Is a Lack of Information Limiting Sanctions Enforcement?*

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

August 25, 2025

Abstract

Sanctions violations remain widespread, even after authorities started requiring third parties to sever all dealings with violators — a measure that was expected to strengthen enforcement substantially. We show that a key reason for continued violations is the difficulty third parties face in detecting violators. Focusing on oil tankers, we find that even advanced machine-learning applied to ship-tracking and satellite data achieves only modest detection accuracy with 7.5% pseudo-R^2. When a maritime AI firm disclosed a suspect-tanker list that raised accuracy to 21.3%, affected vessels' earnings fell by 13%, they started avoiding sanctioned countries, and their resale probability declined by 37%. A dynamic structural model that endogenizes pricing and sanctions compliance under imperfect detection rationalizes these responses, and helps quantify the information elasticity of enforcement. Model calibration implies that the disclosure redirected $1.4 billion annually from violators; but surprisingly, it also reduced earnings for compliant tankers. While improving information on violators leads to significant enforcement gains, counterfactuals reveal that simply increasing penalties — though frequently emphasized by politicians and the media — has little effect when detection is challenging.

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

JEL codes: F51, D82

---

[*] 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, Pulak Ghosh, John Griffin, Minkang Gu, Harrison Hong, Yunzhi Hu, 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 and students at CKGSB, FMA Asia, London Business School, Monash University, Hong Kong University Summer Finance Conference, HKUST, SAIF, EduHK, 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).

In recent decades, sanctions enforcement has evolved from relying exclusively on direct state action and moved towards a more decentralized, market-based approach. In this framework, regulatory authorities require third-party market participants – like banks, insurers, and shipping companies – not only to avoid sanctioned entities, but also to detect and sever ties with those who facilitate sanction violations (see, e.g., Hufbauer and Jung (2020); Norrlöf (2021); Van Genugten (2021)). Acting in tandem with traditional state-imposed penalties, this approach seeks to mobilize these market participants to deny violators access to financial and logistical infrastructure essential to their operations. Yet despite this expanded enforcement architecture, sanctions violations remain widespread and persistent (see, e.g., Fisman et al. (2024); or media accounts, e.g., The New Yorker).

We show that a key reason for such persistence in sanctions violations is the difficulty third-party entities face in detecting violators. This stands in stark contrast to much of the existing literature, which assumes that violators are either publicly designated by the state or otherwise easily detectable. Our findings highlight a need to modify such assumptions in today's world, where enforcement depends on market participants operating under imperfect information.

Our analysis focuses on oil tankers, which are central to the enforcement of sanctions. This is not only because sanctioned states such as Russia, Iran, and Venezuela rely heavily on oil exports for revenue (Brown, 2020), but also because efficient and uninterrupted transport of oil is essential to the functioning of the global economy (Brancaccio, Kalouptsidi and Papageorgiou (2020, 2023), Ready, 2018). Moreover, the oil shipping market also vividly illustrates the challenges third parties face in detecting sanctions violators. This is because U.S. sanctioning authorities added approximately 100 tankers to publicly available sanctioned entity lists between 2021 and 2024; yet the US Congressional Research Service estimates that over 1,600 vessels transported sanctioned oil during the same period. This leaves third parties to detect nearly 1,500 potential violators on their own — a daunting task given the complexity of global shipping operations and the rampant use of modern technology to hide and obfuscate violations. As Daniel Tadros, Chief Operating Officer of the American Club (a major U.S. shipping insurer) noted: "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, or who its owners are."

We start by quantifying the difficulties third parties face in detecting sanctions violators using publicly available data. If detection proves to be very difficult, sanctions may fail regardless of the third parties' intent to comply. However, even if improving detection is necessary, it may not be sufficient for strengthening sanctions enforcement. Sufficiency requires that such an improvement actually hurts violators. This may not happen if third parties do not intend to comply and continue dealing with violators even if they can detect them, e.g., if they believe they can escape punishment. This raises our second question: If violators could somehow be detected more easily, would the market respond by penalizing them?

To address the first question, we utilize a unique dataset provided by an anonymous port agent from a Middle Eastern country listing all oil tankers that violated Iranian sanctions in January 2021. This list contains 33 foreign-flagged (non-Iranian) tankers operating in the Persian Gulf. While we recognize that this dataset is limited in size and scope, it is to our knowledge the only one of its kind that can serve as ground truth on actual sanctions violations, which is a key missing component in the literature (e.g., Wolsing et al. (2022)).[1]

Using this dataset, we assess the predictive accuracy of various machine learning (ML) models designed to detect sanctions violators. These models incorporate factors highlighted by industry experts, such as tankers temporarily disappearing from tracking systems or exhibiting unusual trajectories indicative of route falsification. Our analysis reveals that even a sophisticated third party using these methods would correctly detect sanctions violators with only 7.5% pseudo-R^2. This detection rate is low, suggesting a high burden of compliance – e.g., it means that firms must avoid one-third of their potential counterparties to keep their risk of engaging with violators below 5% (see Trebbi and Zhang (2022) for a broader discussion of compliance costs and their importance). Our findings here challenge the common assumption in the literature that technological advances, like satellites and ship-tracking, make it straightforward to monitor sanctions violators based on publicly available data – even when violators are large, slow-moving oil tankers.

We next examine our second core question: does relaxing the information constraint — i.e., improving the ability of third parties to detect violators — induce the market to penalize implicated vessels? The ideal empirical strategy would involve the exogenous release of a verifiable list of sanctions violators, followed by an evaluation of market outcomes for the vessels on the list relative to a set of comparable, unlisted tankers. While such ground-truth data are rarely available in real time, we approximate this setting using Windward.AI's (LON:WNWD) public disclosure of a global list of suspect tankers on August 16, 2023, via the London Stock Exchange Group's Refinitiv Eikon platform (LSEG: DN207488 (2023)). Although Windward and a few other firms had previously sold proprietary suspect lists to select clients, these datasets were prohibitively expensive for all but a few major players (see, e.g., IHS-ACSS, 2022). Following the disclosure, Windward's suspect tanker module became available for free (for about 6 weeks, following which a nominal fee of €280 per month

---

[1] Given its proprietary nature, we will make this dataset public to aid replicability. Moreover, we use other benchmarks, such as the ability to predict sanctions out-of-sample, etc. to examine robustness.

was set as its subscription price), dramatically lowering the cost of accessing sanctions-risk information. Although the list comprises suspected violators – and not all of the listed tankers may have actually violated sanctions – we document that its release did materially improve the ability of third-party market participants to detect actual violators.

As a simple first comparison, we find that 27 out of 33 tankers on our Iranian sanctions-violator "ground truth" list were classified as high or moderate risk by Windward. In contrast, our machine learning models could correctly detect at most 13 of these tankers at a 95% confidence level. This comparison suggests that Windward's list is significantly more accurate than prior alternatives.[2] One potential reason behind Windward's accuracy is its ability to utilize certain data accessible only to government agencies, which buy its services for other purposes (e.g., their clients include the European Border and Coast Guard Agency (Frontex), and U.S. agencies such as the Drug Enforcement Administration, see Section 3.1). Further, going beyond the Iranian violator list, we also find that Windward's disclosure improved third parties' ability to detect tankers that would be put on sanctions lists in the near future – a key practical goal for sanctions risk management. In particular, combining Windward's list with a machine learning model like the one we used before reduces the false positive rate from 51% to 20%. Overall, our evidence suggests that Windward's disclosure likely represented a sizable positive shock to third-party market participants' ability to detect sanctions-violating tankers.

To assess the effects of this shock, we use a Propensity Score Matched Difference-in-differences (PSM-DiD) estimator, as well as Abadie (2005)'s semiparametric DiD approach. We compare outcomes for vessels classified as high or moderate risk by Windward to a matched set of unflagged vessels before and after the disclosure, and incorporate time×tanker-type fixed effects to absorb tanker-type level variation in outcomes. In particular, we focus on three key variables. First, we find that earnings for high-risk tankers dropped by 13% after the disclosure, relative to counterfactuals. Second, high-risk tankers were 17% less likely than before, compared to other similar tankers, to navigate within 12 nautical miles of sanctioned oil-producing countries like Russia, Iran, and Venezuela. Taken together, these results suggest large losses for flagged vessels, with their fear of heightened scrutiny outweighing profit opportunities from continuing to serve sanctioned regimes. Third, proprietary data from Drewry indicate that the chance of selling a tanker flagged as risky by Windward dropped by 37% within a year after disclosure, likely due to sanctions regimes prohibiting ownership changes for violating tankers.

Given the large enforcement effects of Windward's disclosure, two features of the in-

---

[2] Even a different publicly available list – from the advocacy group United Against Nuclear Iran (UANI) – included only 10 of these tankers as of August 2023.

formation environment are puzzling. First, if better information strengthens enforcement so strongly, then why do authorities not provide more help with detecting likely violators? Second, given that Windward was valued at just $42 million before disclosure, why was the market willing to pay so little for its data, relative to the much larger ex-post savings its release generated?

To shed light on these issues – and, more generally, on the role of information in the economics of sanctions violations – we develop a dynamic structural model of the oil shipping sector, calibrated to market data and our Windward disclosure event-study findings above. To the best of our knowledge, this is the first dynamic structural model studying market-based sanctions enforcement. The model reconciles Windward's ex-ante low valuation with the seemingly substantial changes its disclosure produces, under the assumption that market participants ex-ante beliefs about Windward's accuracy were based largely on learning from its equity market valuation (e.g., Banerjee, 2011; Bond et al., 2012), in the absence of tight priors and little new data to learn from. In this setting, once Windward's valuation drops due to other reasons such as noise trading (De Long et al., 1990), beliefs about it's accuracy can be updated negatively and turn overly pessimistic. This heightened pessimism, in turn, rationalizes low willingness to pay for Windward's data in the future, leading to to continued low valuation. Such an environment could also explain why Windward chose to disclose their suspect list almost for free: they did so to escape the low-valuation trap by demonstrating their accuracy. If such beliefs about higher accuracy spilled over to their other product lines, increase in sales there could offset any loss of revenues from disclosing the suspect tanker list.

Moreover, our model also provides support for a reason often mentioned by authorities for placing on private market participants the burden of detecting violators – doing so incentivizes these participants to leverage their private information, which authorities typically do not have regular access to. For instance, if a company is asked to send high-tech equipment to a location that's actually a pizza shop, it could suspect foul play (see Huneke (2023) for further examples). The model shows that this is true – there is less investment by exporters to generate private information on suspicious tanker behavior after Windward's disclosure, reducing the aggregate amount of information generated on violators. Even so, however, enforcement still improves with Windward's disclosure – as we saw from the event study. This is primarily due to the fact that market participants' private information (which we calibrate using expenditure on compliance) turns out to be much noisier and therefore less valuable for enforcement than the authority's information.

Besides reconciling these seemingly surprising facts, the model enables us to quantify the *information elasticity of enforcement*, i.e., the extent to which market outcomes for actual sanctions-violating tankers (as opposed to suspected violators detected by Windward) and exporters of sanctioned oil respond to changes in detection accuracy by capturing general equilibrium spillovers. Our findings suggest that Windward's disclosure led to sanctioned oil exporters having to pay 1.8% higher freight rates and reduce their overall hiring of tankers by 5.9%. Further, the value of tankers that actually violated sanctions declined by 7%, translating to a loss of approximately \$1.9 million per large 20-year-old tanker – a vessel type commonly used for transporting sanctioned oil. Since 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.

