challenging.

3.6 Does this mean detection is truly difficult?

One might question whether our results indicate a general difficulty in detection or merely reflect limitations in our data or models.

First, closer examination of our data reveals that sanctions-evading predictors exhibit similar patterns for both violating and compliant tankers. E.g., as shown in Figure 5, compliant tankers (top-right) can show irregular trajectories similar to those of violators (top-left), and the dark time distributions for both tanker types overlap significantly. Legitimate factors like weather, mechanical issues, or AIS data errors (Weng et al. (2022)) contribute to this low signal-to-noise ratio, compounded by sparse satellite data.

Second, even if predictive accuracy were doubled with better data, firms would still need to forgo about 20% of clients to maintain a 5% sanctions-risk exposure, as detailed in Internet Appendix C.

Third, evidence of detection challenges abound, even beyond our data and methods. For instance, only 13 of the 33 tankers from our Iran-sanctions violators dataset appear on UANI's "The Ghost Armada" list, as of April 2024, illustrating the broader difficulty in identifying violators.[5]

Overall, detecting sanctions-evading tankers is challenging for third parties. While governments and certain corporations with superior data access (e.g., Starlink) may not face these issues, most third parties lack access to such data but are still required to comply.

4 Public disclosure of suspect tankers

If the challenge of detecting sanctions evaders indeed undermines sanctions' effectiveness, two data features should be evident: (1) evaders are difficult to detect despite honest efforts, and (2) if detection improves – e.g., as through gaining access to a reasonably accurate list of suspected violators – those firms previously engaging with these suspects should change their behavior. We have established (1) in the previous section; here, we address (2).

4.1 The disclosure event

To assess the impact of publicly disclosing sanctions violators, we examine Windward.AI's release of Sanctions Risk Labels for oil tankers on the Refinitiv Eikon platform in August 2023. This classification system categorizes the global fleet into four risk levels–Low, Moderate, High, and Sanctioned–based on factors such as ID manipulation, flag hopping, and dark activity (see Internet Appendix D, and Figure A-2 therein for an example of Windward's data on a particular tanker).

Windward leverages over 100 million daily data points, integrating proprietary satellite

[5] UANI (United Against Nuclear Iran) is a bi-partisan, non-profit organization in the US; they started publishing their list in November 2020, and update it periodically.

imagery and weather data with information from its clients, which include the United Nations, the European Border and Coast Guard Agency (Frontex), and U.S. agencies such as the Drug Enforcement Administration and the Office of Naval Intelligence (Reuters, 17 Mar. 2016, Wired, 18 Mar. 2020). Observers have compared Windward's technology to military-grade signals intelligence adapted for commercial use (RAND, 2017).

Before August 2023, similar risk assessments were available from firms like Windward and Lloyd's List, but their high costs restricted access to large market players. However, after Windward's disclosure, this data became significantly more affordable–a Refinitiv Eikon subscription costs £280 per month.[6]

Beyond expanding access, this disclosure also influenced higher-order beliefs by making risk classifications common knowledge (Aumann (1976)). Previously, firms aware of the suspicious nature of activities of some tankers could not assume that others had the same information. Now, if these tankers were classified as High-risk by Windward, they know that this intelligence is widely accessible. Consequently, even firms with prior access may alter their behavior, responding not just to the data itself but to the broader market dynamics shaped by its broader availability.

4.2 Is this a sizeable information shock?

Windward's suspect tanker list could shape market outcomes if: (1) Market participants were aware of and attentive to the disclosure, and (2) The list provided new, more accurate information on sanctions violators.

First, as Figure A-3 in the Internet Appendix illustrates, Windward.AI's homepage saw a sharp increase in page views following the list's release, signaling heightened market attention.

Second, for the list to introduce new information, it must surpass alternative detection methods significantly in accuracy. We compared its accuracy against our models, as well as against UANI's suspect tankers list, using our Iranian ground truth dataset. Table 2 shows that Windward classified 27 out of 33 known violators as high/moderate risk, whereas our models identified only about half, even at a 90% confidence level. While these results are not directly comparable – since Windward may have flagged tankers based on activity beyond January 2021 – its significantly higher detection rate is notable. Moreover, even UANI's "The

[6] Windward continues to offer tailored products to commercial and government clients, and its overall sales increased following the disclosure, as it secured new contracts by demonstrating its capabilities through these Risk labels (https://www.proactiveinvestors.co.uk/companies/news/1029531/windward-shares-swell-on-ai-maritime-contract-wins-1029531.html). Thus, the release of suspect tanker data likely served to showcase Windward's expertise, rather than as a primary revenue driver for the firm.

Ghost Armada" list – also based on activity not restricted to January 2021 – recognized just 10 of the 33 violators at the time of Windward's disclosure, further underscoring Windward's superior accuracy in identifying sanctions violators.

4.3 Windward's effect on the difficulty of identifying violators

We first assess whether Windward's list helps third parties predict which tankers are about to be sanctioned in the near future. Such predictive ability is crucial for businesses. For example, if a bank ends up unknowingly lending to such a tanker, it might stand to lose the lent amount.

Analyzing tankers sanctioned between August 2023 and April 2024, we compare the predictive power of public data against Windward's list. We focus on the 3,500 tankers labeled high-/moderate-risk by Windward in July 2023 and compare them to low-risk tankers (Internet Appendix Table A-3 presents the sample filtering process in detail).

Using public data as in Section 3, we estimate sanction probabilities through a training-testing procedure, training models on data through January 2023 and testing against sanctions from August 2022 to July 2023. Additionally , predictors here include UANI's list, as well as satellite-based indicators extended to the Black and Baltic Seas for Russian sanctions-related activity.

To compare Windward's predictive ability, we rank tankers by our model's probabilities and select the top 3,500. Among 69 tankers sanctioned about-to-be-sanctioned, i.e., eventually sanctioned post-July 2023, this list of the top-3500 from our model (which also incorporates UANI's list) only contains 33. On the other hand, Windward's list contains 65, demonstrating its substantially superior accuracy.

Finally, combining both approaches improves prediction. Public data alone yields a pseudo-R^2 of 7.5% and a false positive rate of 51%, while Windward's list achieves 17.2% and 26%, respectively. The optimal mix ($0.58 \times$ public data $+ 0.42 \times$ Windward) further improves performance, reaching a pseudo-R^2 of 21% and a false positive rate of 20%, confirming that public data-based ML still has value even when Windward's list is available. Figure A-4 in the Internet Appendix shows full details on how different relative weights on these two approaches affects overall pseudo-R^2s.

4.4 Real effects of the disclosure on suspect tankers

We examine the real effects of Windward's disclosure on suspect ships – critical for sanctions enforcement – by analyzing its impact on tanker rental rates (fixtures), tanker route changes, ownership transitions, and charterer identities.

A. Changes in tanker earnings

We start by estimating the effect of Windward's high/moderate risk classification on affected tankers' fixtures (i.e., freight rates or ship rental prices, negotiated between a shipowner and a charterer) in a Differences-in-differences (DiD) framework. Fixtures are measured in standardized Worldscale units, allowing for easy comparison across contracts (see Internet Appendix E for details). We examine the change in fixtures around Windward's disclosure for tankers labeled as high/moderate risk, and compare them to counterfactual tankers – otherwise identical ones, but labeled by Windward as low risk. The challenge lies in constructing such valid counterfactuals, and we use two methods to address this issue.

First, we use propensity score matching (PSM), introduced by Rosenbaum and Rubin (1983); this method looks to match a treated, i.e., a high/moderate risk tanker, to a low risk one which had the same ex-ante propensity of being classified as high/moderate risk based on public data. We present summary statistics and balance tests – which show that the matched tankers are ex-ante very similar to the treated ones – in Internet Appendix Table A-4. This method is intuitive, and allows for standard DiD plot visualization.

In this framework, treatment effects for high-risk tankers are estimated as follows:

$$fixture_{c,i,k,t} = \left[\sum_{l=-6}^{7} \beta_l^H \times high\_risk_i \times \mathbb{I}_{\{t=l\}}\right] + \alpha_i + \gamma_{k,t} + \epsilon_{c,i,k,t}, \tag{1}$$

where $fixture_{c,i,k,t}$ is the WS rate for contract c, tanker i (with tanker type k) in month t; $high\_risk_i$ equals one for tankers labeled as high-risk in Windward's list, and zero otherwise; α_i is tanker fixed effects, and $\gamma_{k,t}$ is time×tanker-type fixed effects. We are interested in the average treatment effect on the treated (ATT) for high-risk tankers, given by the coefficient series $\{\beta_l^H\}_{l=-6}^{7}$. To keep interpretation simple, we first exclude tankers classified as moderate risk from the sample.

Figure 6 shows that the PSM-DiD estimated coefficients in the pre-period $\{\beta_l^H\}_{l=-6}^{-1}$ are insignificant, indicating a lack of pre-trends. Post-disclosure, fixture rates for high-risk tankers drop immediately in August by about 20 WS units compared to the matched control group, a change that persists until the end of our sample in March 2024.

To assess the impact on tanker earnings, consider that the average fixture for high-risk tankers in July 2023 was 126.8 WS units (Internet Appendix Table A-4). A 20-unit drop in August implies a 15.8% revenue decrease (20/126.8). Note that this decrease reflects the market reaction to the Windward-induced change in the risk of tankers violating sanctions, not actual violations.

Post-period coefficients in Figure 6 are uniformly lower than the pre-period ones, but month-by-month estimation is noisy. To mitigate noise, we aggregate post-period indicators into a single dummy $\mathbb{I}_{\{t\geq 0\}}$, and estimate ATT separately for high-risk and moderate-risk tankers. In a matched sample with high- and low-risk tankers, we run the following regression:

$$fixture_{c,i,k,t} = \beta^H \times high\_risk_i \times \mathbb{I}_{\{t\geq 0\}} + \alpha_i + \gamma_{k,t} + \epsilon_{c,i,k,t}, \tag{2}$$

where β^H is the ATT for high-risk tankers. We perform a similar estimation for moderate-risk tankers by replacing $high\_risk_i$ with $moderate\_risk_i$. Panel A of Table 3 presents the PSM-DiD regression results.

Second, we also employ a semiparametric DiD estimator as in Abadie (2005), which requires weaker identification assumptions than PSM (see Heckman et al. (1997) and Abadie (2005)). Further details are in Internet Appendix E.

Panel B of Table 3 shows results from this method. The sample size in panel B is smaller than in panel A, because Abadie (2005)'s estimation uses a tanker-level sample (not a tanker-month level one), and also excludes tankers without both pre- and post-period fixtures (because the dependent variable is post-period average fixture minus pre-period average fixture).

The results show that, after the information shock, the fixtures of high-risk tankers decreased by an average of 16.45 WS units, ranging from 13.75 to 20.71 units. This translates to a 13% earnings drop (16.45/126.8), relative to the counterfactual. No significant change is observed for moderate-risk tankers.

These results are not only robust across methods and bandwidth choices, but remain similar when we use logistic regression for propensity scores, drop August 2023 (to allow about two weeks for the disclosure effect to come to force), omit Russia, and bootstrap standard errors. We present these robustness tests in Table A-5 in the Internet Appendix.

B. Tanker route changes

Here we examine the impact of the information shock on tanker routes, in particular, we study any change in the relative propensity of treated (i.e., high/moderate risk) tankers to show up near Iran, Russia, or Venezuela (before U.S. sanctions were removed in 2024). Territorial seas, extending 12 nautical miles (nm) from the coast, define the boundaries for these countries (we use data from Flanders Marine Institute).

The dependent variable is a zero-one indicator of AIS signals emitted within the territorial seas of these countries each month. Our test design mirrors the fixtures analysis from the previous section. Figure 7 shows estimated coefficients $\{\beta_l^H\}_{l=-6}^{-1}$, which show probabilities of passing near sanctioned countries. Pre-period coefficients show a lack of pre-trends. Post-period, there is a gradual drop reaching about 10 percentage points at the end of our sample.

Table 4 presents the estimation of ATT on route changes, combining post-period months into a single dummy as in Table 3. As AIS signals are observed much more frequently than fixtures, this route test uses more observations.

The PSM-DiD results in Panel A show that high-risk tankers are about seven percentage points less likely to pass close to sanctioned countries, with no significant effect on moderate-risk tankers. The average probability of signals in these areas is 40 percentage points for high-risk tankers in July, but drops by 25% (10/40) at the end of our sample. Panel B using Abadie (2005)'s semiparametric DiD yields similar results. Overall, our evidence suggests that high-risk tankers avoid approaching sanctioned countries after Windward's disclosure.

C. Tanker ownership changes

Did tankers newly classified as high/moderate-risk become harder to sell after Windward's disclosure? This hypothesis is based on the face that many sanctions regimes explicitly prohibit transactions involving violating entities (e.g., see Council Regulation (EU) 833/2014 on such restrictions, Skadden, 2024).

We use proprietary tanker ownership data from Drewry, showing owner names and countries in six snapshots between December 2022 and March 2024. The dependent variable here equals one if owner names differ across two snapshots and zero otherwise. To account for different time spans between snapshots, we annualize the ownership change variable.

Figure 8 and Table 5 present the results, in the same format as the previous two tables. Our evidence shows that that after the information shock, high-risk (moderate-risk) tankers are 16.4-18.5 (14.6-16.8) percentage points less likely to change owners. This represents a 37% drop in liquidity for high-risk tankers, relative to a pre-disclosure annualized turnover of 45.4 percentage points (17/45.4). Overall, Windward's disclosure has reduced the turnover of high/moderate-risk tankers, making them less desirable.

D. Changes in charterers

Did U.S. and U.S.-allied charterers start avoiding tankers classified as high or moderate risk by Windward? We obtain charterers' names from fixtures data and determine the countries of their headquarters via LinkedIn or company websites. U.S.-allied charterers include companies from the UK, European Union, Australia, New Zealand, or Japan. The dependent variable is the number of tankers of a given risk type that each charterer employs each month. The regression specification is

$$no.\_tankers_{h,t} = \left[\beta^U \times US\_charterer_h + \beta^A \times US\_allied\_charterer_h\right] \times \mathbb{I}_{\{t\geq 0\}} \\ +\alpha_h + \gamma_t + \epsilon_{h,t}, \qquad (3)$$

where h denotes a charterer.

