Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Vessel Profiles in the Shadows: A Binary Logit Analysis of Technical and Operational Divergence in Global Tanker Fleets

Dionysios Polemis
University of Piraeus
Konstantinos Likas
konstantinos.likas@db.com
Deutsche Bank (Germany)

Research Article

Keywords: shadow fleet, crude tankers, sanctions evasion, maritime risk screening, logistic regression, designation risk

Posted Date: April 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9098405/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: Competing interest reported. Konstantinos Likas reports employment with Deutsche Bank AG. Dr. Dionysios Polemis declares no competing interests.

Vessel Profiles in the Shadows: A Binary Logit Analysis of Technical and Operational Divergence in Global Tanker Fleets

Abstract

This article examines whether sanction-designated crude tankers can be distinguished from the wider global fleet through vessel-level characteristics and a cross-validated screening model. Using 2,277 crude oil tankers from Clarksons SIN, of which 382 are positive on a composite designation indicator, the study separates structural inference from operational prediction. Logit and probit models with a spline age term show that designation risk is concentrated in older Aframax tonnage. Relative to Aframax vessels, Suezmaxes and VLCCs exhibit substantially lower odds of designation, while shuttle-tanker status is not robust. The age effect is distinctly non-linear, peaking in the early-to-mid twenties. A screening model that adds pooled flag information attains ROC-AUC 0.9306, PR-AUC 0.7288, and a Brier score of 0.0750. The results support vessel-centred triage, but the model is best interpreted as a designation-risk classifier rather than direct proof of latent shadow-fleet membership.

Keywords: shadow fleet; crude tankers; sanctions evasion; maritime risk screening; logistic regression; designation risk

1 Introduction

The tightening of post-2022 sanctions on Russian seaborne oil exports has transformed tanker markets, insurance arrangements, and maritime compliance practices. Rather than eliminating Russian hydrocarbon exports, the sanctions regime appears to have accelerated the emergence of an alternative shipping ecosystem centred on opaque ownership, permissive registries, non-Western service chains, and ageing tanker tonnage (Grebennikov, 2024;

Poussenkova, 2024; Astrov et al., 2024; Cardoso et al., 2025; Kilian et al., 2024). In policy and practitioner discourse this ecosystem is commonly described as the "shadow fleet", though the term also overlaps with related notions such as dark fleets and grey fleets (Domballe et al., 2024; Hilgenstock et al., 2024; Yusuf, 2024; Eurasia Institute, 2025).

This phenomenon matters beyond sanctions compliance. A shadow-fleet architecture implicates maritime safety, environmental exposure, marine insurance, port-state control, and broader questions of grey-zone coercion and strategic adaptation (Kozmenko, 2024; Antonova & Karatash, 2025; Jones, 2025; Meiertöns, 2025). The fleet is therefore not simply a sanctions-evasion device; it is also a logistics system embedded in geopolitical contestation, regulatory fragmentation, and widening ecological risk (Gözübenli, 2025; Knapp et al., 2025; Rodriguez-Diaz et al., 2025). The literature increasingly connects these developments to hybrid pressure strategies, fragmented commodity governance, and the political economy of sanctions circumvention (Omokaro et al., 2025; Madaki, 2025; Ibadoghlu, 2025).

Existing research has generally emphasised mutable signals of evasion. Scholars and analysts focus on shell-company structures, beneficial-ownership opacity, AIS spoofing, dark activity, deceptive port calls, ship-to-ship transfers, and rapid registry changes (Androjna et al., 2023; Dutt et al., 2024; Domballe et al., 2024; Nainani, 2025). These are undeniably central to sanctions circumvention. Yet they are also strategically adaptive. Names can change, registries can shift, paperwork can be layered, and digital traces can be manipulated. By contrast, a vessel's age, dimensional class, and broad technical profile are slow-moving or effectively fixed. If sanction-designated vessels are drawn disproportionately from a recognisable physical subset of the global fleet, then the shadow-fleet problem is not only administrative or digital. It is also material.

The present article builds on that intuition. It does not deny the importance of ownership opacity, digital deception, or jurisdictional arbitrage. Rather, it asks whether the vessels that appear in the observed designation universe exhibit a durable technical profile. This is useful for two reasons. Analytically, it helps move beyond descriptive claims about "old ships" and "risky registries" towards a structured vessel-level account of designation risk. Operationally, it may provide a more stable first-stage screening basis for regulators, financial institutions, chartering counterparties, and port authorities confronting incomplete or degraded information environments (Eun-Ju Lee, 2025; Rodriguez-Diaz et al., 2025; Knapp et al., 2025).