Further, this model also yields a testable, counterintuitive prediction: While one might intuitively expect that improved classification should benefit compliant ("Good") tankers at the same time as deterring violators, model calibration suggests that these tankers actually face reduced earnings after Windward's disclosure. This occurs because sanctioned ("Rogue") exporters, fearing higher detection probabilities, particularly avoid using "Bad" tankers (those with a history of violations) after the disclosure. As a result, these vessels are pushed out to seek employment with compliant ("Clean") exporters, lowering earnings for compliant tankers in the legitimate market. Although sanctioned 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 4.4, 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. Overall, although better detection improves enforcement, it does not benefit all compliant entities. This particular reallocation from compliant tankers to Clean exporters, however, suggests that disclosing violator identities can also yield unintended benefits – here, lowering shipping costs for legally traded oil.

Finally, our model also enables counterfactual analysis of policies focused on changing compliance incentives through higher penalties. The analysis reveals that such policies have limited impact on enforcement when detection is imperfect. For instance, raising penalties on Clean exporters for using Bad tankers does little to prevent future breaches. Although Clean exporters respond by avoiding high-risk tankers — thereby reducing demand and hence profits for Bad tankers — this creates spillover benefits for sanctioned exporters, who can now hire these tankers at lower rates.

Our paper contributes to multiple strands of literature. First, we relate closely to the literature on sanctions evasion and forensic detection. Prior work has developed creative approaches for uncovering violations using indirect signals—e.g., price gaps in oil markets

(Hsieh and Moretti, 2006), stock return responses to arms embargoes (Dellavigna and La Ferrara, 2010), and rerouted trade flows during geopolitical blockades (Fisman et al., 2024). Unlike these papers, we study how markets respond to improved detection, using a quasi-experimental disclosure and a structural model to trace effects on pricing, routing, and resale. In doing so, we contribute to a broader literature on forensic economics and illegal activity detection (e.g., Dimmock and Gerken, 2012; Fisman and Wei, 2004, 2009; Griffin and Kruger, 2023; Griffin and Maturana, 2016; Griffin and Shams, 2018), while highlighting information frictions as an understudied structural bottleneck to enforcement.

More broadly, we contribute to the literature on the economics of sanctions and embargoes (e.g., Baldwin, 2020; Cipriani et al., 2023; Crozet et al., 2021; Early and Preble, 2019; Eaton and Engers, 1992, 1999; Efing et al., 2023; Huynh et al., 2023; Itskhoki and Mukhin, 2023). These papers typically focus on strategic interactions, incentive compatibility, or political economy considerations and provide foundational analysis to understand sanctions, under the assumption that violations are either observable or formally designated. In contrast, we demonstrate that enhancing detection alone can significantly improve sanctions enforcement. More specifically, we also relate to papers on the effects of recent sanctions or trade restrictions on Iran (Dizaji and van Bergeijk (2013); Haidar (2017)) and Russia (e.g., Babina 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. Similarly, Bai et al. (2025) track the unauthorized shipping of sanctioned crude oil worldwide, and estimate the role played by such flows in reshaping oil markets and macroeconomic dynamics across countries. We differ from this work by focusing on the limits of detection that lead to such unauthorized flows despite regulatory efforts at curbing them, e.g., through market-based sanctions enforcement.

Finally, we add to a large literature on information disclosure and market discipline (e.g., Goldstein and Yang (2019); Liberti, Seru and Vig (2016); Peress (2014)), which analyzes how public and private signals affect contracting and resource allocation. We show that public classification of suspected violators in enforcement-sensitive markets not only shifts prices, but also alters equilibrium dynamics, with spillovers that our model formalizes and quantifies.

Taken together, our paper provides the first unified empirical and structural analysis of sanctions enforcement under imperfect detection. We move beyond traditional concerns of incentive compatibility and strategic signaling to show that informational frictions alone can significantly shape enforcement outcomes.

The rest of the paper is organized as follows: Section 1 provides a brief background on sanctions violators, Section 2 presents ML models detecting sanctions violators, Section 3 examines Windward's sanctions-risk disclosure, Section 4 presents a dynamic structural model and estimates key quantities of interest within its scope, and Section 5 concludes.

1 Background

1.1 Detecting sanctions violators: Do firms actually care?

Here we examine the extent to which third-party businesses invest in sanctions compliance. This matters because even if detecting violators is difficult, the challenge becomes consequential only if third parties are indeed attempting to detect violators and comply.

On one hand, regulators like the U.S. Office of Foreign Assets Control (OFAC) create strong incentives by imposing steep penalties for non-compliance: BNP Paribas, for example, paid \$8.97 billion in 2014. Further, even transactions with entities that violated sanctions years earlier are at risk now, 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 with OFAC for sanctions violations conducted before they bought the company (OFAC, 2020), and Payoneer Inc. was fined for inadvertently processing transactions by entities in sanctioned areas due to weak screening systems (Descartes, 2023).

On the other hand, compliance can be difficult and costly. The U.S. Treasury has acknowledged in its 2021 Sanctions Review that small businesses often lack the resources for full compliance. Even large firms face high costs—sanctions compliance can consume up to 15% of a bank's annual budget (Lloyd's List, May 2022).

To assess the real-world importance of sanctions compliance for third-party firms, we analyze 10-K filings from the 50 largest U.S. financial institutions. These firms often complain that they are at particular risk for inadvertently dealing with counterparties (e.g., borrowers) who may be violating sanctions.

Using LLM-based text analysis (Figure 1), we find that over 75% of them raised sanctions-related concerns in 2023, and 98% cited difficulty in detecting violators as a major challenge. This highlights that detection is a central concern for even large and resource-rich US companies.

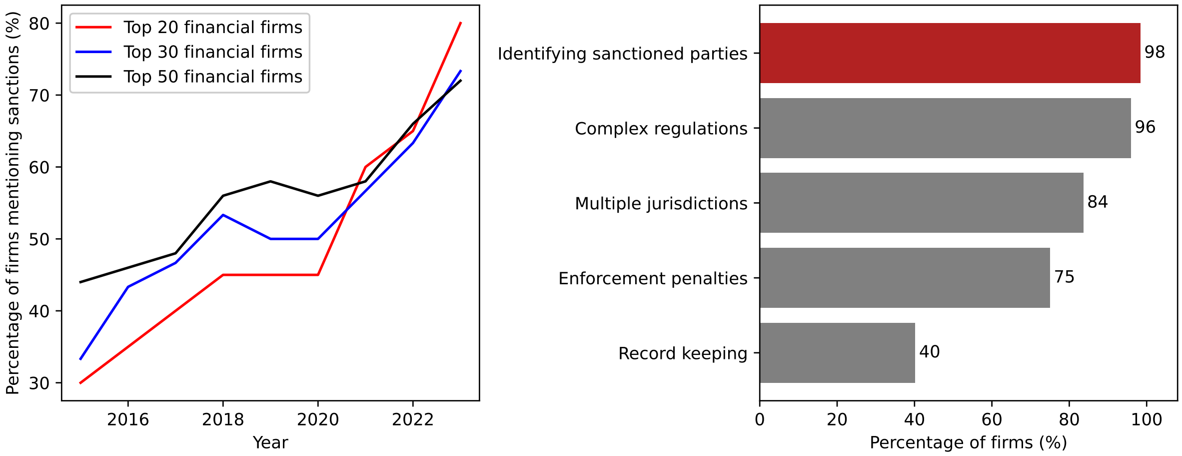

Figure 1: This figure shows the results from a large-language-model-based (GPT 4.0) analysis of 10-K reports of the largest (by total assets) U.S. financial firms between 2015 and 2023. We first ask GPT-4 to identify the reports that mention sanctions enforced by OFAC, or sanctions-related risk, cost, 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.

1.2 Obfuscating Violations: The Case of Oil Transport

Sanctions compliance in the oil shipping sector, which we focus on, is particularly important for market participants because about 20% of oil tankers are estimated to have violated sanctions (CRS, 2024). Here we illustrate typical strategies violators employ to evade detection through ship-tracking data from the Automatic Identification System (AIS).

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 2021 (New York Times, Sep. 2022, Economist, Apr. 2022). Figure 2 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.

2 Detecting sanctions violators

To quantify the challenge of detecting such sanctions violation activity, we take the view of an entity that uses advanced machine learning, combined with radio-signal-based ship-tracking

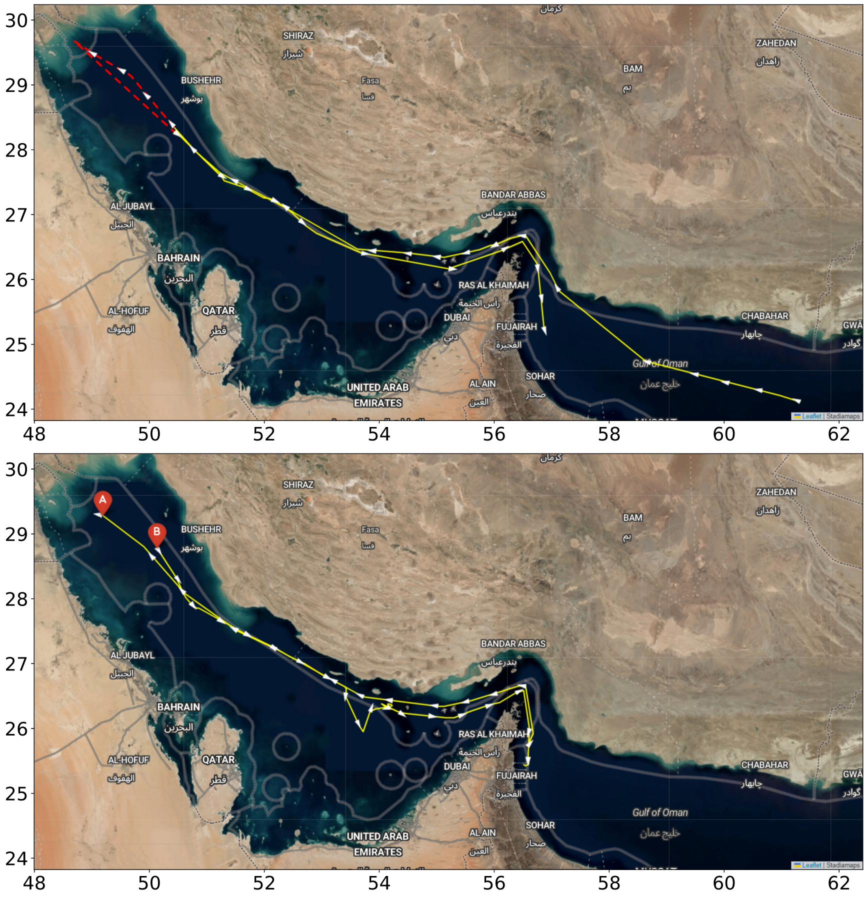

Figure 2: In this figure, solid and dashed lines show the trajectories of two tankers in the Persian Gulf, as given by their AIS signals. White arrows indicate the direction of movement. The top plot shows a case of spoofing, i.e., a tanker which is placed at an Iranian port by our dataset of sanctions-violating tankers, but at the same time emitting falsified AIS signals (red dashed lines) showing it traveling in the northern part of the Gulf near Iraq. 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 while observed at an Iranian port (per our dataset). Marker "A" indicates the last signal before it went dark, emitted at 02:16:12 on 2021-01-23. Marker "B" shows the first signal after its dark period, emitted at 04:54:45 on 2021-01-27 not too far from A and near the Iranian coast, concluding a four-day dark period.

data from AIS, as well as satellite imagery.