Table 6 presents the estimation results for Eq.(3). Similar to our previous results, we find a significant drop in high-risk tankers usage by U.S. charterers, while effects for U.S.-allied charterers are not statistically significant. For perspective, the average U.S. charterer employed 0.652 high-risk tankers in July 2023, and the estimated 0.358 drop represents a 54.9% reduction (0.358/0.652).[7]

Finally, we examine whether selection issues affect our data, specifically if the information shock altered high/moderate-risk tankers' reporting of their fixture contracts to our data vendor. Results in Internet Appendix Table A-7 show no evidence of changes in reporting by risky tankers. Note that our ownership change data does not come from reports to Refinitiv, and is not affected by such potential selection issues.

5 Sanctions and Disclosure through the Lens of a Model

In this section we develop a simple dynamic model of oil shipping, calibrated using market data and our Windward event-study results. This model enables us to quantify the impact of Windward's disclosure on two key actors for whom data is unavailable. First, we estimate its effect on freight rates paid by traders of sanctioned oil (henceforth, "rogue exporters") and the number of shadow tankers they employ. Second, we assess its consequences for earnings and asset values of actual sanctions-violating tankers, by utilizing information on the accuracy of Windward's risk classification (while we observe each tanker's Windward risk labels, we lack information on which vessels have definitively violated sanctions beyond those identified in our Iranian dataset). Additionally, the model helps us evaluate how the disclosure influences the overall cost of legally traded oil shipments, and clarifies the mechanisms driving these effects.

[7] In Internet Appendix Table A-6, we list these U.S. charterers, but find no clear patterns (e.g., private companies vs. public companies) related to changes in their usage of high-risk tankers.

5.1 Theoretical Framework

5.1.1 Environment

The oil shipping sector consists of exporters and tankers. Exporters are classified as either 'Rogue' (R), i.e., dealing in sanctioned oil (e.g., Iranian oil companies), or 'Clean' (C), i.e., dealing in non-sanctioned, legally-traded oil (e.g., US oil companies).

Tankers are categorized as either 'Bad' (B), i.e., those that have previously carried sanctioned oil, or 'Good' (G), i.e., those that have never done so. A Bad tanker cannot revert to Good; but a Good tanker becomes irreversibly Bad if it carries sanctioned oil. This assumption reflects the fact that tankers can be penalized for past violations. Importantly, a Bad tanker is not necessarily sanctioned and on a designated list – it simply has a history of sanction breaches and faces potential future designation. Even so, a Clean exporter employing such a tanker still risks penalties and disruptions – e.g., if the tanker is sanctioned while in use.

Rogue exporters know which tankers are Bad (e.g., because they have employed these tankers previously). Clean exporters only know that a proportion λ of all tankers are Bad (λ is common knowledge). These exporters perform due diligence using a detection technology which labels tankers as 'Low-risk' or 'High-risk'.

This detection technology classifies λ tankers as High-risk (H), and the rest as Low-risk (L). The correlation $\rho \in [0,1]$ between the true tanker type (B or G) and its risk label (H or L) measures Clean exporters' precision in detecting sanction-violating tankers. We define the following conditional probabilities

$$\theta_B \equiv \mathbb{P}[H|B] = \lambda + (1-\lambda)\rho, \quad \theta_G \equiv \mathbb{P}[H|G] = \lambda(1-\rho), \tag{4}$$

from which it also follows that

$$\mathbb{P}[B|H] = \theta_B, \quad \mathbb{P}[B|L] = \theta_G. \tag{5}$$

We assume that detection errors are i.i.d. conditional on tanker type – i.e., an H or L label is assigned each period depending only on a tanker's true type. So once a Good tanker violates sanctions and becomes Bad, its probability of being classified as High-risk increases permanently.

Tankers not only know if they are Bad, but can also infer their risk label from the fixtures quoted to them by Clean exporters. Rogue exporters do not know exact price quotes offered to a tanker by competing Clean exporters (and hence cannot infer risk labels). We assume a stationary environment where the proportion of Bad tankers, λ, is constant.

5.1.2 Sanctions

The sanctioning authority has limited resources and screens a subset of tankers for violations, and sanctions those that they find guilty. We define screening intensity as the probability that a given tanker is investigated by the sanctioning authority. The intensity with which tankers are screened for past violations (w) may be different from that for current-period violations (δ).

Third parties – those not involved in sanctions violations themselves (e.g., a Clean charterers) – account for the fact that the probability of being screened and sanctioned is higher for tankers with risk label H (assuming that their detection technology is not completely uninformative). From their perspective, the probability of a tanker being sanctioned is

$$w_x \cdot \mathbb{I}\{\text{Previously violated sanctions}\} + \delta_x \cdot \mathbb{I}\{\text{Currently carrying sanctioned oil}\}, \tag{6}$$

where w_x (δ_x) is the sanctioning authority's screening intensity for past (current) violations for a tanker with risk label x. Thus, a Bad tanker (i.e., previously violated sanctions), which is labeled H and currently carrying sanctioned oil has an estimated sanction probability $w_H + \delta_H$.

$$w_0 = \int_0^1 w_i di, \quad \delta_0 = \int_0^1 \delta_i di \tag{7}$$

Similarly, we can define

$$w_{GL}, w_{GH}, w_{BL}, w_{BH}, \quad \delta_{GL}, \delta_{GH}, \delta_{BL}, \delta_{BH} \tag{8}$$

we assume

$$w_L = w_{GL} = w_{BL}, \quad w_H = w_{GH} = w_{BH}, \quad \delta_L = \delta_{GL} = \delta_{BL}, \quad \delta_H = \delta_{GH} = \delta_{BH} \tag{9}$$

Therefore,

$$w_0 = \lambda w_H + (1-\lambda) w_L, \quad \delta_0 = \lambda \delta_H + (1-\lambda)\delta_L. \tag{10}$$

where w_x (δ_x) is the probability with which the sanctioning authority is able to gather sufficient evidence to sanction a tanker with risk label x for past (current) violations. Thus, a Bad tanker (i.e., previously violated sanctions), which is labeled H and currently carrying sanctioned oil has an estimated sanction probability $w_H + \delta_H$.

The total screening intensities, weighted by the proportion of H and L tankers, are

$$w_0 = \lambda w_H + (1-\lambda) w_L, \quad \delta_0 = \lambda \delta_H + (1-\lambda)\delta_L. \tag{11}$$

w_0 and δ_0 depend on the resources that the sanctioning authority can spend to uncover violations, which we assume to be constant in the model. Note that the sanctioning authority's information set could be different from the market's. Therefore, w_H, w_L, δ_H and δ_L should be interpreted as the market's best guess of the authority's screening intensities based on the market's information set. Different H-labeled tankers may have different screening intensities under the sanctioning authority's information set, but the market cannot distinguish between them and only expects an average intensity w_H. We assume that the two information sets are correlated and the market expects that the sanctioning authority allocates investigative resources towards H and L tankers according to the proportion of bad tankers across these sets. That is:

$$w_H/w_L = \delta_H/\delta_L = \mathbb{P}[B|H]/\mathbb{P}[B|L]. \tag{12}$$

Under this assumption, if $\rho > 0$, then H labels are more likely associated with Bad tankers, implying $\mathbb{P}[B|H] > \mathbb{P}[B|L]$, and hence $w_H > w_L$ and $\delta_H > \delta_L$. If risk labels are pure noise, i.e., $\rho = 0$ and $\mathbb{P}[B|H] = \mathbb{P}[B|L]$, then $w_H = w_L$ and $\delta_H = \delta_L$, i.e., the H and L labels are not associated with any difference in screening intensities.

Combining equations (5), (11) and (12), we obtain

$$w_H = \frac{\theta_B}{\lambda} w_0, \quad w_L = \frac{\theta_G}{\lambda} w_0, \quad \delta_H = \frac{\theta_B}{\lambda}\delta_0, \quad \delta_L = \frac{\theta_G}{\lambda}\delta_0. \tag{13}$$

E. Timeline

We assume the following sequence of events each period:

1. Clean exporters assign risk labels to tankers using the best existing technology.

2. An aggregate mean-zero shock $\tilde{\epsilon}$ to the oil trading revenue of all exporters occurs, reflecting unpredictable oil price changes that exporters cannot hedge ex-ante. Additionally, each tanker i experiences a tanker-specific shock $\tilde{c}_i$ for is realized, affecting costs when dealing with Rogue exporters.

3. Clean and Rogue exporters simultaneously choose freight rates (fixtures), given their information on tanker labels/types: Clean exporters choose p_L and p_H for L and H

tankers, respectively, and Rogue exporters choose p_G^R and p_B^R for Good and Bad tankers, respectively.

4. Tankers observe these freight rates and decide whether to engage with Clean or Rogue exporters.

5. Each player receives a payoff and proceeds to the next period. Tankers continue to the next period only if not sanctioned for current or past violations.

5.1.3 Exporters' optimization

The per-unit oil trade revenue for Rogue exporters, $\tilde{r}^R$, and for Clean exporters, $\tilde{r}$, are

$$\tilde{r}^R = r^R + \tilde{\epsilon}, \quad \tilde{r} = r + \tilde{\epsilon}, \tag{14}$$

where r^R and r are mean values and $\tilde{\epsilon}$ is an aggregate shock uniformly i.i.d. over $[-\sigma, \sigma]$. Therefore, $\tilde{r}^R$ and $\tilde{r}$ share the common shock $\tilde{\epsilon}$ but can have different means. Since we only model the oil shipping market - and not the oil market itself - we assume for simplicity a fixed ratio $\frac{r^R}{r}$, justified by the stable discounts for sanctioned oil in the months after Windward's disclosure (see, e.g., Inside Shipping, 2024 – there was no increase in the discount on Russian Ural oil until October 2023).

Exporters set freight rates to maximize profits. They do this period by period, given our i.i.d. environment and assuming there is no value to past 'relationships' between exporters and tankers. Their costs include payments to tankers and expected sanctions penalty. For Rogue exporters, this penalty is z^R times the sanction probability.

Let Q_{ij}^R denote the supply of tankers of type i and risk label j dealing with Rogue exporters. Implicitly Q_{ij}^R is a function of p_G^R and p_B^R, which are optimally chosen. After observing $\tilde{r}^R$, Rogue exporters maximize profits:

$$\begin{aligned}\max_{\{p_G^R, p_B^R\}} \; & Q_{GL}^R(\tilde{r}^R - p_G^R - \delta_L z^R) + Q_{GH}^R(\tilde{r}^R - p_G^R - \delta_H z^R) \\ & + Q_{BL}^R[\tilde{r}^R - p_B^R - (\delta_L + w_L)z^R] + Q_{BH}^R[\tilde{r}^R - p_B^R - (\delta_H + w_H)z^R].\end{aligned} \tag{15}$$

The expected sanction penalty for Clean exporters differs from that for Rogue exporters. This is because (i) tankers hired by Clean exporters do not face the additional sanction probability, δ_L or δ_H, as they do not violate sanctions in the current period, and (ii) the penalty z for a Clean exporter using a Bad tanker that gets sanctioned in the current period is smaller than the z^R for Rogue exporters. For instance, a Rogue exporter's entire oil cargo could be impounded if a tanker carrying it is sanctioned, which is significantly more costly than the delays and penalties for insufficient due diligence faced by a Clean exporter.

Let Q_{ij} denote the supply of tankers of type i and risk label j dealing with Clean exporters. Implicitly, Q_{ij} is a function of p_L and p_H, which are optimally chosen. After observing $\tilde{r}$, Clean exporters maximize:

$$\max_{\{p_L, p_H\}} \quad Q_{GL}(\tilde{r} - p_L) + Q_{GH}(\tilde{r} - p_H) + Q_{BL}(\tilde{r} - p_L - w_L z) + Q_{BH}(\tilde{r} - p_H - w_H z). \quad (16)$$

Importantly, z^R is bounded, due to practical limitations to sanction enforcement (at most, the sanctioning authority can impound a tanker carrying sanctioned oil and its cargo). This is the key feature that makes detection accuracy important in the model. If z^R could be increased without bounds, violations would be prevented even with noisy detection, in the spirit of Becker's classic intuition on crime and punishment (Becker, 1968).

5.1.4 Tankers' optimization

Tankers are price-takers – they observe freight rates (fixtures) quoted to them by Clean and Rogue exporters and decide whom to engage with. Because a Good tanker becomes irreversibly a Bad tanker if it engages with Rogue exporters, tankers must consider their continuation values when making decisions.

Let V_G and V_B denote the unconditional values of Good and Bad tankers, and V_{GL}, V_{GH}, V_{BL}, V_{BH} denote the values conditional on risk labels. The Bellman equations for tankers are

$$V_{GL} = \mathbb{E}\left[\max\begin{Bmatrix} p_L + \beta V_G \\ p_G^R - \tilde{c}_i - \delta_L s + \beta V_B \end{Bmatrix}\right], \quad V_{BL} = \mathbb{E}\left[\max\begin{Bmatrix} p_L - w_L s + \beta V_B \\ p_B^R - \tilde{c}_i - (w_L + \delta_L) s + \beta V_B \end{Bmatrix}\right],$$
$$V_{GH} = \mathbb{E}\left[\max\begin{Bmatrix} p_H + \beta V_G \\ p_G^R - \tilde{c}_i - \delta_H s + \beta V_B \end{Bmatrix}\right], \quad V_{BH} = \mathbb{E}\left[\max\begin{Bmatrix} p_H - w_H s + \beta V_B \\ p_B^R - \tilde{c}_i - (w_H + \delta_H) s + \beta V_B \end{Bmatrix}\right],$$

where V_G and V_B are given by

$$V_G = \theta_G V_{GH} + (1 - \theta_G) V_{GL}, \quad V_B = \theta_B V_{BH} + (1 - \theta_B) V_{BL}. \quad (17)$$

When dealing with Rogue exporters, tanker i incurs in each period a random cost $\tilde{c}_i$, assumed to be uniformly i.i.d. over $[0, \bar{c}]$. This cost, observed only by the tanker before choosing an exporter to deal with, allows us to obtain interior solutions for tanker supply. Such a cost may arise from efforts to avoid detection and can reflect a tanker's location, route, and evasion technology or uncertainties concerning financing or other relations with third parties wary of falling under secondary sanctions.