Rigour requires a clear statement of the dependent variable. The positive outcome in this article is not an independently verified census of latent shadow-fleet membership. It is a designation proxy: a vessel is coded positive if it appears on at least one of the sanction or

designation lists in the analytical file. The paper therefore studies *observed designation risk*. That distinction matters because a model may capture the technical correlates of sanctioning behaviour, the correlates of underlying fleet structure, or both. Treating the output as a designation-risk classifier is therefore more defensible than claiming perfect observation of the hidden fleet itself.

Research Objectives

Against that background, the article has four research objectives.

1. To identify the vessel-level technical characteristics associated with observed designation risk within the global crude tanker fleet.
2. To test whether the age effect is linear or whether designation risk is concentrated within a specific ageing window.
3. To evaluate whether an expanded operational screening model materially improves predictive performance once pooled flag information is added.
4. To compare EU, UK, and Ukrainian designation profiles in order to assess whether different sanctioning regimes appear to target the same underlying vessel structure.

The paper argues that designation risk is concentrated in older Aframax tonnage, that the age effect is distinctly non-linear, and that an operational classifier performs substantially better once pooled flag information is introduced. At the same time, the article shows why explanatory technical models and predictive screening models should not be conflated. The former identify structural vessel correlates; the latter also learn the administrative geography of observed sanctions practice.

2 Literature Review

The literature relevant to the shadow fleet is broad, but it can be organised into four main clusters: sanctions economics, maritime-legal geography, operational evasion and digital deception, and maritime safety and environmental risk.

A first cluster examines the macroeconomics and strategic logic of sanctions on energy trade. This literature analyses the oil price cap, export rerouting, discounted crude sales, intermediary markets, and the resilience of Russian hydrocarbon revenues under conditions of

enforcement pressure (Johnson et al., 2023; Kilian et al., 2024; Cardoso et al., 2025; Wang et al., 2025; Omokaro et al., 2025; Astrov et al., 2024). More general work on sanctions as instruments of coercion also shows that economic restrictions frequently generate adaptive market structures rather than simple trade collapse, especially in commodity-intensive sectors with transport flexibility (Madaki, 2025). Within this broader framework, the shadow fleet can be understood as a logistics response to regulatory pressure rather than merely as an incidental by-product.

A second cluster addresses the legal geography and institutional architecture of the fleet. Scholars and policy analysts emphasise ownership opacity, shell-company layers, permissive or opportunistic registries, and the strategic use of loosely governed maritime jurisdictions (Domballe et al., 2024; Hilgenstock et al., 2024; Nainani, 2025; Eun-Ju Lee, 2025; Eurasia Institute, 2025). These studies are crucial because they show how sanctions circumvention is embedded in fragmented governance structures rather than in isolated vessel behaviour. Work on regional grey-zone trade and sanctions circumvention also reinforces the point that shadow shipping is nested in wider networks of intermediary states, financial adaptation, and geopolitical bargaining (Ibadoghlu, 2025; Yusuf, 2024).

A third cluster concerns operational evasion and maritime deception. Research on AIS manipulation, dark activity, and suspicious ship-to-ship operations shows how sanctioned trade increasingly relies upon altered digital signatures, route masking, and signal disruption (Androjna et al., 2023; Dutt et al., 2024). Such work is methodologically innovative and operationally relevant, but it also underscores a strategic weakness: digital and administrative indicators are highly mutable. This creates a strong rationale for complementing them with vessel-centred features that are materially harder to alter.

A fourth cluster focuses on safety, environmental risk, and the ageing structure of the fleet. Researchers have warned that shadow-fleet activity is associated with older tanker tonnage, weaker oversight, elevated casualty risk, and greater ecological exposure, especially in environmentally sensitive sea lanes (Knapp et al., 2025; Gözübenli, 2025; Rodriguez-Diaz et al., 2025; Polayeva & Pustovoit-Dnikov, 2024). This literature is especially important for the present paper because it repeatedly suggests, though often descriptively, that ageing is central to the shadow-fleet problem. What it less often does is test whether age retains explanatory power once class and other technical variables are held constant.

A parallel strategic literature connects the maritime shadow fleet to broader forms of hybrid pressure and seabed insecurity. Work on hybrid warfare, shadow energy logistics, and sea-blindness highlights how commercial shipping can become an instrument of strategic leverage while also exploiting the legal complexity of the maritime domain (Kozmenko, 2024;

Antonova & Karatash, 2025; Jones, 2025; Meiertöns, 2025). This reinforces the need to think of vessel profiles not simply as commercial characteristics, but also as strategic assets within a contested regulatory order.