2.1 Data and predictors for Machine Learning

We utilize three datasets: a proprietary dataset on tankers in Iran in January 2021, AIS data on high-frequency shipping signals, and Sentinel-1 satellite data.

Testing detection accuracy requires a "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), collected via aggregated text messages from ports, brokers, and charterers. While the provenance and scale of the dataset limit its generalizability, it offers a rare opportunity to test detection models against a set of confirmed breaches. We treat this dataset not as conclusive evidence of violations, but as a validation anchor to benchmark model performance. In addition, we cross-check these vessels against other publicly flagged lists (e.g., Windward's list of high-risk tankers, UANI's "The Ghost Armada" list, etc., which we detail in later sections).[3]

Second, we use AIS data.[4] This data has been utilized in academic research before (e.g., Brancaccio et al. (2020)). AIS is also 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).

Third, we employ data from the Sentinel-1 mission, which has two satellites that provide near-real-time, high-resolution images at six-day intervals of most locations on Earth via synthetic aperture radar. Figure A-1 in the Internet Appendix displays three of these images as an example, with further accompanying details. The utility of such imagery for detection is acknowledged by regulators (OFAC, 2023, Economist, Apr. 2022).

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 using AIS data, to be combined non-linearly in our ML analysis. These refer to (i) identity change, (ii) using risky country flags, (iii) ship-to-ship transfers, (iv) irregular trajectories (that are thought to indicate falsified routes), (v) anomalous locations (where tankers typically do not tread), and (vi) dark activity. We also use two satellite-based predictors: (i) satellite detection of falsified locations, and (ii) satellite-based evidence of dark activity. (Internet Appendix B has detailed description of the predictors construction and summary statistics).

2.2 Predictive Models and their Performance

We use Machine learning methods that allow for various versions and non-linear combinations of the above predictors. To detect spoofing, we create a tanker-day sample where the dependent variable is one if a tanker is observed in Iran (per our ground truth dataset), but its AIS signals indicate another location; it is zero otherwise. To detect dark activity, our dependent variable is one if a tanker is observed in Iran during a dark period exceeding 24

---

[3] 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.

[4] AIS was originally designed for collision avoidance. Vessels over 300 gross tonnage must carry AIS equipment (IMO, 2000). AIS data 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.

hours, with zero otherwise.

To optimally combine predictors, we employ decision trees and neural networks. We use 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 from the best cross-validated model using its pseudo-R^2 and implied false positive rate. The latter represents the proportion of business a firm must forgo to keep the likelihood of engaging with a violator below a set threshold, which we set to 95%. Further details are in Internet Appendix C.

Table 1 presents our results. Panel A reveals that our best AIS-based detection models achieve a pseudo-R^2 of about 10% for detecting spoofing and dark activity, with implied false positive rates reaching 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 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 significant here, with incremental false positive rates of 24.5% and 14.8%, respectively.

Note that Bai et al. (2025) (BFLXZ henceforth) also analyze sanctions-violating oil tankers using AIS data, and although this is not the focus of their paper, they find detection somewhat easier than in our main results in this section. Our lower accuracy arises from comparing them to the list of *actual* violators from our ground truth dataset, while BFLXZ compare their list of predicted violators against other *suspect* classifications (e.g., Lloyd's List and S&P Global Market Intelligence). When we also benchmark our predictions against other suspect lists, we achieve comparable accuracy to BFLXZ, likely because detecting tankers already flagged by others based on similar data and technology is easier than detecting actual violators. Also, while BFLXZ mainly focus on detecting going dark, we identify violations through either going dark or spoofing, with the latter being more challenging.

Overall, the detection of sanctions evaders appears challenging, even when combining numerous predictors mentioned in industry and regulatory sources, and using sophisticated methods/models. It is possible that governments and certain corporations with superior data access (e.g., Starlink) do not face these issues, but most third parties lack access to such data; and yet they are still tasked with such detection to comply with sanctions rules.

Table 1: **Performance of ML detection models**

This table shows detection accuracies for spoofing and dark activity, and their simple average ("Total"), using either AIS data alone (Panel A), or together with satellite data (Panel B). The implied false positive rate is calculated from the pseudo-R^2, as described in Internet Appendix C. In Panel C, Δpseudo-R^2 for a given category is the difference between the pseudo-R^2 of a full model, as those in Panel B, and the same model that sets this particular category to zero (note that all our predictors are zero-one indicators). Negative pseudo-R^2 or Δpseudo-R^2 as treated here as zero. We report simple averages of the respective Δpseudo-R^2's across the tree and neural network models, and spoofing and dark detection. The last column of Panel C shows the average increase in 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%

It is also possible that our results reflect limitations in our ground truth dataset or models, rather than a general difficulty in detection for third party market participants – in Internet Appendix C.6 we provide a discussion and further evidence on such detection being generally challenging, even beyond our data or models.

3 Information shocks and delegated enforcement

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

3.1 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-3 therein for an example of Windward's data on a particular tanker).

Windward leverages over 100 million daily data points, integrating proprietary satellite 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, Mar. 2016, Wired, 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/wealthy market players. However, after Windward's disclosure, this data became significantly more affordable–a Refinitiv Eikon subscription costs £280 per month. In spite of the reduced charge for this dataset, the disclosure increased overall sales for Windward by demonstrating the company's capabilities and helping secure new contracts. Following the announcement, Windward shifted its focus toward offering tailored solutions for both commercial and government clients (Proactive).

Of course, for Windward's suspect tanker list to shape market outcomes, two preconditions must have been satisfied:

(1) The list must have provided new, more accurate information on sanctions violators: We compared Windward's 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 detected 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 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 detecting sanctions violators.

(2) Market participants were aware of and attentive to the disclosure: Figure A-7 in the Internet Appendix provides suggestive evidence; it shows how Windward.AI's homepage views saw a sharp increase following the list's release, signaling heightened market attention.

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 2. 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

3.2 Windward's effect on the difficulty of detecting 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; e.g., if a bank ended up unknowingly lending to such a tanker, the loan amount might be at risk if the tanker got sanctioned.

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 2, 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-predicted probabilities and select the top 3,500 – the same number as Windward. Among 69 tankers sanctioned about-to-be-sanctioned, i.e., eventually sanctioned post-July 2023, our list of the top-3500 only contains 33. On the other hand, Windward's list contains 65, demonstrating its substantially superior accuracy.

Finally, combining both approaches improves prediction, suggesting that public data-based ML still has value even after Windward's disclosure. While public data alone yields a pseudo-R^2 of 7.5% (false positive rate of 51%), and Windward's list achieves 17.2% (26% false positive rate), their optimal mix (0.58 × public data + 0.42 × Windward) leads to a pseudo-R^2 of 21.3% (false positive rate of 20%). Figure A-5 in the Internet Appendix shows how different weights on these two approaches affects overall pseudo-R^2s.

3.3 Real effects of the disclosure on suspect tankers

Here we examine the real effects of Windward's disclosure on suspect ships, focusing on (i) tanker rental rates (fixtures), (ii) tanker route changes, and (iii) ownership transitions.

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 counterfactuals, and we use two methods to do so.

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.

Plot A in Figure 3 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.

Next, in Plot B of Figure 3 we examine the impact of the information shock on tanker routes, i.e., the propensity of newly labeled high/moderate risk tankers to show up in Iranian, Russian, or Venezuelan territorial seas (defined using data from the 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. This plot 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.

Finally, in Plot C of Figure 3 we examine whether tankers newly classified as high/moderate-risk become harder to sell after Windward's disclosure. This hypothesis is based on the fact that many sanctions regimes explicitly prohibit transactions involving violating entities (e.g., see Council Regulation (EU) 833/2014 on such restrictions, Skadden, 2024). Here, we use proprietary tanker ownership data from Drewry, showing owner names and countries in six

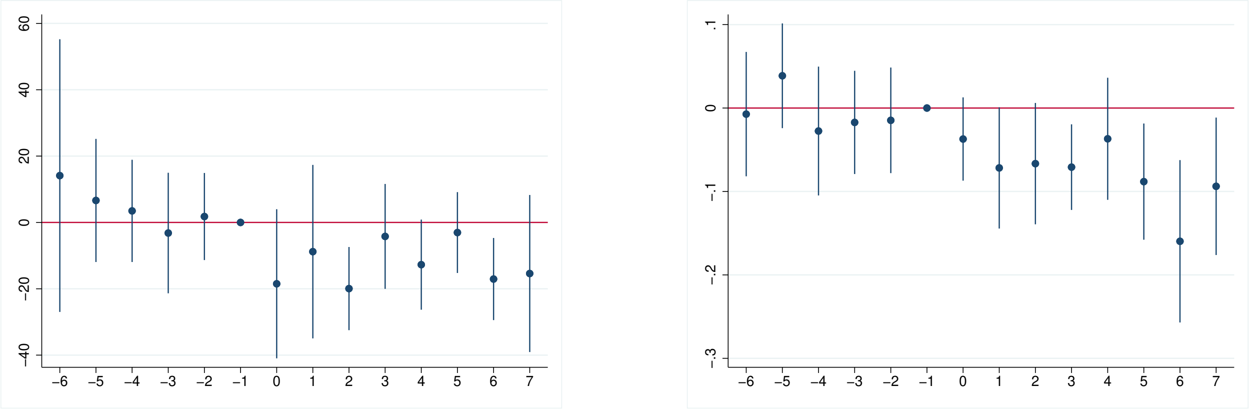

A. Change in fixtures

B. Change in routes

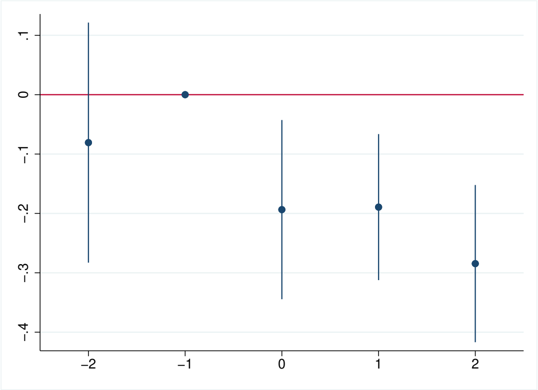

C. Change in ownership

Figure 3: Plot A. shows 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 from Feb 2023 to Mar 2024. 0 indicates Windward's disclosure in Aug 2023. The propensity score is calculated by regressing the high-risk tanker indicator on pre-period tanker characteristics using ML methods: average outputs from a decision tree and neural networks, with variables described in Section 2.1. 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 (trimmed at the 1st and 99th percentiles) on indicators of high-risk tankers interacted with period dummies, controlling for tanker fixed effects and time×tanker-type fixed effects. 95% confidence intervals using standard errors double clustered by tanker and time×tanker-type levels are shown. Plot B. shows analogous coefficients to examine the routes taken by high-risk tankers. The dependent variable equals one in a month when a tanker passes within 12 nautical miles of Iran, Russia, or Venezuela (only before 2024) from Feb 2023 to Mar 2024. 0 indicates Windward's disclosure in Aug 2023. Plot C. shows analogous coeffcients to examine the effect of Windward's disclosure on the probability of tanker owner changes for high-risk tankers. 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.

snapshots between December 2022 and March 2024. The dependent variable 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. Plot C shows that Windward's disclosure has indeed reduced the turnover of high/moderate-risk tankers (which we merge together in this panel due to small sample sizes here), making them less desirable.

While this visual evidence is suggestive, and post-period coefficients in Figure 3 are uniformly lower than the pre-period ones, 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 $moderate\_risk_i$. The first three columns in Panel A of Table 3 presents the PSM-DiD regression results for fixtures.

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.