The penalty imposed on a Bad tanker when sanctioned is s, assumed for simplicity to equal the tanker's entire value. Given discounting, this penalty is $s = \beta V_B$. Critical values of $\tilde{c}_i$ at which tanker i would be indifferent between dealing with a Clean or Rogue exporter are as below:

$$\begin{aligned} c_{GL} &= p_G^R - p_L - \delta_L s + \beta(V_B - V_G), \quad & c_{BL} &= p_B^R - p_L - \delta_L s, \\ c_{GH} &= p_G^R - p_H - \delta_H s + \beta(V_B - V_G), \quad & c_{BH} &= p_B^R - p_H - \delta_H s. \end{aligned}$$

For $\tilde{c}_k$ below (above) such a critical value the tanker prefers to deal with a Rogue (Clean) exporter. Using these critical values, we derive that

$$V_G = \frac{\theta_G \mathbb{E}\left[\frac{(c_{GH})^2}{2\bar{c}} + p_H\right] + (1-\theta_G)\mathbb{E}\left[\frac{(c_{GL})^2}{2\bar{c}} + p_L\right]}{1-\beta}, \tag{18}$$

$$V_B = \frac{\theta_B \mathbb{E}\left[\frac{(c_{BH})^2}{2\bar{c}} + p_H\right] + (1-\theta_B)\mathbb{E}\left[\frac{(c_{BL})^2}{2\bar{c}} + p_L\right] - [\theta_B w_H + (1-\theta_B) w_L] s}{1-\beta}, \tag{19}$$

where the expectation is over next period's aggregate shock $\tilde{\epsilon}$.

5.1.5 Equilibrium

A stationary equilibrium is characterized by: (I) tankers' optimal decision rule based on the critical values $c_{GL}(\tilde{\epsilon})$, $c_{GH}(\tilde{\epsilon})$, $c_{BL}(\tilde{\epsilon})$, $c_{BH}(\tilde{\epsilon})$ where $\tilde{\epsilon}$ refers to the aggregate shock in the current period, and, (II) Rogue and Clean exporters' optimal freight rates $p_G^R(\tilde{\epsilon})$, $p_B^R(\tilde{\epsilon})$, $p_L(\tilde{\epsilon})$, $p_H(\tilde{\epsilon})$. In equilibrium, (i) tankers maximize their values and exporters maximize their profits, (ii) exporters' beliefs about tanker types or risk labels are consistent with tankers' decisions in aggregate, and (iii) markets clear, where the market clearing conditions for Rogue exporters (Q_{ij}^R) and Clean exporters (Q_{ij}) are

$$Q_{ij}^R = A_{ij}\left(\frac{c_{ij}}{\bar{c}}\right), \quad Q_{ij} = A_{ij}\left(1 - \frac{c_{ij}}{\bar{c}}\right), \quad ij \in \{GL, GH, BL, BH\}, \tag{20}$$

where $A_{GL} = (1-\lambda)(1-\theta_G)$, $A_{GH} = (1-\lambda)\theta_G$, $A_{BL} = \lambda(1-\theta_B)$, and $A_{BH} = \lambda\theta_B$.

5.2 Model calibration

The average number of trips per year in our fixtures sample is 3.4. We set the discount factor β to 0.9, accounting for both time discounting and tanker depreciation.[8] Based on (CRS, 2024), we set the proportion λ of Bad tankers to 0.213. We normalize Clean exporters' mean oil trade revenue (more precisely, profits before shipping costs) r per voyage to one. This revenue amounts to 15.5 million for an average tanker (i.e., 0.8 million barrels), at \$70 per barrel, with a 27.7% before-shipping profit margin (this was the oil industry's operating margin for 2023 Q2, as per CSI Market, 2023). For Rogue exporters, assuming a \$4 per barrel discount (as for Russian exports to India in the second half of 2023, as per Reuters, Sept. 2023), their mean revenue is $r^R = (70 \times 0.277 - 4)/(70 \times 0.277) = 0.794$.

Noting that OFAC's civil monetary penalties typically depend on the amount of the sanctions violating transaction (CFR, Appendix A to Part 501), we set the sanction penalty z for Clean exporters dealing with Bad tankers equal to the freight cost they pay, i.e., around \$4 million in 2023 (e.g., Inside Shipping, 2024), which implies $z = 4/15.5 = 0.258$. We set the sanction penalty z^R for Rogue exporters based on the median value from actual cases of seized oil (e.g., US Attorney's Office, 2024), resulting in $z^R = 50/15.5 = 3.226$. The correlation ρ in the model, which measures Clean exporters' classification precision in detecting sanctions-violating tankers, is obtained by converting the pseudo-R^2 of sanction predictability in Figure A-4 in the Internet Appendix. The pseudo-R^2 is 7.5% based on public data, (implying $\rho^{pre} = 0.296$), while it is 17.2% after adding Windward's risk labels to the information set (implying $\rho^{post} = 0.453$). In what follows superscripts *pre* and *post* denote the pre- and post-disclosure periods.

We calibrate the screening intensities (w_0 and δ_0), revenue volatility (σ), and maximum cost ($\bar{c}$) by matching six moments: (i) proportion of sanctioned tankers, (ii) annualized volatilities of low-risk and (iii) high-risk tankers' fixtures, relative to their mean, (iv) annualized volatility of top US exporter/charterers' revenue relative to its mean (proxy for that of Clean exporters), (v) disclosure's effect on high-risk tankers' fixtures (from the DiD estimation in Table 3), and (vi) disclosure's effect on high-risk tankers' routes (from the DiD estimation in Table 4). We construct model analogs of these moments as follows:

First, the number of sanctioned tankers in the model is $Q_{Sanc} = Q_{BL}w_L + Q_{BH}w_H + Q^R_{GL}\delta_L + Q^R_{GH}\delta_H + Q^R_{BL}(w_L + \delta_L) + Q^R_{BH}(w_H + \delta_H)$ (since we normalize the total number of tankers to one, this will give us the proportion of sanctioned tankers). The volatilities of low-risk and high-risk tankers' fixtures are $\text{std}(p_L)$ and $\text{std}(p_H)$, respectively.

Second, Clean exporters' total revenue in the model equals $(Q_{GL} + Q_{GH} + Q_{BL} + Q_{BH})\tilde{r}$.

[8] Assuming a time-discounting rate of 0.95, a 10% scrap value, and a 15-year life cycle gives $0.95 \times \exp(\log(0.1)/(15 \times 3.4)) \approx 0.9$.

We get its empirical equivalent by aggregating the revenues of the top 12 publicly traded US oil companies from their 10-K reports. The full list of large US companies involved in oil shipping, including those that we use and also some private ones, is in the Internet Appendix Table A-6.

Third, the model analog for the disclosure effect on high-risk tankers' fixtures is

$$\frac{\mathbb{E}[p^{post} - p^{pre}|\mathrm{H}^{post}]}{\mathbb{E}[p^{pre}|\mathrm{H}^{post}]}, \tag{21}$$

where p is fixture and H^{post} denotes a tanker that is labeled as high-risk in the post-period. The average pre-period fixture, conditional on a post-period high-risk label is

$$\mathbb{E}[p^{pre}|\mathrm{H}^{post}] = \mathbb{P}[\mathrm{L}^{pre}|\mathrm{H}^{post}] \times \mathbb{E}[p_L^{pre}] + \mathbb{P}[\mathrm{H}^{pre}|\mathrm{H}^{post}] \times \mathbb{E}[p_H^{pre}], \tag{22}$$

where $\mathbb{E}[p_L^{pre}]$ and $\mathbb{E}[p_H^{pre}]$ are price averages taken over the pre-period aggregate shock $\tilde{\epsilon}$ and $\mathbb{P}[.|.]$ is conditional probability. Noting that $\mathbb{P}[\mathrm{L}^{pre}|\mathrm{H}^{post}] = 1 - \mathbb{P}[\mathrm{H}^{pre}|\mathrm{H}^{post}]$, we calculate $\mathbb{P}[\mathrm{H}^{pre}|\mathrm{H}^{post}]$ by splitting it into three cases based on tanker types: (i) G^{pre} and G^{post}, (ii) G^{pre} and B^{post}, and (iii) B^{pre} and B^{post}. The formula is shown in the Internet Appendix.

Finally, we use the probability of sending AIS signals from areas near sanctioned countries (see Table 4) as a proxy for the probability of dealing with Rogue exporters in the model. Defining the probabilities for H and L tankers of dealing with Rogue exporters as:

$$\begin{aligned} b_H &= \mathbb{P}[\mathrm{B}|\mathrm{H}] \times \mathbb{P}[\tilde{c}_i < c_{BH}] + \mathbb{P}[\mathrm{G}|\mathrm{H}] \times \mathbb{P}[\tilde{c}_i < c_{GH}] = \theta_B \frac{c_{BH}}{\bar{c}} + (1-\theta_B)\frac{c_{GH}}{\bar{c}}, \\ b_L &= \mathbb{P}[\mathrm{B}|\mathrm{L}] \times \mathbb{P}[\tilde{c}_i < c_{BL}] + \mathbb{P}[\mathrm{G}|\mathrm{L}] \times \mathbb{P}[\tilde{c}_i < c_{GL}] = \theta_G \frac{c_{BL}}{\bar{c}} + (1-\theta_G)\frac{c_{GL}}{\bar{c}}, \end{aligned}$$

the disclosure's effect on high-risk tankers' routes in the model is

$$\frac{\mathbb{E}[b^{post} - b^{pre}|\mathrm{H}^{post}]}{\mathbb{E}[b^{pre}|\mathrm{H}^{post}]}, \tag{23}$$

where

$$\mathbb{E}[b^{pre}|\mathrm{H}^{post}] = \mathbb{P}[\mathrm{L}^{pre}|\mathrm{H}^{post}] \times \mathbb{E}[b_L^{pre}] + \mathbb{P}[\mathrm{H}^{pre}|\mathrm{H}^{post}] \times \mathbb{E}[b_H^{pre}]. \tag{24}$$

$\mathbb{E}[b_L^{pre}]$ and $\mathbb{E}[b_H^{pre}]$ are the averages of the respective probabilities taken over the pre-period aggregate shock $\tilde{\epsilon}$ and the conditional probability $\mathbb{P}[\mathrm{H}^{pre}|\mathrm{H}^{post}]$ is as specified above.

We solve the model by discretizing the aggregate shock $\tilde{\epsilon}$ on a grid.

Panel B in Table 7 shows that the model fits the data well. The data values vs. model analogs for the proportion of sanctioned tankers, annualized volatilities of low-risk and high-risk tankers' fixtures, and volatility of top US charterers' revenue are 0.5% vs. 0.5%, 0.77 vs.

0.77, 0.81 vs. 0.83, and 0.30 vs. 0.21, respectively. The model also captures the disclosure effect on high-risk tankers' fixture changes (13% drop in the data vs. 11% in the model) and route changes (17% drop in the data vs. 20% in the model).

Importantly, the variable's values in the model align with their real-world values, obtained from the media, financial market data, or other sources. First, the volatility of oil trade revenue derived from moments matching is $\sigma = 0.365$, consistent with the Oil Volatility Index OVX around 0.40 over 2022-2023 (OVX measures oil price volatility, and hence, earnings volatility). Second, the model implied freight rate premium paid by Rogue exporters relative to Clean exporters (i.e., $(\mathbb{E}[p_B^R]+\mathbb{E}[p_G^R])/(\mathbb{E}[p_H]+\mathbb{E}[p_L])-1$) is around 110%-125%, per Panel D in Table 7, consistent with the sanctions premium reported in the media for Russian oil delivered in 2023 to Indian and Chinese ports (e.g., Inside Shipping, 2024, which reports calculations from Argus Media). Third, model-based tanker values between \$37 and \$44 million, per Panel D in Table 7, also match actual values – a 15-year-old tanker had an average value of \$45 million in the second-hand market in 2023 (using values for VLCC, Suezmax and Aframax tankers as per Xclusiv, 2023), and a 20-year old tanker had a value of \$30 million.

Overall, model-based estimates closely match the respective values from the media, providing assurance that quantities generated within the model can be taken seriously.

5.3 Model implications

Figure 9 illustrates several implications of increased classification precision from our model. The vertical dashed lines show the pre- and post-Windward precisions, as measured by pseudo-R^2, at 7.5% and 17.2%, respectively.

Panel A of Figure 9 shows that as classification precision increases, Rogue exporters pay higher (lower) freight rates for hiring Good (Bad) tankers - i.e., the fixtures $\mathbb{E}[p_G^R]$ ($\mathbb{E}[p_B^R]$) increase (decrease), with expectations over the current-period aggregate shock. This is because Bad tankers now face on average higher screening intensity of past activity (they are more likely to be labeled as High-risk, and w_H is now higher), which makes them less attractive. Good tankers do not have past violations and hence become even more attractive.

Panel B shows the probabilities of tankers dealing with Rogue exporters, which are higher for Bad tankers than for Good tankers. Within either Good or Bad tankers, High-risk ones are less likely to deal with Rogue exporters than Low-risk ones due to higher screening intensity ($w_H > w_L$ and $\delta_H > \delta_L$ when $\rho > 0$), which implies a higher expected cost of violating sanctions for these tankers. Furthermore, as classification precision increases, the probability of dealing with Rogue exporters decreases more for High-risk tankers because their screening intensity increases more, per Equation (12) ($\mathbb{P}[B|H]$ increases, and $\mathbb{P}[B|L]$ decreases with classification precision).

The next two panels of Figure 9 shed light on aggregate implications for Rogue exporters. Panel C shows that the per-tanker spending of Rogue exporters increases, which is due mostly to the increase in freight rates for Good tankers paid by Rogue exporters (from Panel A in the figure). This increase is not offset by the drop in freight rates that they pay to Bad tankers because, overall, Rogue exporters employ more Good tankers – these are in total about four times more than the Bad tankers (λ is 0.213), while Bad tankers are at most twice as likely to be employed by the Rogue exporters (from Panel B).

Panel D shows that with higher classification precision, the total number of oil tankers engaging with Rogue exporters decreases. This follows from Panel B. Quantitatively, the information shock accounts for an 11.7% drop in Rogue exporters' total hiring of tankers, when pseudo-R^2 increases from the pre-Windward to post-Windward precision. This shortfall, of course, can translate into a drop in exports of sanctioned oil itself, but the magnitude of that drop may not be one-for-one. This is because some of these tankers can be replaced by other means, e.g., through oil pipelines – which will now have to overwork relative to the pre-period equilibrium, potentially leading to higher costs of transport and higher depreciation of these alternative transportation assets.

Next, Panel E of Figure 9 shows the relation between classification precision and the equilibrium freight rates (i.e., the fixtures $\mathbb{E}[p_L]$ and $\mathbb{E}[p_H]$) for Clean exporters. Both these rates decrease in response to better classification, because now fewer tankers overall deal with Rogue exporters, increasing tanker supply for Clean exporters. At the same time, $\mathbb{E}[p_H]$ decreases more than $\mathbb{E}[p_L]$. This is because as precision increases, the H label indicates a higher probability of a Bad tanker, so Clean exporters require a larger discount for hiring such a tanker. The left-most point of the panel takes this intuition to the extreme: when pseudo-R^2 is zero and classification is pure noise, Clean exporters set $\mathbb{E}[p_H]$ and $\mathbb{E}[p_L]$ to be equal.

Panel F shows how expected freight rates from Clean exporters change with classification precision. These rates for Good and Bad tankers are given by

$$\mathbb{E}[p|G] = \theta_G \mathbb{E}[p_H] + (1-\theta_G)\mathbb{E}[p_L], \quad \mathbb{E}[p|B] = \theta_B \mathbb{E}[p_H] + (1-\theta_B)\mathbb{E}[p_L].$$

As classification becomes more precise, the decrease in θ_G and increase in θ_B (see Equation 5) makes G (B) tankers' expected price closer to $\mathbb{E}[p_L]$ ($\mathbb{E}[p_H]$).

This result – that the increase in classification precision diminishes the freight rates received also by Good tankers, as shown in Panel F – is counter-intuitive. This arises through the following mechanism: from Panel B, that with higher precision Rogue exporters increasingly avoid Bad tankers, fearing a higher chance of detection – their usage of BH tankers drops while BL increases slightly; however, there are more BH tankers than BL (as long as $\rho > 0$). Therefore, more Bad tankers now seek employment with Clean exporters, which boosts the supply available for the latter and lowers the earnings for Good tankers who work with them. This drop is only partly offset by Rogue exporters' willingness to now pay more for these Good tankers (Panel A), because the Rogue market is considerably smaller and continues to shrink (Panel D). Tanker values show a correspondingly declining pattern (from Table 7, Panel D), because they reflect discounted values of future expected freight rates.

Finally, in Table 8, we estimate the aggregate (annual) effect of the information shock on the four relevant parties in terms of dollar values. We focus on price effects and fix quantities (i.e., the Qs in the model) to pre-Windward levels. Alternatively, we could have assumed that Rogue exports drop one-for-one with the fall in their usage of oil tankers, but we think such an assumption might exaggerate effects if Rogue exporters ship the same oil using alternative means, e.g., pipelines. Of course, one could estimate the effect of such changes in transporting methods on their cost of exports if one knew how much these alternative means cost – but we do not have this information. So we estimate costs and benefits by fixing overall export quantities, and let the information disclosure affect the market only through prices.

We analyze the changes in costs/benefits for exporters and tankers between the pre- and post-Windward periods. We find that the shipping cost for Rogue exporters increases by 2% (from 9.0 to 9.2 mln per tanker-trip) due to the increase in p_G^R, while the cost for Clean exporters decreases by 6.3% (from 4.2 to 4.0 mln per tanker-trip) due to the decreases in both p_L and p_H. On the tanker side, the fixture income decreases by 5.7% (from 4.9 to 4.7 mln per tanker-trip) and 3.3% (from 5.0 to 4.9 mln per tanker-trip) for Bad and Good tankers, respectively. To estimate the aggregate annual effect, we do the following back-of-the-envelope calculation: we multiply the above dollar values by the number of tankers and 3.4 trips per year (as explained earlier). Therefore, Windward's disclosure leads to \$780 million and \$1,536 million annual losses for Rogue exporters and Bad tankers, respectively. These losses accrue as gains for the sanctions-compliant agents, with a \$2.3 billion overall gain for Clean exporters and Good tankers combined.

Taken together, these results suggest that Rogue exporters and Bad tankers are indeed adversely impacted by Windward's disclosure. As this is a key goal for sanctions enforcement, our exercise here makes a clear case for disclosing more precise information on tanker sanction risks to the market.

5.4 Testing a key counter-intuitive implication of the model

We wish to test the counter-intuitive model prediction that after Windward's disclosure, low-risk tankers receive lower fixtures from Clean exporters due to increased competition in that market (as shown in Panel E of Figure 9).

We conduct tests exploiting charterer-level variations. Ideally, the change in charterer-level competition should be measured by the incremental tanker supply to each charterer after the information shock, but this quantity is unobservable. We circumvent this issue in designing our test by relying on the intuition of "shift-share instruments" (Bartik, 1991). Let g_k, for $k \in 1, \ldots, K$, are the varying exposures to increased competition (i.e., "shifts") of K groups of tankers. Let also s_{ik} be the proportion (i.e., "share") of group k tankers that charterer i deals with ($\sum_k s_{ik} = 1$). Then the share-weighted sum of exposures $\sum_k s_{ik} g_k$ serves as an ex-ante measure of charterer-level competition.[9]

To ensure variation in the exposures g_k, we group tankers by product type – i.e., refined products (e.g., gasoline and diesel), unrefined (e.g. crude oil), and unclassified products, hence $K = 3$. Among the refined-product tankers in Windward's list, 24% are high-/moderate-risk tankers, while this proportion is 36% for unrefined, and 34% for unclassified products tankers. This variation suggests that charterers with more unrefined or unclassified products are more likely to experience increased supply from risky tankers after Windwards' disclosure. Note that this strategy nets out other common/aggregate changes in fixtures by exploiting charterer-level variation in exposures to supply shocks.

Using this measure, we focus on low-risk tankers and run the following regression:

$$fixture_{c,i,k,t} = \beta_0 \times \mathbb{I}_{\{\text{High-competition}\}} + \beta_1 \times \mathbb{I}_{\{\text{High-competition}\}} \times \mathbb{I}_{\{t \geq 0\}} + \text{FE} + \epsilon_{c,i,k,t}. \quad (25)$$

The dependent variable is low-risk tankers' fixtures, and $\mathbb{I}_{\{\text{High-competition}\}}$ equals one if the charterer-level shift-share competition measure exceeds certain cross-sectional cutoff (e.g., top 50%, 40%, or 30%), and is zero otherwise. Our model predicts a negative DiD coefficient β_1.

Estimation results in Table 9 are consistent with our model predictions: higher competition corresponds to significantly lower freight rates for low-risk tankers after Windward's disclosure. The results are robust across various cutoffs for the high-competition indicator, with estimated effects ranging from 12.9 to 15.9 WS units. Note that our model is not forced to generate this implication and the consistent empirical evidence makes our model

[9]For this measure to work, one would need stickiness in the charterer-tanker relationship. We verify this in the data – 40% of the tankers engage with charterers that they have worked with in the preceding two years.

reassuring.

6 Conclusion

Our results indicate that Windward's sudden disclosure of information identifying suspected violators negatively impacted both rogue exporters and sanctions-violating tankers. Whether these findings suggest a need to reform sanctions-related information policy depends, of course, on whether our estimates point to a market failure. After all, Windward successfully aggregated various types of information to more accurately identify suspects, and eventually made this information public. Does this imply that the market was able to solve its own problems with identifying sanctions violators, meaning no policy changes are necessary?

One way to approach this question is by comparing the value of Windward's information to the market versus its value to sanctioning authorities. On the one hand, Windward's total market capitalization was $47 million just before its list of suspected sanctions-violating tankers was made public. This $47 million reflects the discounted value of Windward's future profits, which in turn represents the market's willingness to pay for its information. On the other hand, following Windward's disclosure, sanctioned exporters faced an increase in shipping costs of approximately $780 million per year, and sanction-violating tankers experienced a reduction in freight rates exceeding $1.5 billion (from Section 5).

The key question then becomes: how much should a sanctioning authority that is focused on enforcement be willing to pay to impose a $2.3 billion annual cost on violators? If this amount is more than the $47 million the market was willing to pay, then our findings indeed suggest a need for change in information provision.

References

Abadie, A. (2005), 'Semiparametric Difference-in-Differences Estimators', *Review of Economic Studies* **72**(1), 1–19.

Aumann, R. J. (1976), 'Agreeing to disagree', *Annals of Statistics* **4**(6), 1236–1239.

Babina, T., Hilgenstock, B., Itskhoki, O., Mironov, M. and Ribakova, E. (2023), 'Assessing the impact of international sanctions on Russian oil exports', *Available at SSRN* .

Bartik, T. J. (1991), *Who Benefits from State and Local Economic Development Policies?*, W.E. Upjohn Institute for Employment Research.

Becker, G. (1968), 'Crime and punishment: An economic approach', *Journal of Political Economy* **76**, 169–217.

Brancaccio, G., Kalouptsidi, M. and Papageorgiou, T. (2020), 'Geography, transportation, and endogenous trade costs', *Econometrica* **88**(2), 657–691.

Brancaccio, G., Kalouptsidi, M. and Papageorgiou, T. (2023), 'The impact of oil prices on world trade', *Review of International Economics* **31**(2), 444–463.

Brown, P. (2020), *Oil market effects from US economic sanctions: Iran, Russia, Venezuela*, Vol. 3, Congressional Research Service Washington, DC, USA.

Clayton, C., Maggiori, M. and Schreger, J. (2023), A framework for geoeconomics, Working Paper 31852, National Bureau of Economic Research.

Crozet, M., Hinz, J., Stammann, A. and Wanner, J. (2021), 'Worth the pain? Firms' exporting behaviour to countries under sanctions', *European Economic Review* **134**.

Dellavigna, S. and La Ferrara, E. (2010), 'Detecting illegal arms trade', *American Economic Journal: Economic Policy* **2**(4), 26–57.

Dimmock, S. G. and Gerken, W. C. (2012), 'Predicting fraud by investment managers', *Journal of Financial Economics* **105**(1), 153–173.

Dizaji, S. F. and van Bergeijk, P. A. G. (2013), 'Potential early phase success and ultimate failure of economic sanctions: A VAR approach with an application to Iran', *Journal of Peace Research* **50**, 721–736.

Duggan, M. and Levitt, S. D. (2002), 'Winning isn't everything: Corruption in sumo wrestling', *American Economic Review* **92**(5), 1594–1605.

Early, B. (2021), Making sanctions work: promoting compliance, punishing violations, and discouraging sanctions busting, *in* P. A. van Bergeijk, ed., 'Research Handbook on Economic Sanctions', E. Elgar Publishing, Northhampton, MA, USA, pp. 167–186.

Early, B. R. and Peterson, T. M. (2022), ‘Does punishing sanctions busters work? sanctions enforcement and us trade with sanctioned states’, *Political Research Quarterly* **75**(3), 782–796.

Early, B. R. and Preble, K. (2019), ‘Enforcing Economic Sanctions: Analyzing How OFAC Punishes Violators of U.S. Sanctions’, *Working Paper* .

Eaton, J. and Engers, M. (1992), ‘Sanctions’, *Journal of Political Economy* **100**, 899–928.

Eaton, J. and Engers, M. (1999), ‘Sanctions: Some simple analytics’, *American Economic Review* **89**, 409–414.

Efing, M., Goldbach, S. and Nitsch, V. (2023), ‘Freeze! Financial sanctions and bank responses’, *Review of Financial Studies* **36**, 4417–4459.

Farzanegan, M. R. (2011), ‘Oil revenue shocks and government spending behavior in Iran’, *Energy Economics* **33**(6), 1055–1069.

Fisman, R. J., Marcolongo, G. and Wu, M. (2024), ‘The undoing of economic sanctions: Evidence from the Russia-Ukraine conflict’, *Working Paper* .

Fisman, R. and Wang, Y. (2017), ‘The Distortionary Effects of Incentives in Government: Evidence from China’s “Death Ceiling” Program’, *American Economic Journal: Applied Economics* **9**, 202–218.

Fisman, R. and Wei, S.-J. (2004), ‘Tax rates and tax evasion: evidence from “missing imports” in China’, *Journal of Political Economy* **112**(2), 471–496.

Fisman, R. and Wei, S.-J. (2009), ‘The smuggling of art, and the art of smuggling: Uncovering the illicit trade in cultural property and antiques’, *American Economic Journal: Applied Economics* **1**(3), 82–96.

Griffin, J. and Kruger, S. (2023), ‘What is Forensic Finance?’, *Available at SSRN* .

Griffin, J. M. and Maturana, G. (2016), ‘Who facilitated misreporting in securitized loans?’, *The Review of Financial Studies* **29**(2), 384–419.

Griffin, J. M. and Shams, A. (2018), ‘Manipulation in the VIX?’, *Review of Financial Studies* **31**(4), 1377–1417.

Haidar, J. I. (2017), ‘Sanctions and export deflection: evidence from Iran’, *Economic Policy* **32**, 319–355.

Hamilton, J. (1983), ‘Oil and the macroeconomy since World War II’, *Journal of Political Economy* **38**, 228–248.

Heckman, J. J., Ichimura, H. and Todd, P. E. (1997), ‘Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme’, *Review of Economic Studies* **64**(4), 605–654.

Hong, H. and Stein, J. (1999), ‘A unified theory of underreaction, momentum trading, and overreaction in asset markets’, *Journal of Finance* **54**, 2143–2184.

Hong, H., Torous, W. and Valkanov, R. (2007), ‘Do industries lead stock markets?’, *Journal of Financial Economics* **83**, 367–396.

Hong, H. and Yogo, M. (2012), ‘What does futures market interest tell us about the macroeconomy and asset prices?’, *Journal of Financial Economics* **105**, 473–490.