The first three columns in Panel B of Table 3 shows results from this method for fixtures. 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.

Table 3: **Difference-in-differences analysis**

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 (i) fixtures, (ii) routes, and (iii) ownership for high/moderate-risk tankers after Windward's disclosure. The sample period is from Feb 2023 to Mar 2024, and the post-period starts from Aug 2023. In panel A, we match tankers within each 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, with standard errors double clustered at the tanker and time×tanker-type levels. In panel B, we implement Abadie's method by collapsing our sample into two cross-sections by calculating each tanker's pre- and post-period averages of the dependent variables. Since Abadie's estimator is derived without directly accounting for macro trends, we manually subtract 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, and derive standard errors as in Abadie (2005). *, **, *** denote significance at the 10%, 5%, and 1% level, respectively.

Panel A: PSM-DiD

Dependent variable	I. Fixtures			II. Showing AIS signals near sanctioned countries			III. Owner change		
Bandwidth	0.01	0.03	0.05	0.01	0.03	0.05	0.01	0.03	0.05
ATT for High Risk	-16.06***	-15.15**	-14.61**	-0.073***	-0.072***	-0.071***	-0.183	-0.169	-0.164***
	[-2.69]	[-2.39]	[-2.31]	[-3.69]	[-3.76]	[-3.79]	[-3.25]	[-2.91]	[-2.87]
Obs. (tanker-months)	5,378	5,574	5,574	39,246	39,246	39,246	11,402	11,402	11,402
ATT for Mod. Risk	-0.68	-1.20	-1.49	-0.016	-0.015	-0.015	-0.168*	-0.159	-0.152**
	[-0.17]	[-0.32]	[-0.41]	[-1.01]	[-1.02]	[-1.05]	[-3.30]	[-2.79]	[-2.50]
Obs. (tanker-months)	6,081	6,083	6,083	43,820	43,820	43,820	11,776	11,776	11,776
Time × Type FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Tanker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Abadie (2005) Semiparametric DiD

	raw	by time	by time × type	raw	by time	by time × type	raw	by time	by time × type
ATT for High Risk	-18.42***	-20.71***	-13.75***	-0.067	-0.067	-0.060***	-0.185	-0.184	-0.170***
	[-3.56]	[-4.47]	[-3.16]	[-4.01]	[-4.03]	[-3.60]	[-3.30]	[-3.29]	[-3.04]
Obs. (tankers)	1,036	1,036	1,036	2,543	2,543	2,543	2,346	2,346	2,346
ATT for Mod. Risk	-3.74	-5.09	-0.51	-0.012	-0.012	-0.008	-0.168	-0.167	-0.146***
	[-0.88]	[-1.40]	[-0.15]	[-0.97]	[-0.96]	[-0.62]	[-3.20]	[-3.18]	[-2.77]
Obs. (tankers)	1,190	1,190	1,190	3,139	3,139	3,139	2,421	2,421	2,421

Of note here is that while matching on observables reduces imbalance, we recognize that unobserved differences between risky and low-risk tankers may still bias our estimates. To mitigate such concerns, we assess the robustness of our results. Our evidence suggests that they remain similar, not only across multiple methods and bandwidth choices, but also when we use logistic regression for propensity scores, drop August 2023 (to allow about two weeks for the disclosure effect to come into force), omit Russia, and use bootstrap standard errors. We present these robustness tests in Table A-5 in the Internet Appendix.

The middle columns in Table 3 present the estimation of ATT on route changes, combining post-period months into a single dummy, as we did for fixtures. 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, it drops by 17% (6.8/40) in the post-period. 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.

The last three columns in Table 3 present results on tanker ownership changes in a regression setting similar to the other columns. 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).

In one further exercise (reported in Internet Appendix E.5) we examine whether U.S. and U.S.-allied charterers start avoiding tankers classified as high or moderate risk by Windward. Our evidence shows a significant drop in high-risk tankers usage by U.S. charterers, while effects for U.S.-allied charterers are not statistically significant.

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.

4 Sanctions and Disclosure through the Lens of a Model

In this section, we develop a dynamic structural model of intermediaries and sanctions, and calibrate it using data from the oil shipping market and our Windward event-study results from the preceding section. The model (i) enables us to quantify the impact of

Windward's disclosure on two key actors for whom data is unavailable – actual sanctions-violating tankers (as opposed to suspected violators detected by Windward) and exporters of sanctioned oil. Additionally, (ii) the model suggests one way to reconcile Windward's ex-ante low valuation with the seemingly substantial changes its disclosure produces. This particular modeling choice also provides one explanation for why Windward might have been stuck in a low-valuation trap, which might, in turn, have led them to disclose the list at low cost to publicize their accuracy. Finally, (iii) we use the model to generate a counterfactual where sanction penalties are increased further, and empirically test an important counter-intuitive result which it produces.

4.1 Theoretical Framework

4.1.1 Environment

We model the oil-shipping sector. Our economy consists of oil exporters and tankers.

- 'Rogue' (R) exporters deal in sanctioned products (e.g., Iranian oil companies), while 'Clean' (C) exporters deal in non-sanctioned, legally-traded products (e.g., US oil companies).

- Tankers are either 'Good' (G), i.e., those that have never carried sanctioned oil, or 'Bad' (B), i.e., those that have done so at least once in the past. A Bad tanker cannot revert back to being Good, reflecting the fact that tankers can be sanctioned for past violations. Once a Bad tanker is sanctioned and put on a designated list it cannot be hired by any exporter. However, some Bad tankers are not sanctioned yet, and these are the main source of risk for an exporter in our model, who faces penalties and disruptions if such tanker is sanctioned while in their employ.

- Third parties like banks, insurers etc. who provide intermediary services face similar risks if engaging with Bad tankers. Therefore, we do not explicitly include them in the model for simplicity; one can think of them as all being folded into the Clean exporter category.

- Rogue exporters know which tankers are Bad because they have employed them previously. Clean exporters only know that a proportion λ of all tankers are Bad (λ is common knowledge). They use a noisy detection technology that classifies a portion λ of tankers as High-risk (H), and the rest as Low-risk (L). We assume a stationary environment where the proportion of Bad tankers, λ, is constant.

- Tankers know if they are Bad, and can also infer their risk label (H or L) from the fixture quoted to them by a randomly matched Clean exporter (we assume, for simplicity, that search frictions prevent tankers from shopping for rates from different Clean exporters). Rogue exporters do not know the fixtures quoted by Clean exporters (who compete with them for tankers), and hence cannot infer risk labels.

4.1.2 Timeline

We assume the following sequence of events each period:

1. Clean exporters optimally choose detection technology and assign risk labels to tankers.
2. An aggregate i.i.d. mean-zero shock $\tilde{\epsilon}$ to the oil trading revenue of all exporters realizes, reflecting unpredictable oil price changes that exporters cannot hedge ex-ante.
3. All exporters choose simultaneously freight rates to offer to tankers: 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.
4. Each tanker receives two freight rates: one from a randomly matched Clean exporter and another from Rogue exporters, and decides whom to engage with.
5. The sanctioning authority sanctions (a subset of) violators.
6. Each player receives a payoff and proceeds to the next period (tankers continue to the next period only if not sanctioned in this period).

4.1.3 Clean exporters' detection technology

Signals

There are N Clean exporters, indexed by i. To detect whether a given tanker j is Bad, Clean exporters rely on three potential sources of information (we omit here the tanker index j for brevity).

(i) A common signal $s^{C,com}$, based on public information, available at zero cost and with fixed precision σ_ζ^{-2}:

$$s^{C,com} = \mathbb{I}(\text{Bad Tanker}) + \zeta, \quad \zeta \sim \mathbb{N}(0, \sigma_\zeta^2), \quad \zeta \perp\!\!\!\perp u. \tag{3}$$

(ii) A private signal $s_i^{C,pri}$, with precision $\sigma_{\xi,i}^{-2}$:

$$s_i^{C,pri} = \mathbb{I}(\text{Bad Tanker}) + \xi_i, \quad \xi_i \sim \mathbb{N}(0, \sigma_{\xi,i}^2). \quad \xi_i \perp\!\!\!\perp u, \zeta, \tag{4}$$

Clean exporter i can reduce $\sigma_{\xi,i}$ by paying a (continuous) cost:

$$\Omega(\sigma_{\xi,i}) = \kappa\sigma_{\xi,i}^{-2}, \tag{5}$$

where κ is a constant calibrated by matching moments.[5]

(iii) A signal s^W that can be bought for a fixed price P^W from a specialized data provider like Windward. In our model, such a profit-motivated data intermediary plays a key role, as its data release represents the information shock that we examine.

As Windward integrates information from multiple sources, that include the United Nations, European Border and Coast Guard Agency (Frontex), and various U.S. agencies, among others, we model Windward's signal as a noisy version of a signal s^A, which is only available to the sanctioning authority:

$$s^W = s^A + v, \quad v \sim \mathbb{N}(0, \sigma_v^2), \quad v \perp\!\!\!\perp u, \zeta, \xi. \tag{6}$$

s^A determines the probability of sanctioning the given tanker, as specified below.

The common signal in Eq. (3) reflects due diligence conducted using public data, of the type presented in Section 2.

The private signal in Eq. (4) captures the requirement imposed by the sanctioning authority that market participants do their own due diligence, which is motivated by its belief that these participants have private information on suspicious sanctions-violating behavior. Although we lack details on the exact private information Clean exporters have, industry reports reveal the equilibrium cost they pay for due diligence, enabling us to calibrate relevant parameters.

Windward's signal s^W in Eq. (5) – unlike the common signal $s^{C,com}$ – contains information about the authority's signal which determines sanctions. However, Clean exporters do not know the precision of this signal ex-ante, i.e., they do not know σ_v^2 before Windward's disclosure, but form beliefs about it. In Section 4.2.3, we specify how their beliefs can be calibrated from Windward's observed equity market value. For now, let $\tilde{\sigma}_{v,i}^{-2}$ denote Clean exporter i's posterior belief on Windward's precision.

From signals to risk labels:

Clean exporters optimally choose whether to buy Windward's signal, to what extent they want to improve the precision of their private signal by paying the associated cost, and how

---

[5] This functional form satisfies the no-arbitrage condition in information acquisition: the cost of independently acquiring two signals with (im)precision σ_A and σ_B and then optimally combining them, $\Omega(\sigma_A) + \Omega(\sigma_B)$, is the same as the cost of directly acquiring a signal with the combined (im)precision, $\Omega(1/\sqrt{\sigma_A^{-2} + \sigma_B^{-2}})$.

to combine the signals. For the latter, we assume that they weigh each signal by the inverse of its noise variance.

Let $\chi = 1$ ($\chi = 0$) indicate a Clean exporter's decision to buy (not buy) Windward's information (we study a symmetric equilibrium in which all Clean exporters choose the same action). Before Windward's disclosure, the expected precision of the combined signal is $\sigma_\zeta^{-2}+\sigma_\xi^{-2}+\chi\cdot(\sigma_u^2+\tilde{\sigma}_v^2)^{-1}$, and the cost is $\Omega(\sigma_\xi)+\chi\cdot P^W$. After Windward's disclosure, clean exporters have free access to s^W and can observe its precision. Accordingly, the combined signal's precision becomes $\sigma_\zeta^{-2} + \sigma_\xi^{-2} + \chi \cdot (\sigma_u^2 + \sigma_v^2)^{-1}$, and the cost reduces to $\Omega(\sigma_\xi)$.