Hsieh, C.-T. and Moretti, E. (2006), ‘Did Iraq cheat the United Nations? Underpricing, bribes, and the oil for food program’, *Quarterly Journal of Economics* **121**, 1211–1248.

Huynh, L. D. T., Hoang, K. and Ongena, S. (2023), ‘The impact of foreign sanctions on firm performance in Russia’, *Swiss Finance Institute Research Paper* (23-115).

Itskhoki, O. and Mukhin, D. (2022), Sanctions and the exchange rate, Technical report, National Bureau of Economic Research.

Jacob, B. A. and Levitt, S. D. (2003), ‘Rotten apples: An investigation of the prevalence and predictors of teacher cheating’, *Quarterly Journal of Economics* **118**(3), 843–877.

Kaplow, L. (1990), ‘Optimal Deterrence, Uninformed Individuals, and Acquiring Information about Whether Acts Are Subject to Sanctions’, *Journal of Law, Economics, and Organization* **6**(1), 93–128.

Kilpatrick, R. (2022), ‘Monitoring sanctions compliance at sea’, *Northwestern Journal of International Law and Business* **42**, 221–251.

Lacy, D. and Niou, E. M. S. (2004), ‘A theory of economic sanctions and issue linkage: The roles of preferences, information, and threats’, *Journal of Politics* **66**(1), 25–42.

Lastauskas, P., Proškutė, A. and Žaldokas, A. (2023), ‘How do firms adjust when trade stops?’, *Journal of Economic Behavior & Organization* **216**, 287–307.

Liberti, J. M., Seru, A. and Vig, V. (2016), ‘‘Information, Credit, and Organization’’.

Mamonov, M. and Pestova, A. (2023), ‘The price of war: Macroeconomic and cross-sectional effects of sanctions on Russia’.

Marion, J. and Muehlegger, E. (2008), ‘Measuring illegal activity and the effects of regulatory innovation: Tax evasion and the dyeing of untaxed diesel’, *Journal of Political Economy* **116**(4), 633–666.

Mulder, N. (2022), ‘The Sanctions Weapon’, *Finance & Development June 2022* .

Pape, R. (1998), ‘Why economic sanctions still don’t work’, *International Security* **23**, 66–77.

Peress, J. (2014), ‘The media and the diffusion of information in financial markets: Evidence from newspaper strikes’, *Journal of Finance* **69**(5), 2007–2043.

Ready, R. C. (2018), 'Oil consumption, economic growth, and oil futures: The impact of long-run oil supply uncertainty on asset prices', *Journal of Monetary Economics* **94**, 1–26.

Rosenbaum, P. R. and Rubin, D. B. (1983), 'The central role of the propensity score in observational studies for causal effects', *Biometrika* **70**(1), 41–55.

Siddiquee, S. H. M. and van Bergeijk, P. (2012), Reconsidering Economic Sanctions Reconsidered: Detailed Analysis of the Peterson Institute Sanction Database, ISS Working paper, Erasmus University.

Weng, J., Li, G. and Zhao, Y. (2022), 'Detection of abnormal ship trajectory based on the complex polygon', *Journal of Navigation* **75**(4), 966–983.

Wolsing, K., Roepert, L., Bauer, J. and Wehrle, K. (2022), 'Anomaly Detection in Maritime AIS Tracks: A Review of Recent Approaches', *Journal of Marine Science and Engineering* **10**, 1–19.

Table 1: Performance of ML detection models

The hyper-parameters of the models are optimally determined by 10-fold cross-validation and reported in Internet Appendix Table A-2. Panel A (Panel B) reports the model performance using only AIS data (both AIS and satellite data). We separately report the detection of spoofing and dark activity, and "total" reports the simple average of these two. The implied false positive rate is calculated from the pseudo-R^2, with $\bar{p} = 2.23\%$ as estimated from our data and a risk level $1-\beta = 5\%$. Panel C illustrates the importance of individual categories of predictors. Δpseudo-R^2 for a category is the difference between the pseudo-R^2 of a full model, as those reported in panel B, and the same model that sets this particular category to zero (note that all our predictors are zero-one indicators). When calculating Δpseudo-R^2, we treat negative pseudo-R^2 or negative Δpseudo-R^2 as zero. We report the simple averages of the respective Δpseudo-R^2's for the tree and neural network models, as well as spoofing and dark detection. The last column of this panel shows the average increase in the implied false positive rate when setting the respective category to zero.

Panel A: Model Performance Using Only AIS Data			
Detect	Model	Pseudo-R^2	Implied false positive rate
Total	Tree	**9.1%**	**52.9%**
	NN	**10.2%**	**49.8%**
Spoofing	Tree	6.7%	60.5%
	NN	7.2%	58.9%
Dark	Tree	11.5%	46.3%
	NN	13.3%	42.2%
Panel B: Model Performance Using AIS and Satellite Data			
Detect	Model	Pseudo-R^2	Implied false positive rate
Total	Tree	**15.2%**	**38.0%**
	NN	**16.3%**	**35.8%**
Spoofing	Tree	11.8%	45.7%
	NN	13.8%	40.9%
Dark	Tree	18.7%	31.6%
	NN	18.8%	31.4%
Panel C: Predictor Importance			
Category		ΔPseudo-R^2	ΔImplied false positive rate
Satellite Detection		6.2%	14.7%
Identity Change		3.7%	8.0%
Risky Flag		2.8%	6.1%
Irregular Trajectory		0.2%	0.3%
Ship-to-ship Transfer		6.2%	14.8%
DBSCAN Outlier		5.3%	12.2%
Dark Activity		9.4%	24.5%

Table 2: Comparison between Windward's disclosure and ML-model detection

The first two lines in this table show the total number of tankers that violated Iranian sanctions during 2021 January, as per our proprietary dataset, and how many among them are reported on Windward's list as high/moderate risk. The bottom panel shows the performance of our decision tree and neural network models. We use either only AIS-based predictors or combine them with predictors based on satellite data. The set of predictors is as described in Section 3. The optimal hyper-parameters of the models, e.g., tree depth or number of layers of a neural network, are determined via 10-fold cross-validation. We show separately the number of sanctions-violating tankers that engage in spoofing and going dark that our models detect; we also report the total number of detections for both types of behavior – the totals can be smaller than the sum of the two types of behavior because a tanker can engage in both. Results are shown for different confidence levels (cl = 99%, 95%, and 90%) – for example, under a 95% confidence level no more than 5% of complying tankers can be mislabelled as violators.

Total Violators:		33			
on Windward's list:		**27**			
			Number of violators detected		
Detect	Model	Predictors	$cl = 99\%$	$cl = 95\%$	$cl = 90\%$
Total	Tree	Satellite + AIS	**5**	**12**	**17**
	NN	Satellite + AIS	**5**	**13**	**15**
	Tree	Only AIS	3	8	14
	NN	Only AIS	0	5	13
Spoofing	Tree	Satellite + AIS	3	5	6
	NN	Satellite + AIS	4	5	6
	Tree	Only AIS	1	2	5
	NN	Only AIS	0	3	4
Dark	Tree	Satellite + AIS	3	9	14
	NN	Satellite + AIS	1	10	12
	Tree	Only AIS	2	6	10
	NN	Only AIS	0	4	11

Table 3: Difference-in-differences analysis of fixtures changes

This table reports the coefficients (i.e., average treatment effects on the treated, ATT) estimated with PSM-DiD and Abadie (2005)'s semiparametric DiD to examine the change in fixtures for high/moderate-risk tankers after Windward's disclosure. The sample period is from Feb 2023 to Mar 2024. The post-period starts from Aug 2023. The propensity score is calculated by regressing the high/moderate-risk tanker indicator on pre-period tanker characteristics using ML methods: average outputs from a decision tree and neural networks. The variables used to construct the propensity score are shown in Table ??. In panel A, we match tankers within each tanker type and calculate weights based on propensity scores and a Gaussian kernel with a bandwidth of 0.01, 0.03, or 0.05. Then, we do DiD estimation in the matched sample, controlling for tanker fixed effects and time×tanker-type fixed effects. The standard errors in panel A are double clustered at the tanker and time×tanker-type levels. In panel B, we implement Abadie's method. We first collapse our sample into two cross-sections by calculating each tanker's pre- and post-period average fixtures and then use the Stata command "absdid", which needs to specify the dependent variable (i.e., fixture changes from pre- to post-period), treatment dummy (i.e., high/moderate-risk tanker indicator), and propensity score. Since Abadie's estimator is derived without directly accounting for macro trends through time fixed effects, we manually subtract the cross-sectional mean to account for time fixed effects or the cross-sectional mean by tanker type to account for time×tanker-type fixed effects. We trim the fixtures at the 1st and 99th percentiles each month to avoid the influence of outliers. The standard errors in panel B are as derived in Abadie (2005). *, **, *** denote significance at the 10%, 5%, and 1% level, respectively.

Panel A: PSM-DiD			
	Dependent variable: Fixtures		
	Bandwidth = 0.01	Bandwidth = 0.03	Bandwidth = 0.05
ATT for High Risk	-16.063***	-15.150**	-14.605**
	[-2.69]	[-2.39]	[-2.31]
Obs. (tanker-month)	5,378	5,574	5,574
ATT for Moderate Risk	-0.684	-1.197	-1.493
	[-0.17]	[-0.32]	[-0.41]
Obs. (tanker-month)	6,081	6,083	6,083
Time × Tanker Type FE	Yes	Yes	Yes
Tanker FE	Yes	Yes	Yes

Panel B: Abadie (2005) Semiparametric DiD			
	Raw fixtures	Fixtures demeaned by time	Fixtures demeaned by time × tanker type
ATT for High Risk	-18.417***	-20.709***	-13.749***
	[-3.56]	[-4.47]	[-3.16]
Obs. (tanker)	1,036	1,036	1,036
ATT for Moderate Risk	-3.736	-5.087	-0.507
	[-0.88]	[-1.40]	[-0.15]
Obs. (tanker)	1,190	1,190	1,190

Table 4: Difference-in-differences analysis of route changes

This table reports the coefficients (i.e., average treatment effects on the treated, ATT) estimated with PSM-DiD and Abadie (2005)'s semiparametric DiD to examine the change of the routes of high/moderate-risk tankers after Windward's disclosure. The dependent variable is a zero-one indicator of passing during each month within 12 nautical miles (i.e., the boundary of the territorial sea) of Iran, Russia, or Venezuela (only before 2024). The sample period is from Feb 2023 to Mar 2024. The post-period starts from Aug 2023. The propensity score is calculated by regressing the high/moderate-risk tanker indicator on pre-period tanker characteristics using ML methods. The variables used to construct the propensity score are shown in Table A1 in the Internet Appendix. In Panel A, we match tankers within each tanker type and calculate weights based on propensity scores and a Gaussian kernel with a bandwidth of 0.01, 0.03, or 0.05. Then, we do DiD estimation in the matched sample, controlling for tanker fixed effects and time×tanker-type fixed effects. The standard errors in panel A are double clustered at the tanker and time×tanker-type levels. In Panel B, we implement Abadie's method in the same way as mentioned in Table 3. The standard errors in panel B are as derived in Abadie (2005). *, **, *** denote significance at the 10%, 5%, and 1% level, respectively.

Panel A: PSM-DiD			
	Dependent variable: Showing AIS signals near sanctioned countries		
	Bandwidth = 0.01	Bandwidth = 0.03	Bandwidth = 0.05
ATT for High Risk	-0.073***	-0.072***	-0.071***
	[-3.69]	[-3.76]	[-3.79]
Obs. (tanker-month)	39,246	39,246	39,246
ATT for Moderate Risk	-0.016	-0.015	-0.015
	[-1.01]	[-1.02]	[-1.05]
Obs. (tanker-month)	43,820	43,820	43,820
Time × Tanker Type FE	Yes	Yes	Yes
Tanker FE	Yes	Yes	Yes
Panel B: Abadie (2005) Semiparametric DiD			
	Probability	Probability demeaned by time	Probability demeaned by time × tanker type
ATT for High Risk	-0.067***	-0.067***	-0.060***
	[-4.01]	[-4.03]	[-3.60]
Obs. (tanker)	2,543	2,543	2,543
ATT for Moderate Risk	-0.012	-0.012	-0.008
	[-0.97]	[-0.96]	[-0.62]
Obs. (tanker)	3,139	3,139	3,139

Table 5: Difference-in-differences analysis of ownership changes

This table reports the coefficients (i.e., average treatment effects on the treated, ATT) estimated with PSM-DiD and Abadie (2005)'s semiparametric DiD to examine the effect of Windward's disclosure on the probability of tanker owner changes (i.e., liquidity). Our proprietary tanker ownership data contains six snapshots: Dec 2022, Mar, Jul, Sep, Dec 2023, and Mar 2024. The dependent variable is a zero-one indicator of owner changes by comparing the owner names across two snapshots. Note that the time spans between consecutive snapshots are different, so we make them comparable by annualizing the owner change variable. The pre-period owner change includes Dec 2022 - Mar 2023 and Mar - Jul 2023. The post-period owner change includes Jul - Sep 2023, Sep - Dec 2023, and Dec 2023 - Mar 2024. The propensity score is calculated by regressing the high/moderate-risk tanker indicator on pre-period tanker characteristics using ML methods. The variables used to construct the propensity score are shown in Table A1 in the Internet Appendix. In panel A, we match tankers within each tanker type and calculate weights based on propensity scores and a Gaussian kernel with a bandwidth of 0.01, 0.03, or 0.05. Then, we do DiD estimation in the matched sample, controlling for tanker fixed effects and time×tanker-type fixed effects. The standard errors in panel A are double clustered at the tanker and time×tanker-type levels. In Panel B, we implement Abadie's method in the same way as mentioned in Table 3. The standard errors in panel B are as derived in Abadie (2005). *, **, *** denote significance at the 10%, 5%, and 1% level, respectively.

Panel A: PSM-DiD			
	Dependent variable: Owner change		
	Bandwidth = 0.01	Bandwidth = 0.03	Bandwidth = 0.05
ATT for High Risk	-0.183***	-0.169***	-0.164***
	[-3.25]	[-2.91]	[-2.87]
Obs. (tanker-month)	11,402	11,402	11,402
ATT for Moderate Risk	-0.168***	-0.159**	-0.152**
	[-3.30]	[-2.79]	[-2.50]
Obs. (tanker-month)	11,776	11,776	11,776
Time × Tanker Type FE	Yes	Yes	Yes
Tanker FE	Yes	Yes	Yes
Panel B: Abadie (2005) Semiparametric DiD			
	Owner change	Owner change demeaned by time	Owner change demeaned by time × tanker type
ATT for High Risk	-0.185***	-0.184***	-0.170***
	[-3.30]	[-3.29]	[-3.04]
Obs. (tanker)	2,346	2,346	2,346
ATT for Moderate Risk	-0.168***	-0.167***	-0.146***
	[-3.20]	[-3.18]	[-2.77]
Obs. (tanker)	2,421	2,421	2,421

Table 6: Difference-in-differences analysis of changes of charterers

In this table, we examine whether U.S. or U.S.-allied companies tend to stop chartering high/moderate-risk tankers after Windward's disclosure. U.S.-allied companies are those from the UK, European Union, Australia, New Zealand, or Japan. The dependent variable is the number of tankers of a given risk type that each charterer employs in a month. The sample period is from Feb 2023 to Mar 2024. The post-period starts from Aug 2023. We regress the dependent variable on the indicators of U.S. and U.S.-allied charterers interacted with the post-period indicator controlling for charterer fixed effects and time fixed effects. The coefficients of the DiD estimation are reported. The standard errors are double clustered at the charterer and time level. *, **, *** denote significance at the 10%, 5%, and 1% level, respectively.

	Dependent Variable: No. tankers for each charterer-month		
	No. high-risk tankers	No. mod.-risk tankers	No. low-risk tankers
U.S. charterer $\times \mathbb{I}_{\{t\geq 0\}}$	-0.358*	-0.131	0.065
	[-1.91]	[-0.67]	[0.09]
U.S.-allied charterer $\times \mathbb{I}_{\{t\geq 0\}}$	0.013	-0.030	0.136
	[0.08]	[-0.12]	[0.31]
Obs. (charterer-month)	1,820	1,820	1,820
Time FE	Yes	Yes	Yes
Charterer FE	Yes	Yes	Yes

Table 7: Model parameters, moments matching, and implied quantities

This table presents the directly-calibrated model parameters in Panel A, the moments that we match in Panel B, the parameters derived from moments matching in Panel C, the imputed quantities in Panel D, and the model-implied effects due to price changes (i.e., we analyze the effects by fixing Qs in the models to pre-Windward levels) in Panel E. These results are discussed in Sections 5.2 and 5.3.

Panel A: Parameters directly calibrated

Variable	Calibration method	Symbol	Value
Discount factor	time-discounting (0.95) × depreciation	β	0.9
Proportion of Bad tankers	estimation in CRS (2024)	λ	0.213
Mean oil trade revenue (Clean exporters)	15.5 mln, normalized to one	r	1
Mean oil trade revenue (Rogue exporters)	4 dollars per barrel discount	r^R	0.794
Sanction penalty (Clean exporters)	4 mln, comparable to shipping costs	z	0.258
Sanction penalty (Rogue exporters)	50 mln, based on the median from actual cases of seized oil	z^R	3.226
Classification precision (pre-period)	implied by a pseudo-R^2 of 7.5% (public data's predictability in Figure A-4 in Internet Appendix)	ρ_{pre}	0.296
Classification precision (post-period)	implied by a pseudo-R^2 of 17.2% (Windward's predictability in Figure A-4 in Internet Appendix)	ρ_{post}	0.453

Panel B: Matched moments

Moment	Model analog	Data	Model
Proportion of sanctioned tankers (average of pre- and post-Windward periods)	$\frac{Q^{pre}_{Sanc}+Q^{post}_{Sanc}}{2}$	0.5%	0.5%
Annualized volatility of low-risk tankers' fixtures before Windward's disclosure (divided by mean)	$\frac{\text{std}[p^{pre}_L]}{\mathbb{E}[p^{pre}_L]}$	0.77	0.77
Annualized volatility of high-risk tankers' fixtures before Windward's disclosure (divided by mean)	$\frac{\text{std}[p^{pre}_H]}{\mathbb{E}[p^{pre}_H]}$	0.81	0.83
Annualized volatility of top US charterers' revenue before Windward's disclosure (divided by mean)	$\frac{\text{std}[(Q^{pre}_{GL}+Q^{pre}_{GH}+Q^{pre}_{BL}+Q^{pre}_{BH})\tilde{r}]}{\mathbb{E}[(Q^{pre}_{GL}+Q^{pre}_{GH}+Q^{pre}_{BL}+Q^{pre}_{BH})\tilde{r}]}$	0.30	0.21
Disclosure effect on high-risk tankers' fixtures (DiD estimation in Table 3)	$\frac{\mathbb{E}[p^{post}-p^{pre}\mid\text{Post H}]}{\mathbb{E}[p^{pre}\mid\text{Post H}]}$	-0.13	-0.11
Disclosure effect on high-risk tankers' routes (DiD estimation in Table 4)	$\frac{\mathbb{E}[b^{post}-b^{pre}\mid\text{Post H}]}{\mathbb{E}[b^{pre}\mid\text{Post H}]}$	-0.17	-0.20

Panel C: Parameters derived from moments matching

Variable	Symbol	Value
Screening intensity (previous violation)	w_0	0.76%
Screening intensity (current violation)	δ_0	1.60%
Volatility of oil trade revenue	σ	0.365
Max operation cost when dealing with Rogue exporters	$\bar{c}$	0.872

Panel D: Imputed quantities (in mln)

Period	$\mathbb{E}(p_H)$	$\mathbb{E}(p_L)$	$\mathbb{E}(p^R_B)$	$\mathbb{E}(p^R_G)$	V_B	V_G
Pre-Windward	4.0	4.3	7.8	9.5	41.0	44.0
Post-Windward	3.7	4.0	7.5	9.9	36.9	41.0

Table 8: Model-implied effects due to price changes

This table reports the aggregate annual effect of the information shock on the four relevant parties – Rogue exporters, Clean exporters, Bad tankers, and Good tankers (in million dollars). We calculate the effects due to price changes by fixing quantities (i.e., the Qs in the model) to pre-Windward levels. For example, we report $\omega_G^{pre}\mathbb{E}(p_G^{R,pre}) + \omega_B^{pre}\mathbb{E}(p_B^{R,pre})$ as the per-tanker-trip cost for Rogue exporters in the pre-Windward period and $\omega_G^{pre}\mathbb{E}(p_G^{R,post}) + \omega_B^{pre}\mathbb{E}(p_B^{R,post})$ as the one in the post-Windward period, where $\omega_G^{pre} = [Q_{GL}^{R,pre} + Q_{GH}^{R,pre}]/[Q_{GL}^{R,pre} + Q_{GH}^{R,pre} + Q_{BL}^{R,pre} + Q_{BH}^{R,pre}]$ and $\omega_B^{pre} = 1 - \omega_G^{pre}$. We estimate the aggregate annual effects by multiplying per-tanker-trip costs/benefits by the number of tankers and 3.4 trips per year (the average number of trips per tanker from our fixtures data).

	Sanctions violators		Sanctions-compliant agents	
	R exporters (price paid)	Bad tankers (price received)	C exporters (price paid)	Good tankers (price received)
Pre-Windward (per tanker-trip)	9.0	4.9	4.2	5.0
Post-Windward (per tanker-trip)	9.2	4.7	4.0	4.9
Percentage change	2.0%	-5.7%	-6.3%	-3.3%
Aggregate effect (annual, mln) (-/+ indicates loss/gain)	-780	-1,536	5,662	-3,346
Total aggregate effect (annual, mln) (-/+ indicates loss/gain)	-2,316		2,316	

Table 9: Testing a counter-intuitive model implication: Competition effects

This table reports empirical tests on the model's implication that Windward's disclosure reduces low-risk tankers' freight rates (competition effects) – we exploit charterer-level variation in a DiD framework. Since tankers carrying unrefined and unclassified products are more likely to be classified as risky by Windward, the charterers with a larger proportion of such tankers are expected to face relatively higher supply of low-risk tankers after Windwards' disclosure, leading to higher competition among these tankers and lower freight rates received by them. The dependent variable is low-risk tankers' fixtures. The explanatory variables are the high-competition indicator and its interaction with the post-period indicator. The high-competition indicator $\mathbb{I}_{\{\text{High-competition}\}}$ equals one if the charterer-level shift-share competition measure, as described in Section 5.4, exceeds certain cross-sectional cutoff (top 50, 40, or 30%), and zero otherwise. The standard errors are double clustered at the tanker and time×tanker-type levels. *, **, *** denote significance at the 10%, 5%, and 1% level, respectively.

	Dependent variable: Low-risk tankers' fixtures		
	Cutoff = top 50%	Cutoff = top 40%	Cutoff = top 30%
$\mathbb{I}_{\{\text{High-competition}\}} \times \mathbb{I}_{\{t \geq 0\}}$	-12.893**	-14.304**	-15.866**
	[-2.15]	[-2.48]	[-2.62]
Obs. (tanker-month)	4,999	4,999	4,999
Time × Tanker Type FE	Yes	Yes	Yes
Tanker FE	Yes	Yes	Yes

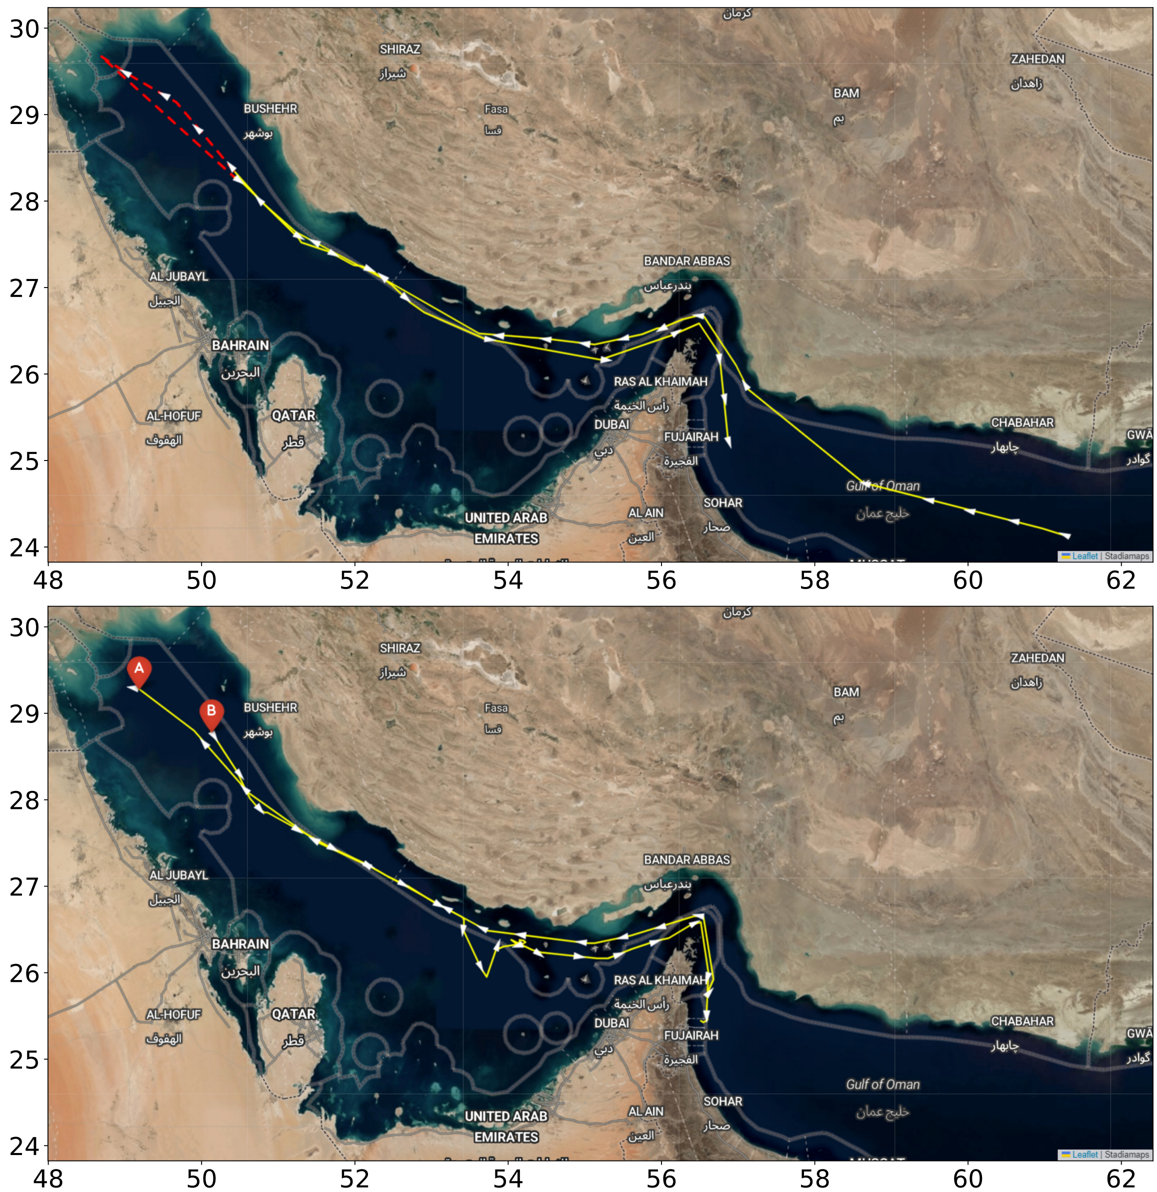

Figure 1: In this figure, solid and dashed lines show the trajectories of two tankers in the Persian Gulf, as given by their AIS signals (AIS is the Automatic Identification System for ship-tracking). White arrows indicate the direction of movement. The top plot shows a case of spoofing, i.e., at a time when the tanker was physically seen at an Iranian port (as per our dataset of sanctions-violating tankers), its AIS signals were showing it traveling in the northern part of the Gulf near Iraq (red dashed lines). Therefore, the tanker was manipulating its signals. The yellow solid lines show this tanker's path before and after its Iranian port visit. The bottom plot shows a case of dark activity, i.e., the tanker stopped emitting signals in a period when it was observed at an Iranian port (again as per our dataset). The last signal before it went dark is indicated by marker "A" and was emitted at 02:16:12 on 2021-01-23. The first signal after emerging from its dark period is indicated by marker "B", not too far from A and near the Iranian coast, and was emitted at 04:54:45 on 2021-01-27, ending a four-day dark period.