Let s^C denote the optimally combined signal. Clean exporters assign risk labels to tankers by applying a threshold K to s^C: a tanker is classified as High-risk (H) if s^C exceeds K, and as Low-risk (L) otherwise. The threshold K is chosen to satisfy

$$\mathbb{E}(s^C \geq K) = \lambda. \tag{7}$$

4.1.4 Sanctions

The sanctioning authority allocates to each tanker j monitoring resources a_j, which depend on a noisy signal s_j^A that it receives about the tanker's type:

$$a_j = f(s_j^A), \quad \text{and} \quad f'(\cdot) > 0. \tag{8}$$

where

$$s_j^A = \mathbb{I}_j(\text{Bad Tanker}) + u_j, \quad u_j \sim \mathbb{N}(0, \sigma_u^2). \tag{9}$$

To simplify analysis, we specify

$$f(s_j^A) = \frac{\exp(s_j^A)}{\mathbb{E}[\exp(s_j^A)]}, \tag{10}$$

so that a_j is always positive and $\mathbb{E}[a_j] = 1$.

The probability that a given tanker j is sanctioned is:

$$\mathbb{P}_j(\text{sanctioned}) = a_j \cdot [\pi_1 \cdot \mathbb{I}_j(\text{Bad tanker}) + \pi_2 \cdot \mathbb{I}_j(\text{deal with R})], \tag{11}$$

where $\mathbb{I}_j$(Bad tanker) equals one if this tanker is Bad and zero otherwise, $\mathbb{I}_j$(deal with R) equals one if the tanker is currently transporting sanctioned oil and zero otherwise, π_1 is the probability per unit of resource that the tanker is sanctioned for past violations (i.e., for being Bad), and π_2 is this probability of a current violation (i.e., transporting sanctioned oil from Rogue exporters in the current period). The parameters, π_1 and π_2, reflect legal and other technical constraints that the authority faces in sanctioning tankers.

4.1.5 Beliefs about exposure to sanctions-related penalties

Tankers: Depending on its type and risk label, each tanker belongs to one of four groups – GL, GH, BL, or BH. Each tanker j knows its group and the sanctioning authority's resource allocation mechanism as in Eq.(8), but does not know a_j (i.e., exactly how intensely the authority is monitoring it). Given Eq.(11), we can obtain the expected sanction probabilities for tankers in each group, conditional on their information set. Adding to the group index a 0 or 1, depending on whether the tanker is currently transporting sanctioned oil (i.e. whether $\mathbb{I}(\text{deal with R})$ is 0 or 1), these conditional probabilities are:

$$\begin{aligned}
\bar{w}_{GL0} &= 0, & \bar{w}_{GL1} &= \mathbb{E}[a|G,L]\cdot\pi_2,\\
\bar{w}_{GH0} &= 0, & \bar{w}_{GH1} &= \mathbb{E}[a|G,H]\cdot\pi_2,\\
\bar{w}_{BL0} &= \mathbb{E}[a|B,L]\cdot\pi_1, & \bar{w}_{BL1} &= \mathbb{E}[a|B,L]\cdot[\pi_1+\pi_2]\\
\bar{w}_{BH0} &= \mathbb{E}[a|B,H]\cdot\pi_1, & \bar{w}_{BH1} &= \mathbb{E}[a|B,H]\cdot[\pi_1+\pi_2],
\end{aligned} \tag{12}$$

where we again omit the tanker index j for brevity. Note that $\bar{w}_{GL0} = \bar{w}_{GH0} = 0$ because Good tankers know that they will not be sanctioned due to past violations; this assumes that these tankers are always "cleared" if investigated for past violations.

Clean exporters: They calculate the expected probabilities of L and H tankers being sanctioned as follows:

$$\begin{aligned}
\bar{w}_L^C &= \frac{Q_{GL}}{Q_{GL}+Q_{BL}}\bar{w}_{GL0} + \frac{Q_{BL}}{Q_{GL}+Q_{BL}}\bar{w}_{BL0},\\
\bar{w}_H^C &= \frac{Q_{GH}}{Q_{GH}+Q_{BH}}\bar{w}_{GH0} + \frac{Q_{BH}}{Q_{GH}+Q_{BH}}\bar{w}_{BH0},
\end{aligned} \tag{13}$$

where Q_{GL}, Q_{GH}, Q_{BL}, and Q_{BH} are the number of tankers in each group that Clean exporters hire in equilibrium. These probabilities imply that Clean exporters realize that the proportion of Good vs. Bad tankers among the tankers *that they hire* is likely to be different from the proportion of Good vs. Bad tankers among *all* tankers, as tankers of specific types could be more likely to engage with them.[6]

Rogue exporters: They only know tanker types but not their risk labels. So they calculate

---

[6] Note that here each party's beliefs depend on quantities that they do not know ex-ante (e.g., Clean exporters never observe Q_{GL} and Q_{BL} separately – they only observe the sum, as they cannot distinguish between G and B tankers). In a rational expectations equilibrium, however, each party will make guesses about these unknown components such that their beliefs are consistent with model solutions (and therefore with each other).

the expected sanction probabilities for the G and B tankers that they hire as:

$$
\begin{aligned}
\bar{w}_G^R &= \frac{Q_{GL}^R}{Q_{GL}^R + Q_{GH}^R}\bar{w}_{GL1} + \frac{Q_{GH}^R}{Q_{GL}^R + Q_{GH}^R}\bar{w}_{GH1}, \\
\bar{w}_B^R &= \frac{Q_{BL}^R}{Q_{BL}^R + Q_{BH}^R}\bar{w}_{BL1} + \frac{Q_{BH}^R}{Q_{BL}^R + Q_{BH}^R}\bar{w}_{BH1}.
\end{aligned} \tag{14}
$$

4.1.6 Exporters' optimization

The objective of the exporters is to maximize profits. The per-unit oil trade revenue for Clean exporters, $\tilde{r}$, and for Rogue exporters, $\tilde{r}^R$, are

$$
\tilde{r} = r + \tilde{\epsilon}, \qquad \tilde{r}^R = r^R + \tilde{\epsilon}, \tag{15}
$$

where r and r^R are mean values and $\tilde{\epsilon}$ is the aggregate shock, uniformly distributed over $[-\sigma, \sigma]$. For simplicity, we assume a fixed ratio $\frac{r^R}{r}$, which implies a constant discount on sanctioned oil.[7]

Exporters' costs include (i) information acquisition cost (only for Clean exporters), (ii) payments to tankers, and (iii) penalties if the tanker they employ is sanctioned. The sanction penalties are z^C and z^R for Rogue and Clean exporters, respectively. We assume that $z^R > z^C$, as a Rogue exporter's entire oil cargo could be impounded if a tanker carrying it is caught and sanctioned, which is significantly more costly than the delays and penalties facing a Clean exporter for hiring a Bad tanker that now gets sanctioned but is currently carrying unsanctioned oil.

Clean exporters: They first choose their detection technology (i.e., σ_ξ and χ), and then set prices p_L and p_H. Given the i.i.d. environment, the decisions on detection technology and prices are taken period by period.

We first consider the Clean exporters' price decision. Since tankers and exporters are randomly matched in a pair (Section 4.1.2), each Clean exporter occupies a $\frac{1}{N}$ portion of the market, and does not compete with other Clean exporters. For a given detection technology, each Clean exporter optimally chooses prices to maximize profits:

$$
\tilde{\Pi} = \max_{\{p_L, p_H\}} \frac{1}{N}(Q_{GL} + Q_{BL})(\tilde{r} - p_L - \bar{w}_L^C z^C) + \frac{1}{N}(Q_{GH} + Q_{BH})(\tilde{r} - p_H - \bar{w}_H^C z^C), \tag{16}
$$

where $\bar{w}_L^C$ and $\bar{w}_H^C$ are expected sanction probabilities defined in Eq.(13). The Q's are functions of the prices, which are derived from tankers' optimization (as described below).

---

[7] Consistent with the data in the months surrounding Windward's disclosure (see, e.g., Inside Shipping, 2024).

We use backward induction to derive Clean exporters' detection technology. Let Π denote the expected equilibrium profit, given these exporters' belief about Windward's precision $\tilde{\sigma}_v^{-2}$ (again, symmetry implies same equilibrium beliefs):

$$\Pi(\sigma_\xi, \chi; \tilde{\sigma}_v^2) = \mathbb{E}[\tilde{\Pi} | \tilde{\sigma}_v^2] \tag{17}$$

The optimization problem is then

$$\max_{\{\sigma_\xi, \chi \in \{0,1\}\}} \quad \Pi(\sigma_\xi, \chi; \tilde{\sigma}_v^2) - \Omega(\sigma_\xi) - \chi \cdot P^W, \tag{18}$$

with optimality conditions

$$\frac{\partial \Pi(\sigma_\xi, \chi)}{\partial \sigma_\xi} = \frac{\partial \Omega(\sigma_\xi)}{\partial \sigma_\xi}, \quad \chi = \begin{cases} 0, & \text{if } \Pi(\sigma_\xi, 1) - \Pi(\sigma_\xi, 0) < P^W, \\ 1, & \text{if } \Pi(\sigma_\xi, 1) - \Pi(\sigma_\xi, 0) > P^W, \\ 0 \text{ or } 1, & \text{if } \Pi(\sigma_\xi, 1) - \Pi(\sigma_\xi, 0) = P^W. \end{cases} \tag{19}$$

Rogue exporters: They solve the optimization problem only for prices:

$$\max_{\{p_G^R, p_B^R\}} \; (Q_{GL}^R + Q_{GH}^R)(\tilde{r}^R - p_G^R - \bar{w}_G^R z^R) + (Q_{BL}^R + Q_{BH}^R)(\tilde{r}^R - p_B^R - \bar{w}_B^R z^R), \tag{20}$$

where $\bar{w}_G^R$ and $\bar{w}_B^R$ are expected sanction probabilities defined in Eq.(14).

4.1.7 Tankers' optimization

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

Let V_G and V_B denote the values of Good and Bad tankers, and V_{GL}, V_{GH}, V_{BL}, V_{BH} denote the values further conditional on risk labels. Let $\theta_B = \mathbb{P}[H|B]$ and $\theta_G = \mathbb{P}[H|G]$. The tanker values satisfy the following relations:

$$V_G = \theta_G V_{GH} + (1 - \theta_G) V_{GL}, \quad V_B = \theta_B V_{BH} + (1 - \theta_B) V_{BL}. \tag{21}$$

The tankers' Bellman equations are

$$\begin{aligned}
V_{GL} &= \mathbb{E}\Big[\max\Big\{p_L - \bar{w}_{GL0}\cdot z + \beta V_G, \quad p_G^R - \tilde{c} - \bar{w}_{GL1}\cdot z + \beta V_B\Big\}\Big],\\
V_{GH} &= \mathbb{E}\Big[\max\Big\{p_H - \bar{w}_{GH0}\cdot z + \beta V_G, \quad p_G^R - \tilde{c} - \bar{w}_{GH1}\cdot z + \beta V_B\Big\}\Big],\\
V_{BL} &= \mathbb{E}\Big[\max\Big\{p_L - \bar{w}_{BL0}\cdot z + \beta V_B, \quad p_B^R - \tilde{c} - \bar{w}_{BL1}\cdot z + \beta V_B\Big\}\Big],\\
V_{BH} &= \mathbb{E}\Big[\max\Big\{p_H - \bar{w}_{BH0}\cdot z + \beta V_B, \quad p_B^R - \tilde{c} - \bar{w}_{BH1}\cdot z + \beta V_B\Big\}\Big]. \qquad (22)
\end{aligned}$$

where $\bar{w}_{GL0}$, $\bar{w}_{GL1}$, etc., are expected sanction probabilities defined in Eq.(12), and $\tilde{c}$ is a random cost that a tanker incurs each period when dealing with Rogue exporters. Observed only by the tanker before choosing an exporter to deal with, this cost is assumed to be uniformly i.i.d. over $[0, \bar{c}]$ and allows us to obtain interior solutions for tanker supply. Such a cost may arise from efforts to avoid detection, reflecting tanker-specific variables like its location, route, and evasion technology.