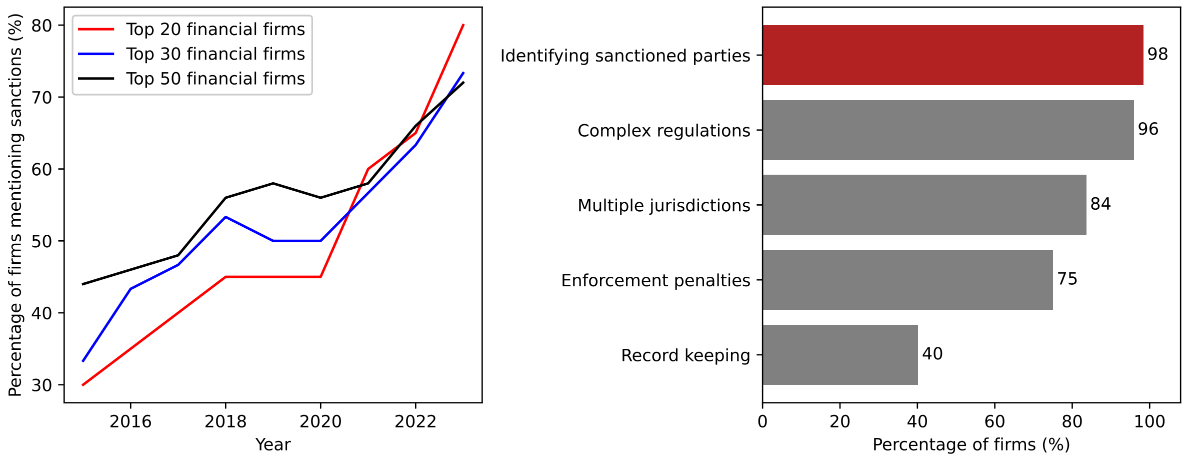

Figure 2: This figure shows the results from a large-language-model-based analysis of 10-K reports of the largest (by total assets) U.S. financial firms between 2015 and 2023. We first ask GPT-4 to identify the reports that mention sanctions enforced by OFAC, or sanctions-related risk, cost, or uncertainty. Then, we ask it to identify the reasons why the firms think sanctions compliance is challenging (further details are in Internet Appendix A). The left plot shows the percentage of firms among the 20, 30, and 50 largest financial firms, respectively, that mention economic sanctions-related risk (or cost or uncertainty) in their 10-K reports over the 2015-2023 period. The right plot shows the top five challenges in sanctions compliance for the 50 largest financial firms.

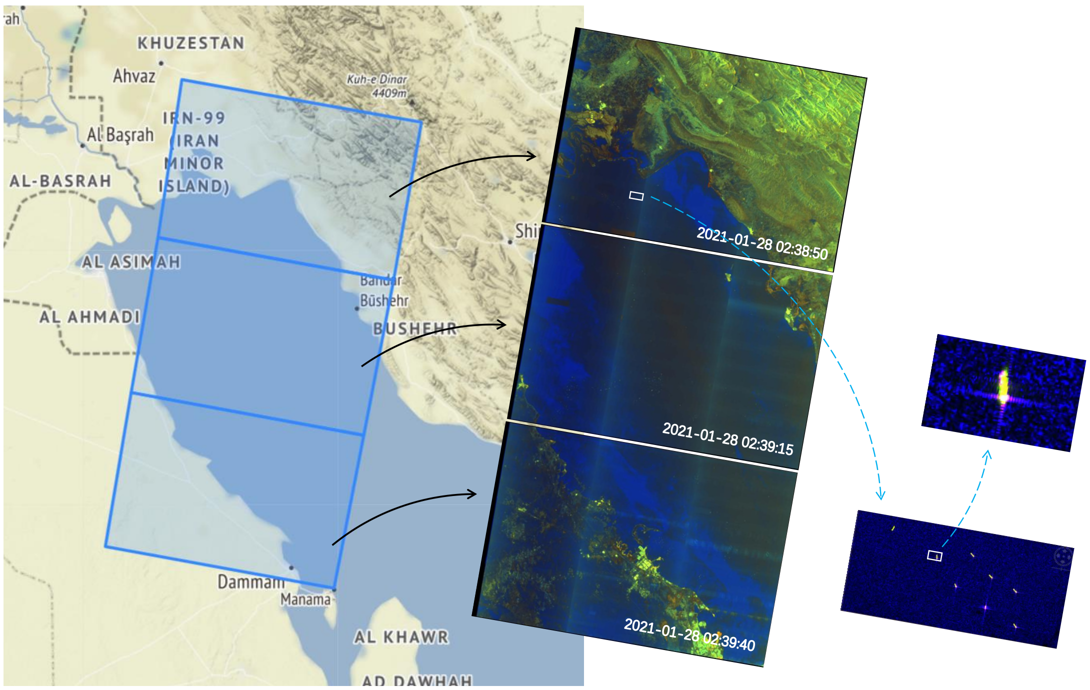

Figure 3: This figure shows examples of satellite images. These three images were taken a little after 2:30 a.m. on the 28th of January 2021 in the Persian Gulf area. Each image covers a rectangular sector, 250 kilometers wide. One insert on the right shows eight ships as little yellow objects at sea, while the other insert shows at the largest enlargement level one of these ships, with individual pixels clearly seen.

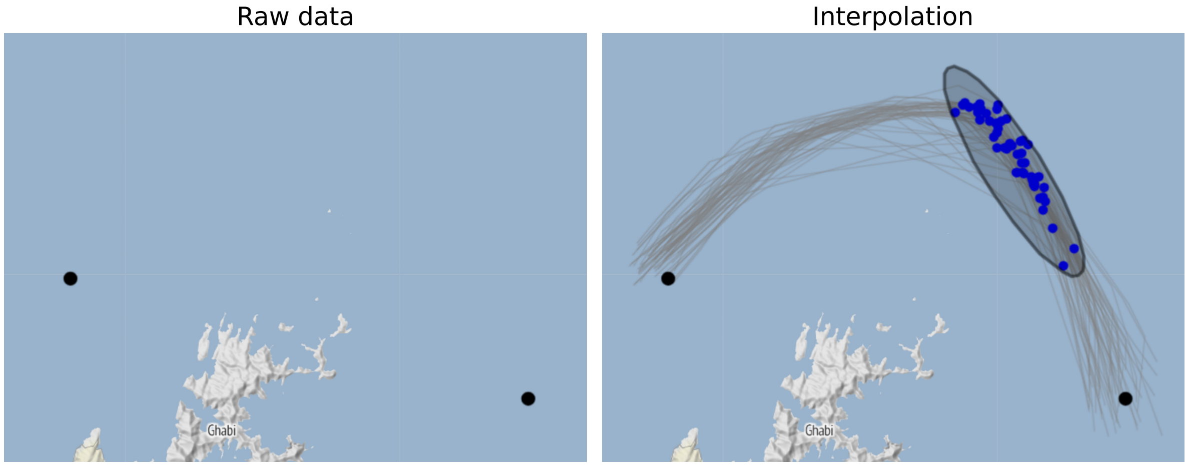

Figure 4: This figure illustrates the construction of the search area via interpolation. The black dots show the locations of a given tanker (call it A) at time t_0 (t_1) of the closest AIS signal before (after) the time t^I when a satellite could see the plotted area (the Strait of Hormuz). Let x_{t_0} and x_{t_1} denote tanker A's locations as per the AIS signals at t_0 and t_1. We require that t_0 and t_1 are within six hours of t^I. Tanker A was moving between these two times, and the distance between x_{t_0} and x_{t_1} is large. We determine the search area using the trajectories of other comparable tankers that have passed near both x_{t_0} and x_{t_1}, in the same direction and over a similar time interval, and hence at similar speed. We require a comparable tanker to have sent a signal within five kilometers from x_{t_0}, and then another one within the same distance from x_{t_1}, whereby the time it took to travel between these two signals is within 30 minutes of the difference $t_1 - t_0$. To increase the number of matches, we search for comparable tankers in the three-month period between December 2020 and February 2021. If multiple pairs of signals from a given tanker satisfy these matching criteria, we chose the pair with starting and ending locations closest to x_{t_0} and x_{t_1}. The grey lines in the figure show the trajectories of comparable tankers. Then we find by interpolation the location of each comparable tanker at its time corresponding to t^I for tanker A. These locations are shown as blue dots in the figure. Finally, we define the search area as the smallest ellipse that contains all of these interpolated locations (i.e., the blue dots), except those which are more than five standard deviations away from the center of the interpolation locations. We use ellipses to reduce the size of the search area.

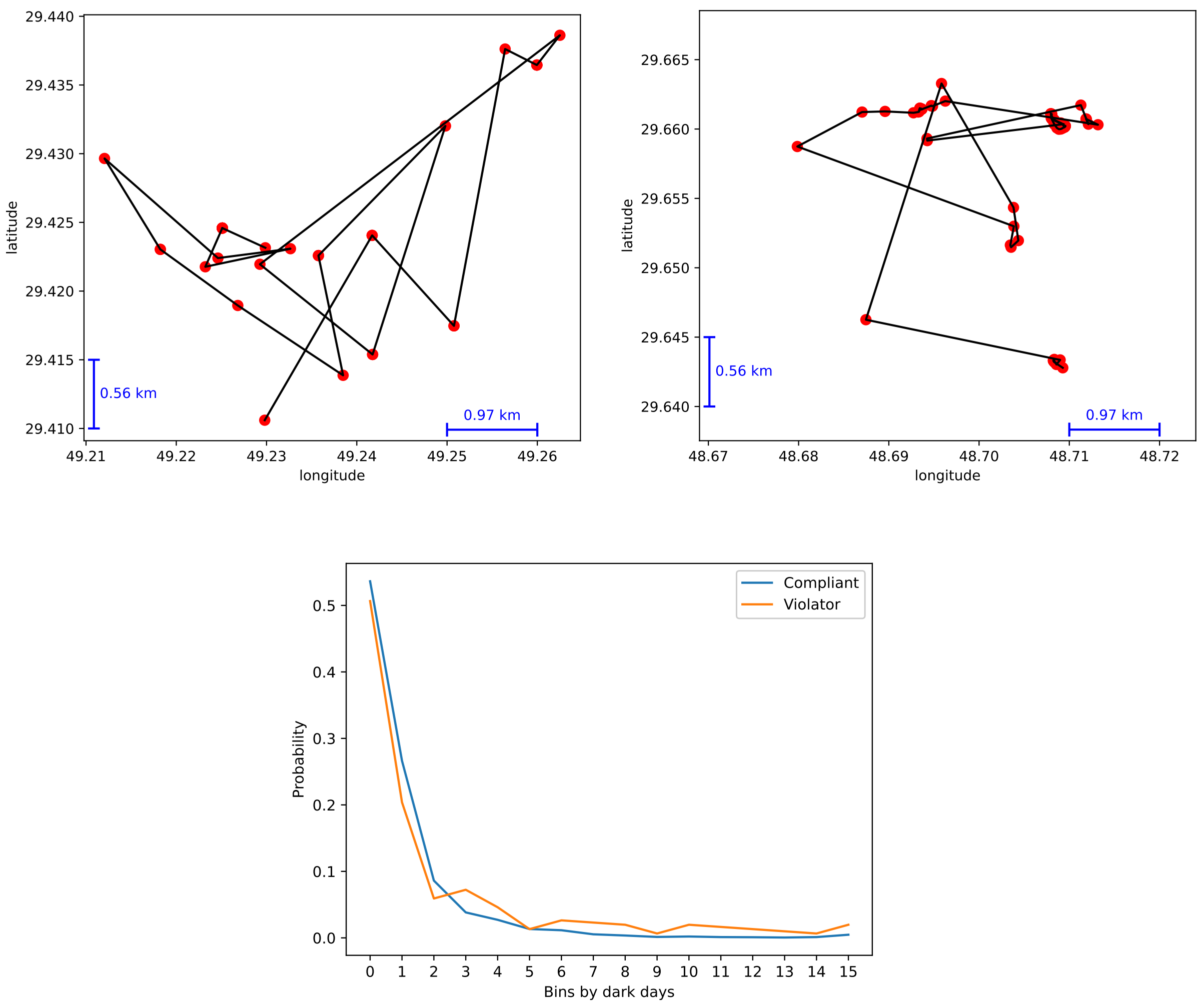

Figure 5: This figure shows irregular trajectories and distributions of dark days from sanctions-violating tankers and sanctions-compliant tankers. The top-left plot is for a sanctions-violating tanker from our proprietary dataset, and the top-right plot is for a sanctions-compliant tanker. The axes of these two plots show degrees of latitude and longitude, and the red dots show the coordinates of each AIS signal emitted by each tanker. The bottom plot shows the distribution of dark days in Jan 2021 for tankers entering the Persian Gulf conditional on the tanker going dark for at least half a day. Each bin in x-axis represents a range of dark time: bin 0 is from 12 hours to 24 hours; bin x ($x = 1, 2, \cdots, 14$) is from x to $x+1$ days; bin 15 is equal or greater than 15 days. In this plot, the blue (yellow) line represents sanctions-compliant (sanctions-violating) tankers.