The penalty imposed on a Bad tanker when sanctioned is z, assumed for simplicity to equal the tanker's entire value. Given discounting, this penalty is

$$z = \beta V_B. \qquad (23)$$

The critical values $\tilde{c}_j$ at which tanker j would be indifferent between dealing with a Clean or Rogue exporter are then:

$$\begin{aligned}
c_{GL} &= p_G^R - p_L - (\bar{w}_{GL1} - \bar{w}_{GL0})z - \beta(V_G - V_B), \qquad & c_{BL} &= p_B^R - p_L - (\bar{w}_{BL1} - \bar{w}_{BL0})z,\\
c_{GH} &= p_G^R - p_H - (\bar{w}_{GH1} - \bar{w}_{GH0})z - \beta(V_G - V_B), \qquad & c_{BH} &= p_B^R - p_H - (\bar{w}_{BH1} - \bar{w}_{BH0})z.
\end{aligned}$$

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

$$V_G = \frac{\theta_G \mathbb{E}\left[\frac{(c_{GH})^2}{2\bar{c}} + p_H\right] + (1-\theta_G)\mathbb{E}\left[\frac{(c_{GL})^2}{2\bar{c}} + p_L\right] - [\theta_G w_{GH0} + (1-\theta_G) w_{GL0}]z}{1-\beta}, \qquad (24)$$

$$V_B = \frac{\theta_B \mathbb{E}\left[\frac{(c_{BH})^2}{2\bar{c}} + p_H\right] + (1-\theta_B)\mathbb{E}\left[\frac{(c_{BL})^2}{2\bar{c}} + p_L\right] - [\theta_B w_{BH0} + (1-\theta_B) w_{BL0}]z}{1-\beta}, \qquad (25)$$

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

4.1.8 Equilibrium

A stationary equilibrium is characterized by: (i) Clean exporters' choice of detection technology (i.e., σ_ξ and χ); (ii) Clean and Rogue exporters' freight rates $p_G^R(\tilde{\epsilon})$, $p_B^R(\tilde{\epsilon})$, $p_L(\tilde{\epsilon})$, $p_H(\tilde{\epsilon})$, where $\tilde{\epsilon}$ refers to the aggregate shock in the current period; and (iii) tankers' decision rule based on the critical values $c_{GL}(\tilde{\epsilon})$, $c_{GH}(\tilde{\epsilon})$, $c_{BL}(\tilde{\epsilon})$, $c_{BH}(\tilde{\epsilon})$, which satisfy

- Optimality: Clean exporters optimally choose detection technology; Clean and Rogue exporters optimally set prices to maximize their profits; tankers maximize their values

- Market clearing: the market clears for all. In particular, for Clean exporters (Q_{xy}) and Rogue exporters (Q^R_{xy}) this implies, respectively

$$Q_{xy} = A_{xy}\left(1 - \frac{c_{xy}}{\bar{c}}\right), \quad Q^R_{xy} = A_{xy}\left(\frac{c_{xy}}{\bar{c}}\right), \quad xy \in \{GL, GH, BL, BH\},$$

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

The model solution is described in detail in Internet Appendix F1.

4.2 Model calibration and fit

4.2.1 Directly calibrated parameters

We focus on changes in market equilibrium in response to Windward's disclosure. We directly calibrate seven parameters. The discount factor β is set 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 r to one – this is Clean exporters' mean oil trade revenue per voyage, or, more precisely, their profits before shipping costs. 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 (i.e., just before Windward's disclosure) as per CSI Market, 2023). The mean revenue of Rogue exporters is $r^R = (70 \times 0.277 - 4)/(70 \times 0.277) = 0.794$, assuming a \$4 per barrel discount (as for Russian exports to India in the second half of 2023, Reuters, Sept. 2023).

Given 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^C 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 $z^R = 50/15.5 = 3.226$.

We calibrate the precision of the public signal, σ_ζ^{-2}, by converting the pseudo-R^2 of public data's predictability in Figure A-5 in the Internet Appendix. The 7.5% pseudo-R^2 implies σ_ζ to be 1.111. The values of directly calibrated parameters are in Table 4 Panel A.

---

[8] The average number of trips per year in our fixtures sample is 3.4. Assuming a time-discounting rate of 0.95, a 10% scrap value, and a 15-year life cycle gives $0.95 \times \exp(\log(0.1)/(15 \times 3.4)) \approx 0.9$.

4.2.2 Parameters derived from moments matching

We calibrate the remaining eight parameters by matching nine moments. The eight parameters are: the per-unit-resource sanction probabilities π_1 and π_2, max operation cost when dealing with Rogue exporters $\bar{c}$, volatility of oil trade revenue σ, precision of the authority's signal σ_u^{-2}, precision of the Windward's signal σ_v^{-2}, Clean exporters' belief on Windward's signal $\tilde{\sigma}_v$, and coefficient on information acquisition cost κ.

To identify the sanction probabilities π_1, π_2 and precision of the authority's signal, we mainly rely on three moments: actual proportion of sanctioned tankers, disclosure's effect on tankers' fixtures, and disclosure's effect on tankers' routes (these come from the DiD estimations in Table 3).

We identify the maximum operation cost $\bar{c}$ and the volatility of oil trade revenue σ mainly by matching three moments: the volatility-to-mean ratios for low- and high-risk tankers, and the volatility-to-mean ratio for top US charterers' revenue.

We use the pseudo-R^2 of the Windward-and-public-data combined signal's predictability (i.e., 21.3% in Figure A-5 in the Internet Appendix) to pin down the precision of the Windward's signal. We use Windward's market value to uncover Clean exporters' belief $\tilde{\sigma}_v$ about Windward's σ_v. Finally, we identify the coefficient on information acquisition cost κ by assuming this cost to be the major component of compliance costs and then match the ratio of non-compliance costs to compliance costs (see Secureframe (2025)). The moments and the parameters derived from moments matching are in Panels B and C of Table 4 .

4.2.3 Calibrating $\tilde{\sigma}_v$

Recall that in Section 4.1.3 we left flexible the Clean exporters' beliefs about Windward's precision σ_v^{-2}. The most straightforward assumption would be to impose full rational expectations, i.e., $\tilde{\sigma}_v = \sigma_v$.

However, adopting this assumption results in a significant divergence between the model-implied values and the observed data, most notably with respect to Windward's valuation. Specifically, our model would imply the valuation of its signal about suspect tankers to be \$5,405 mln in July 2023, whereas in fact the total valuation of Windward at that time was \$42 mln and the estimated upper bound of the valuation of its signal was \$27 mln (in the Internet Appendix we derive this number as an upper bound of the valuation specific to Windward's suspect tanker list, given that it has other business lines).

This divergence is so substantial that it is robust to many alternative modeling assumptions. Intuitively, a data provider like Windward's valuation depends on the value of the signal they provide, which, in turn, depends on market participants' expected incremental profits from buying that signal, which we estimate to be at most \$27 million in July 2023. This suggests that no market participant believed that buying Windward's signal would increase their profits beyond this amount (i.e., a high $\tilde{\sigma}_v$). However, as demonstrated in Section 3.3, the revenue of high-risk tankers declined by 13% following the Windward disclosure, and these tankers started avoiding sanctioned exporters. Both of these findings indicate a large spillover of high-risk tankers in the Clean oil shipping market, reducing Clean exporters' shipping costs. Such a reduction, in turn, would imply incremental profits for them that are an order of magnitude higher than \$27 mln. But had these exporters anticipated such large incremental profits (i.e., had they correctly believed $\tilde{\sigma}_v = \sigma_v$), they should have been willing to pay much more for Windward's data, and its valuation should have correspondingly been an order of magnitude higher.

Moving away from the equality $\tilde{\sigma}_v = \sigma_v$, however, brings forth the challenge of pinning down what exact value to assign to $\tilde{\sigma}_v$ in the calibration among the many possibilities. Next we explain how we do this simply by using Windward's actual valuation.

Clean exporter i's willingness-to-pay for Windward's signal, $\tilde{\delta}_i$, depends on her belief about Windward's precision $\tilde{\sigma}_v^{-2}$ as follows:

$$\tilde{\delta}_i(\tilde{\sigma}_v^2) = \Pi(\sigma_{\xi,i}, 1; \tilde{\sigma}_v^2) - \Pi(\sigma_{\xi,i}, 0; \tilde{\sigma}_v^2), \tag{26}$$

where $\Pi(\cdot; \tilde{\sigma}_v^2)$ is the expected profit as in Eq.(18) and its argument one or zero denotes whether Windward's signal is bought or not. Let the true value of Windward's information be δ. We assume that Clean exporters' prior belief about δ is distributed $\mathbb{N}(\mu_\delta, \Sigma_\delta)$. Additionally, each Clean exporter receives a noisy signal (e.g., from tankers newly sanctioned in every period or from Windward providing a trial version of their list before purchase) about the true value of δ: $x_i = \delta + \eta_i$, with independent noise $\eta_i \sim \mathbb{N}(0, \Sigma_\eta)$. Each Clean exporter i rationally updates her willingness-to-pay based on her signal x_i and Windward's market value M^W:

$$\tilde{\delta}_i = \mathbb{E}[\delta | x_i, M^W]. \tag{27}$$

The absence of near-arbitrage opportunities implies that

$$M^W = \sum_{i=1}^{N} \tilde{\delta}_i + \epsilon^W, \tag{28}$$

where $\epsilon^W \sim \mathbb{N}(0, \Sigma_\epsilon)$ comes from noise traders.[9] Solving Eq.(27) and Eq.(28) gives

$$\tilde{\delta}_i = \mu_\delta + A(x_i - \mu_\delta) + B\left(\frac{M^W}{N} - \mu_\delta\right), \tag{29}$$

---

[9] For example, if $M^W << \sum_{i=1}^N \tilde{\delta}_i$, all Clean exporters can buy Windward together to realize the surplus.

where A solves $\frac{\Sigma_\eta}{\Sigma_\epsilon}(N-1)(N+\frac{\Sigma_\eta}{\Sigma_\delta})A^3 + A - 1 = 0$ and $B = \frac{N(N-1)\Sigma_\eta A^2}{\Sigma_\epsilon + N(N-1)\Sigma_\eta A^2}$. The derivation details are in the Internet Appendix.

Note that, since Windward is a start-up with very short history and sanction rules are constantly updated yet little new information on violators is provided every period, a Clean exporter who has not yet decided to buy Windward's signal (and maybe has a trial version with a small subset of the data to decide whether to buy) would likely face an information environment characterized by (i) diffuse priors on Windward's accuracy ($\Sigma_\delta/\Sigma_\eta \to \underline{c} > 0$), and (ii) a very noisy signal x_i ($\Sigma_\epsilon/\Sigma_\eta \to 0$). In the Internet Appendix, we show that under these circumstances $A \to 0$ and $B \to 1$, and hence

$$\tilde{\delta}_i \to \frac{M^W}{N}. \tag{30}$$

Eq.(30) combined with $\tilde{\delta}_i(\tilde{\sigma}_v^2) = \Pi(\sigma_\xi, 1; \tilde{\sigma}_v^2) - \Pi(\sigma_\xi, 0; \tilde{\sigma}_v^2)$ allows us to uncover Clean exporters' belief from Windward's market value. In such an equilibrium, Clean exporters are indifferent between $\chi = 0$ and $\chi = 1$ (i.e., buying or not buying Windward's signal), as per Eq.(19), assuming that Windward is competitively priced (i.e., it's value reflects the entire surplus it generates).

4.2.4 Implications of $\tilde{\sigma}_v \neq \sigma_v$

The fact that Clean exporters' ex-ante $\tilde{\sigma}_v$ did not equal the true σ_v reveals an important friction in the market. Had these exporters been able to recognize Windward's true accuracy in detecting violators, they would have expected much higher incremental profits from its information, and hence would have been willing to pay for Windward accordingly. A significantly higher valuation would have motivated the creation of similar companies who could charge appropriately for detecting suspected violators in different domains. This would represent an endogenous market-based resolution of the detection challenges.

However, with diffuse priors and inadequate additional information to update beliefs about Windward's precision over time, it is rational for Clean exporters to try to infer it by learning from the company's market value. And this means that a wide range of beliefs and associated valuations can be justified in equilibrium. If the market somehow underestimates Windward's value at some point, noisy learning will prevent convergence to the truth. Such a situation, where Windward gets stuck in a value trap, could then justify the company releasing their signal at very low cost to demonstrate its true accuracy (and thus benefit its other lines of business). Not only is this precisely what happened in the disclosure event we study, but Windward's stock market valuation also more than doubled in the six months after disclosure. It was eventually taken private at a valuation of \$280 mln in December

2024.

4.2.5 Model fit

We report the moment matching results in Table 4. The model captures the effect of Windward's on high-risk tankers' fixtures (13% drop in both data and model) and routes (17% drop in both data and model).

Importantly, the values of other model variables not targeted in our calibration exercise align well 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 $\sigma = 0.241$, consistent with the Oil Volatility Index OVX around 28% in 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 168%-187%, per Panel D in Table 4, consistent with the sanctions premium reported in the media for Russian oil delivered to Indian and Chinese ports (e.g., Inside Shipping, 2024, which reports calculations from Argus Media). Third, model-based tanker values are around \$30 million, which also matches actual values for a 20-year-old tanker as per Xclusiv (2023).

Overall, our model-based estimates appear to align reasonably well with their corresponding values reported by alternative sources, offering some confidence in the reliability of the quantities generated by the model.

4.3 Model implications

In Figure 4, we present the model implications of Windward's disclosure by varying the precision of Windward's signal, σ_v^{-2}. We focus on the pre-period equilibrium where almost no Clean exporters buy Windward's signal (as Clean exporters are indifferent between buying or not buying Windward's signal) – which we model as Windward's signal precision being zero before disclosure – and the post-period where everyone has access for free. The two vertical dashed lines indicate these two equilibria, respectively.

Panel A of Figure 4 illustrates tankers' expectations about the authority's resource allocation. First, as the authority allocates resources based on a noisy signal about tanker types, Bad tankers rationally anticipate that they tend to face more monitoring than Good tankers. Second, as the precision of Windward's signal increases, High-risk (Low-risk) tankers expect that they will face more (less) monitoring. This is because in the post-period, risk labels are derived based on Windward's signal, and Windward's signal contains information about the authority's signal. Therefore, as the precision of Windward's signal increases, tankers can

Table 4: **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. These results are discussed in Sections 4.2 and 4.3.

Panel A: Directly calibrated parameters

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)	\$15.5m, normalized to one	r	1
Mean oil trade revenue (Rogue exporters)	\$4 per barrel discount	r^R	0.794
Sanction penalty (Clean exporters)	\$4m, comparable to shipping costs	z^C	0.258
Sanction penalty (Rogue exporters)	\$50m, based on the median value from actual cases of seized oil	z^R	3.226
(Im)precision of the public signal	implied by a pseudo-R^2 of 7.5% (public data's predictability in Figure A4)	σ_ζ	1.111

Panel B: Matched moments

Moment	Model analog	Data	Model
Proportion of sanctioned tankers (average of pre- and post-Windward periods)	$\frac{\mathbb{E}[Q_{Sanc}^{pre}+Q_{Sanc}^{post}]}{2}$	0.7%	0.7%
Disclosure effect on high-risk tankers' fixtures (DiD estimation in Table 3)	$\frac{\mathbb{E}[p^{post}-p^{pre}\vert\text{Post H}]}{\mathbb{E}[p^{pre}\vert\text{Post H}]}$	-0.13	-0.130
Disclosure effect on high-risk tankers' routes (DiD estimation in Table 4)	$\frac{\mathbb{E}[b^{post}-b^{pre}\vert\text{Post H}]}{\mathbb{E}[b^{pre}\vert\text{Post H}]}$	-0.17	-0.170
Annualized volatility of low-risk tankers' fixtures before Windward's disclosure (divided by mean)	$\frac{\text{std}[p_L^{pre}]}{\mathbb{E}[p_L^{pre}]}$	0.77	0.774
Annualized volatility of high-risk tankers' fixtures before Windward's disclosure (divided by mean)	$\frac{\text{std}[p_H^{pre}]}{\mathbb{E}[p_H^{pre}]}$	0.81	0.807
Annualized volatility of top US charterers' revenue before Windward's disclosure (divided by mean)	$\frac{\text{std}[(Q_{GL}^{pre}+Q_{GH}^{pre}+Q_{BL}^{pre}+Q_{BH}^{pre})\tilde{r}]}{\mathbb{E}[(Q_{GL}^{pre}+Q_{GH}^{pre}+Q_{BL}^{pre}+Q_{BH}^{pre})\tilde{r}]}$	0.30	0.141
Precision of the combined signal in the post-period (implied by a pseudo-R^2 of 21.3%)	$\sigma_\zeta^{-2}+(\sigma_u^2+\sigma_v^2)^{-1}$	2.345	2.345
Valuation of the Windward's signal (based on share price one-day before disclosure)	$\Pi(\sigma_\xi^*,1)-\Pi(\sigma_\xi^*,0)$	\$27m	\$27m
Non-compliance costs divided by compliance costs (see Secureframe (2025))	$\frac{\mathbb{E}[(Q_{GL}^{pre}+Q_{BL}^{pre})\bar{w}_L^C z^C+(Q_{GH}^{pre}+Q_{BH}^{pre})\bar{w}_H^C z^C]}{\Omega(\sigma_\xi^*)}$	2.71	2.70

Panel C: Parameters derived from moments matching

Variable	Symbol	Value
Sanction probability per unit of resource (previous violation)	π_1	0.39%
Sanction probability per unit of resource (current violation)	π_2	4.18%
Max operation cost when dealing with Rogue exporters	$\bar{c}$	0.952
Volatility of oil trade revenue	σ	0.241
(Im)precision of the authority's signal	σ_u	0.485
(Im)precision of the Windward's signal (true value)	σ_v	0.645
(Im)precision of the Windward's signal (Clean exporters' ex-ante belief)	$\bar{\sigma}_v$	13.019
Coefficient on information acquisition cost (×10^3)	κ	4.428

Panel D: Imputed quantities (in mln)

Period	$\mathbb{E}(p_H)$	$\mathbb{E}(p_L)$	$\mathbb{E}(p_B^R)$	$\mathbb{E}(p_G^R)$	V_B	V_G
Pre-Windward	2.71	2.83	6.27	8.57	26.64	29.77
Post-Windward	2.43	2.75	6.09	8.77	24.79	28.52

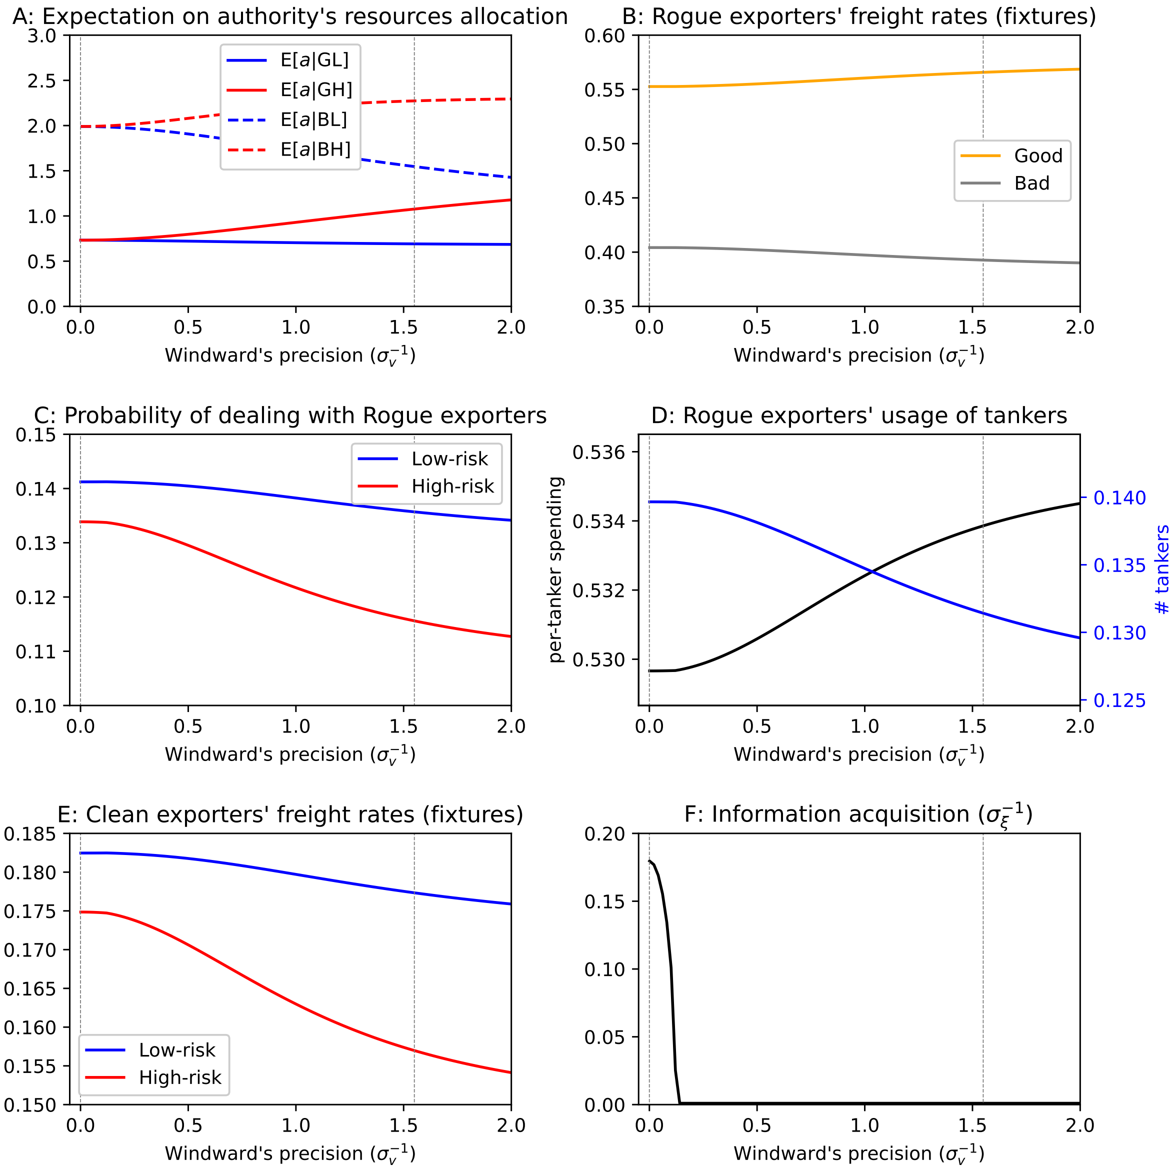

Figure 4: This figure presents the predictions from the calibrated model from Section 4. We solve the model by discretizing the aggregate shock $\tilde{\epsilon}$ on a grid. We plot the average model outcomes for the values of $\tilde{\epsilon}$ as functions of pseudo-R^2, which is backed out from the 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 21.3%, respectively.

learn more about the authority's resource allocation from risk labels. This learning effect is the driving force of the disclosure effect on route changes.