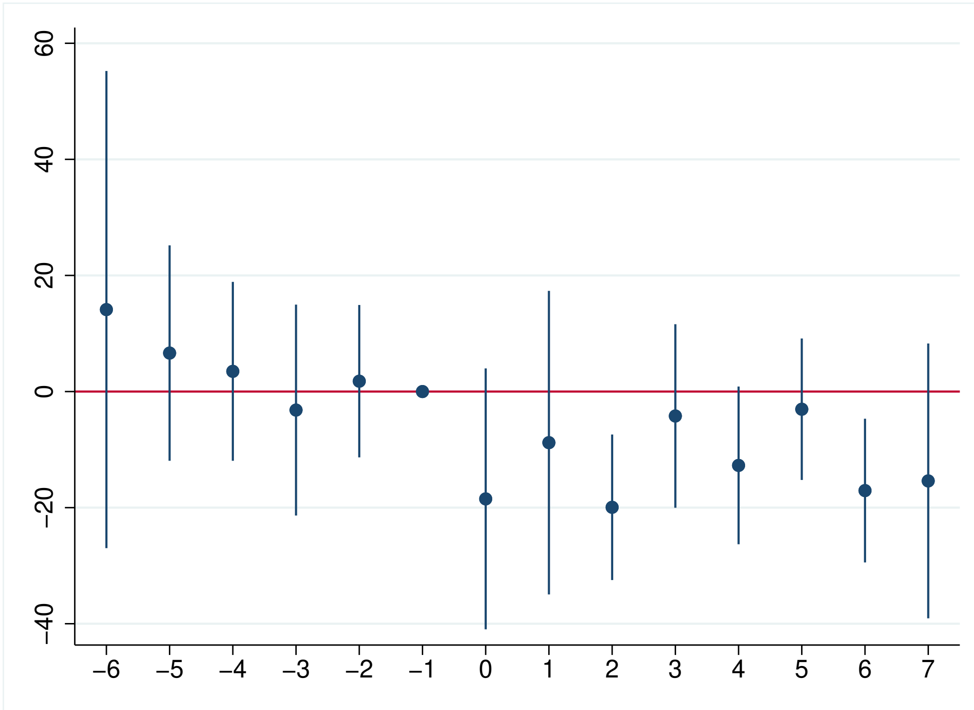

Figure 6: This figure plots the time series of coefficients (i.e., average treatment effects on the treated, ATT) estimated with PSM-DiD to examine the change in fixtures for high-risk tankers only (i.e., we remove moderate-risk tankers from the sample). The sample period is from Feb 2023 to Mar 2024. Period 0 indicates the month of Windward's disclosure, i.e., Aug 2023. The propensity score is calculated by regressing the high/moderate-risk tanker indicator on pre-period tanker characteristics using ML methods: average outputs from a decision tree and neural networks, using variables described in Section 3.3. We match tankers within each tanker type and calculate weights based on propensity scores and a Gaussian kernel with a bandwidth of 0.01. Then, we regress fixtures on indicators of high-risk tankers interacted by period dummies controlling for tanker fixed effects and time×tanker-type fixed effects. We trim the fixtures at the 1st and 99th percentiles each month to avoid the influence of outliers. Vertical bars represent 95% confidence intervals. The standard errors are double clustered at the tanker and time×tanker-type levels.

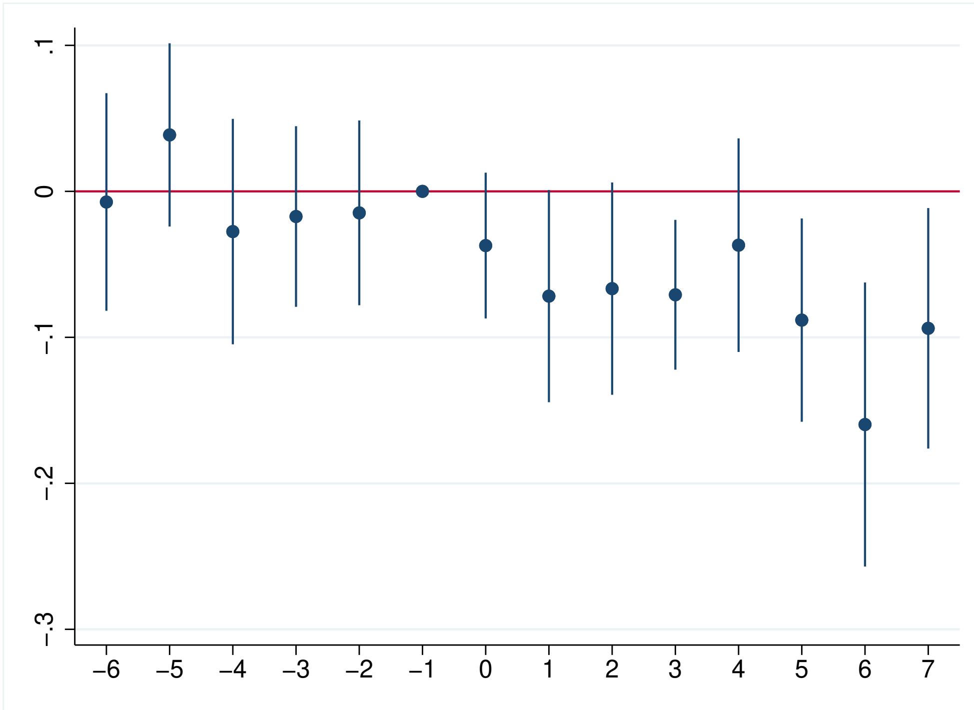

Figure 7: This figure plots the time series of coefficients (i.e., average treatment effects on the treated, ATT) estimated with PSM-DiD to examine the change of the routes of high-risk tankers only (i.e., we remove moderate-risk tankers from the sample). The dependent variable is a zero-one indicator of passing during each month within 12 nautical miles (i.e., the boundary of the territorial sea) of Iran, Russia, or Venezuela (only before 2024). The sample period is from Feb 2023 to Mar 2024. Period 0 indicates the month of Windward's disclosure, i.e., Aug 2023. The propensity score is calculated by regressing the high/moderate-risk tanker indicator on pre-period tanker characteristics using ML methods. The variables used to construct the propensity score are shown in Table A1 in the Internet Appendix. We match tankers within each tanker type and calculate weights based on propensity scores and a Gaussian kernel with a bandwidth of 0.01. Then, we regress the dependent variable on indicators of high-risk tankers interacted by period dummies controlling for tanker fixed effects and time×tanker-type fixed effects. Vertical bars represent 95% confidence intervals. The standard errors are double clustered at the tanker and time×tanker-type levels.

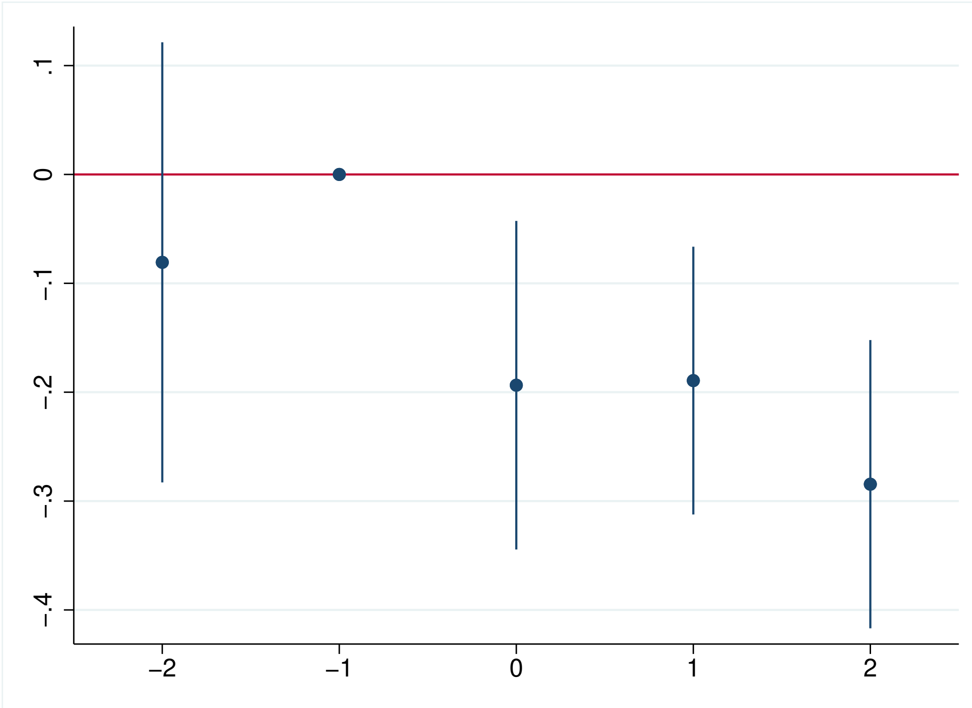

Figure 8: This figure plots the time series of coefficients (i.e., average treatment effects on the treated, ATT) estimated with PSM-DiD to examine the effect of Windward's disclosure on the probability of tanker owner changes for high-risk tankers only (i.e., we remove moderate-risk tankers from the sample). Our proprietary tanker ownership data contains six snapshots: Dec 2022, Mar, Jul, Sep, Dec 2023, and Mar 2024. The dependent variable is a zero-one indicator of owner changes by comparing the owner names across two snapshots. Since the time spans between consecutive snapshots are different, we make them comparable by annualizing the owner change variable. The pre-period includes Dec 2022 - Mar 2023 and Mar - Jul 2023. The post-period includes Jul - Sep 2023, Sep - Dec 2023, and Dec 2023 - Mar 2024. The propensity score is calculated by regressing the high/moderate-risk tanker indicator on pre-period tanker characteristics using ML methods. The variables used to construct the propensity score are shown in Table **??**. We match tankers within each tanker type and calculate weights based on propensity scores and a Gaussian kernel with a bandwidth of 0.01. Then, we regress the dependent variable on indicators of high-risk tankers interacted by period dummies controlling for tanker fixed effects and time×tanker-type fixed effects. Vertical bars represent 95% confidence intervals. The standard errors are double clustered at the tanker and time×tanker-type levels.

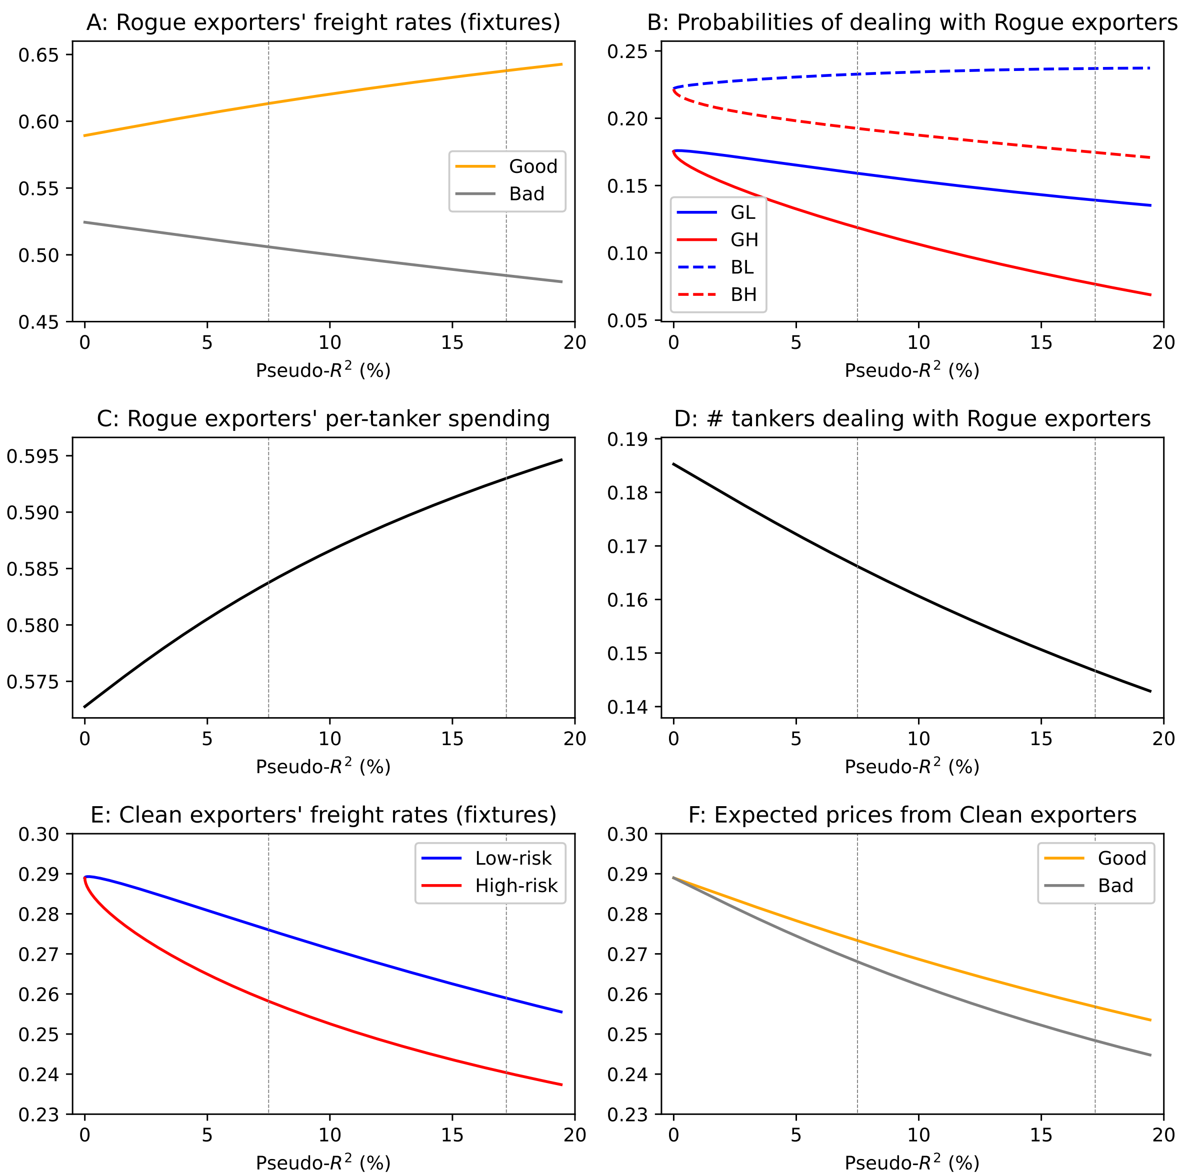

Figure 9: This figure presents the predictions from the calibrated model from Section 5. We solve the model by discretizing the aggregate shock $\tilde{\epsilon}$ on a grid. We plot the average model outcomes for the values of $\tilde{\epsilon}$ as functions of pseudo-R^2, which is backed out from the classification precision ρ in the model. In each subplot, two vertical dashed lines indicate the pre- and post-Windward precisions, which are 7.5% and 17.2%, respectively.