Panel B shows that as Windward's precision improves, Rogue exporters pay marginally higher freight rates to hire Good tankers (i.e., $\mathbb{E}[p_G^R]$ increases) and lower rates for Bad tankers (i.e., $\mathbb{E}[p_B^R]$ decreases).

Panel C shows the probabilities of tankers engaging in transactions with Rogue exporters. High-risk tankers are less likely to engage with Rogue exporters. Importantly, as Windward's precision improves, the likelihood of engagement with Rogue exporters declines more sharply for High-risk tankers, as those tankers learn from risk labels that the sanctioning authority now allocates more monitoring resources to them. These results align with our empirical findings on route changes.

Panel D sheds light on aggregate implications for Rogue exporters. The black line 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 B in the figure). This increase is, of course, not offset by the drop in freight rates that they pay to Bad tankers because there are many more Good tankers than Bad tankers.

The blue line in Panel D shows that with higher classification precision, the total number of oil tankers engaging with Rogue exporters decreases. This follows from Panel C. Quantitatively, the information shock accounts for a 5.9% drop in Rogue exporters' total hiring of tankers. 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 – although this likely involves higher costs of transportation (if they were lower cost, then ROgue exporters would have used them rather than tankers in the pre-period).

Panel E shows the relation between Windward's precision and the equilibrium freight rates (i.e., the fixtures $\mathbb{E}[p_L]$ and $\mathbb{E}[p_H]$) from Clean exporters. First, as precision increases, High-risk tankers' freight rates sharply decrease. 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.

Second, Panel E reveals a counterintuitive result: as precision increases, Low-risk tankers' freight rates also decrease. This outcome arises from two opposing forces, which we derive from the first-order conditions of exporters' optimization. The first is an information effect: improved precision makes the L label indicate lower sanction risk and hence more desirable. The second is a competition effect: as precision increases, High-risk tankers that are increasingly excluded from the Rogue market now relocate to the Clean market, intensifying competition and putting downward pressure on freight rates for Low-risk tankers. Our quantitative analysis indicates that the competition effect outweighs the information effect, resulting in a net decline in freight rates for Low-risk tankers. In the following section, we empirically test this counterintuitive prediction using a shift-share design.

Panel F shows that as Windward's precision improves, Clean exporters have less incentive to acquire private information. Our quantitative analysis here shows that as Windward's precision increases, Clean exporters quickly switch their information source to Windward's signal. The substitution is because of the decreasing marginal benefits of information for Clean exporters. Moreover, Clean exporters prefer Windward's signal as it contains information about the authority's signal, which allows Clean exporters to infer sanction risk not only through tanker types, but also through the noise term in the authority's signal. This result highlights the key tension between Clean exporters' incentives and sanction enforcement.

Finally, in Table 5, 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.[10]

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 1.8% (from 8.21 to 8.35 mln per tanker-trip) due to the increase in p_G^R, while the cost for Clean exporters decreases by 4.3% (from 2.80 to 2.68 mln per tanker-trip) due to reductions in both p_L and p_H. On the tanker side, fixture income decreases by 5.4% (from 3.13 to 2.96 mln per tanker-trip) for Bad tankers and decreases by 1.7% (from 3.67 to 3.61 mln per tanker-trip) for Good tankers. 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 (from the data). Therefore, Windward's disclosure leads to \$512 (\$919) million annual losses for Rogue exporters (Bad tankers). These losses accrue as gains for the sanctions-compliant agents, with a \$1.4 billion overall gain for Clean exporters and Good tankers combined.

4.4 Testing the counter-intuitive model implication

Here we test the main 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 4). We conduct tests exploiting charterer-level variations.

Ideally, we would want to measure 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 $k \in 1, \ldots, K$, be the varying exposures to increased competition (i.e., "shifts")

---

[10] Alternatively, we could have assumed that sanctioned oil 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.

Table 5: **Model-implied overall effects**

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)	8.21	3.13	2.80	3.67
Post-Windward (per tanker-trip)	8.35	2.96	2.68	3.61
Percentage change	1.8%	-5.4%	-4.3%	-1.7%
Aggregate effect (annual, mln) (-/+ indicates loss/gain)	-512	-919	2,674	-1,243
Total aggregate effect (annual, mln) (-/+ indicates loss/gain)	-1,431		1,431	

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.[11]

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:

$$fixture_{c,i,k,t} = \beta_0 \times \mathbb{I}_{\{\text{High-competition}\}} + \beta_1 \times \mathbb{I}_{\{\text{High-competition}\}} \times \mathbb{I}_{\{t \geq 0\}} + \text{FE} + \epsilon_{c,i,k,t}. \qquad (31)$$

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 a cross-sectional cutoff (e.g., top

---

[11] 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.

Table 6: **Testing the counter-intuitive model implication: Competition spillovers**

We test the model's implication that Windward's disclosure reduces low-risk tankers' freight rates (competition effects), exploiting 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 $\mathbb{I}_{\{\text{High-competition}\}}$ and its interaction with the post-period indicator. $\mathbb{I}_{\{\text{High-competition}\}}$ equals one if the charterer-level shift-share competition measure, as described in Section 4.4, 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

50%, 40%, or 30%), and is zero otherwise. Our model predicts a negative DiD coefficient β_1.

Estimation results in Table 6 are consistent with model predictions: higher competition corresponds to 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. These results not only provide an out-of-sample test for our model, but suggest that policymakers should be mindful of negative spillovers on compliant parties in equilibrium.

4.5 Counterfactual analysis: Increasing penalties under noisy detection

Finally, to further explore implications for enforcement design, we use our model to study a key counterfactual: what would happen if the sanctioning authority increased the penalty (z) on Clean exporters who were found to be using sanctions-violating tankers.[12] This scenario aligns with much of the focus in political and media discussions, as well as literature suggesting that aligning the incentives of all parties with that of the sanctioning authority should enhance enforcement.

---

[12] Recall that for Rogue exporters, we already assumed that the sanctioning authority impounds the entire oil cargo – i.e., the penalty z^R cannot be increased any further. Hence we focus on Clean exporters here.

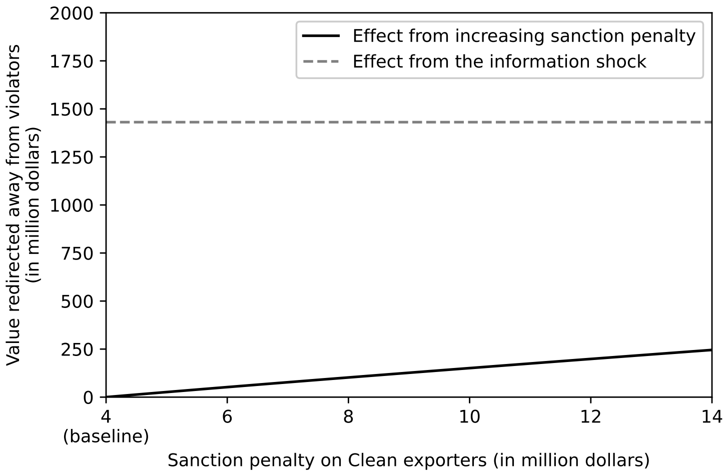

Figure 5: This figure presents the counterfactual analysis of increasing the sanction penalty. The x-axis is the sanction penalty for Clean exporters (i.e., parameter z) expressed in million dollars. The y-axis is the model-implied effects of the value redirected away from violators, as calculated in Table 5. For comparison, we plot the effect from the information shock using a dotted line.

In Figure 5, we conduct this analysis by varying z from 0.258 (4 mln dollars) to 0.903 (14 mln dollars) and fixing detection precision at the pre-Windward level. We focus on the effect of the value redirected away from violators, as calculated in Table 5.

The figure reveals that increasing sanction penalties in the presence of noisy detection has little effect on enforcement. This is due to two reasons: first, as Clean exporters cannot perfectly observe tanker compliance status, increasing penalties raises the expected cost of inadvertently hiring a violating tanker. In response, Clean exporters reduce their overall demand for tanker services to manage this risk. This hurts both Good and Bad tankers. Of course, Bad tankers are hurt more, as Clean exporters start especially avoiding high-risk labeled ones that are more likely to be Bad. However, this shrinkage in demand from the compliant sector increases tanker supply and lowers cost for Rogue exporters, making the total value transferred away from violators in response to higher penalties negligible.

5 Conclusion

Our results show that public disclosure of suspected sanctions violators strengthens market-based sanctions enforcement by reducing profits from violations. One implication of our results is that some amount of public investment in helping the market detect violators may be required to strengthen sanctions enforcement, at least in the current scenario.

References

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

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* .

Bai, X., Fernández-Villaverde, J., Li, Y., Xu, L. and Zanetti, F. (2025), 'The (Un)Intended Consequences of Oil Sanctions Through the Dark Shipping of Sanctioned Oil', *Working Paper* .

Baldwin, D. A. (2020), 'Economic statecraft: New edition', *Princeton University Press* .

Banerjee, S. (2011), 'Learning from prices and the dispersion in beliefs', *The Review of Financial Studies* **24**(9), 3025–3068.
URL: *https://doi.org/10.1093/rfs/hhr050*

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

Bond, P., Edmans, A. and Goldstein, I. (2012), 'The real effects of financial markets', *Annu. Rev. Financ. Econ.* **4**(1), 339–360.

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.

Cipriani, M., Goldberg, L. S. and La Spada, G. (2023), 'Financial sanctions, swift, and the architecture of the international payment system', *Journal of Economic Perspectives* **37**(1), 31–52.

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

De Long, J. B., Shleifer, A., Summers, L. H. and Waldmann, R. J. (1990), 'Noise trader risk in financial markets', *Journal of political Economy* **98**(4), 703–738.

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.

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.

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 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.

Goldstein, I. and Yang, L. (2019), ‘Good disclosure, bad disclosure’, *Journal of Financial Economics* **131**(1), 118–138.

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.

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.

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.

Hufbauer, G. C. and Jung, E. (2020), ‘What’s new in economic sanctions?’, *European economic review* **130**, 103572.

Huneke, M. (2023), The Blind Men and the Elephant, Hughes Hubbard & Reed.

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. (2023), International sanctions and limits of Lerner symmetry, *in* ‘AEA Papers and Proceedings’, Vol. 113, pp. 33–38.

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

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’’.

LSEG: DN207488 (2023), DN207488 - Windward Sanction Compliance data to become Fee Liable and Orderable - Customer Action Required, Data Notification.

Norrlöf, C. (2021), Economic statecraft: Finance and Money, Report series, Atlantic Council.

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.

Trebbi, F. and Zhang, M. B. (2022), ‘The cost of regulatory compliance in the United States’.

Van Genugten, J. (2021), ‘Conscripting the global banking sector: Assessing the importance and impact of private policing in the enforcement of U.S. Economic sanctions’, *Berkeley Business Law Journal* **18**, 136–164.

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.