The present article contributes to these literatures in three ways. First, it shifts attention from descriptive claims about obscure or ageing tonnage to an explicit vessel-level estimation strategy. Secondly, it separates technical explanation from operational screening, which the existing literature often treats together. Thirdly, it compares regime-specific designation patterns across the EU, UK, and Ukraine, thereby moving beyond aggregate discussion of the shadow fleet towards a more disaggregated account of observed sanctioning behaviour.

3 Methodology

3.1 Data and Sample Construction

The analytical file, acquired from Clarksons Shipping Intelligence Network (SIN), contains 2,277 crude oil tankers and 22 vessel-level fields, including vessel name, type, class, gross tonnage, deadweight, length overall, build year, builder, flag state, main engine fuel type, power type, current global zone, destination port country or region, IMO number, and four sanction or designation indicators for the European Union, the United Kingdom, Ukraine, and the United States. The empirical design is deliberately vessel-centred. This follows from a recurring theme in the shadow-fleet literature: many of the most visible indicators of sanctions evasion, such as ownership layering, AIS irregularities, flag switching, and deceptive routing, are strategically mutable and therefore analytically useful but unstable as baseline identifiers (Androjna et al., 2023; Dutt et al., 2024; Domballe et al., 2024; Nainani, 2025). By contrast, class, age, and broad technical characteristics change slowly or not at all.

The composite dependent variable is coded as one if a vessel appears on at least one of the four regime lists and zero otherwise. This yields 382 positive observations and a composite positive rate of 16.7765%. Because one vessel has missing fuel information, the effective estimation sample for the inference and prediction models is 2,276. The class distribution consists of 694 Aframax tankers, 675 Suezmax tankers, and 908 VLCCs. Mean vessel age is 14.24 years, the median is 15 years, and the sample maximum is 35 years.

This coding choice should be interpreted carefully. The positive class does not represent a complete census of latent shadow-fleet membership; rather, it identifies vessels that fall within the *observed designation universe*. That distinction is important because several stud-

ies suggest that shadow-fleet activity extends beyond formally designated entities and that designation practice itself is shaped by legal capacity, evidentiary standards, and strategic priorities (Hilgenstock et al., 2024; Domballe et al., 2024; Yusuf, 2024; Eun-Ju Lee, 2025). The models therefore estimate designation risk, not metaphysical membership in a hidden fleet.

Table 1: Sample structure and descriptive statistics

Statistic	Value	Comment
Total vessels in file	2,277	Raw analytical dataset
Inference sample	2,276	One observation excluded due to missing fuel data
Prediction sample	2,276	Same effective sample as the inference model
Composite positives	382	Positive on at least one regime list
Composite positive rate	16.7765%	Moderately imbalanced classification problem
Mean age	14.24 years	Snapshot-based measure
Median age	15 years	Central tendency of ageing distribution
Maximum age	35 years	Oldest vessel in sample

3.2 Variable Design

The explanatory variables are chosen to privilege physical and slow-moving vessel characteristics over highly mutable administrative identifiers. This reflects both substantive theory and the practical enforcement logic found in the shadow-fleet literature. Research on sanctions circumvention repeatedly stresses that evasive actors exploit opacity in corporate structure, registry use, and digital signals (Androjna et al., 2023; Dutt et al., 2024; Nainani, 2025). A vessel-level specification therefore offers a more durable baseline from which to assess whether the observed designation universe has a recognisable material profile.

- **Age**: defined as the reference snapshot year minus year built.
- **Class**: categorical variable taking the values Aframax, Suezmax, and VLCC.
- **Simplified fuel**: grouped into IFO 380, VLS IFO, and Other/Alternative.
- **Shuttle tanker**: binary variable equal to one for shuttle tankers and zero otherwise.
- **Pooled flag state**: an operational variable used only in the screening model, retaining the largest registry groups separately and collapsing the remainder into an "Other" category.

Several other vessel dimensions, such as deadweight, gross tonnage, and length overall, are highly collinear with class and therefore omitted from the core models in order to preserve interpretability. Likewise, pooled flag state is excluded from the technical-inference model because it captures administrative geography as well as vessel structure. This matters analytically. The literature shows that registry choice is often part of the evasive strategy itself, rather than a purely exogenous vessel attribute (Domballe et al., 2024; Eun-Ju Lee, 2025; Nainani, 2025). Flag is therefore highly relevant for screening, but less appropriate as a primary explanatory variable in a model intended to isolate structural technical correlates.

3.3 Descriptive Diagnostics

The descriptive plots already reveal strong differentiation across the sample. These figures do not replace multivariate analysis, but they establish why a non-linear and vessel-centred modelling strategy is warranted.

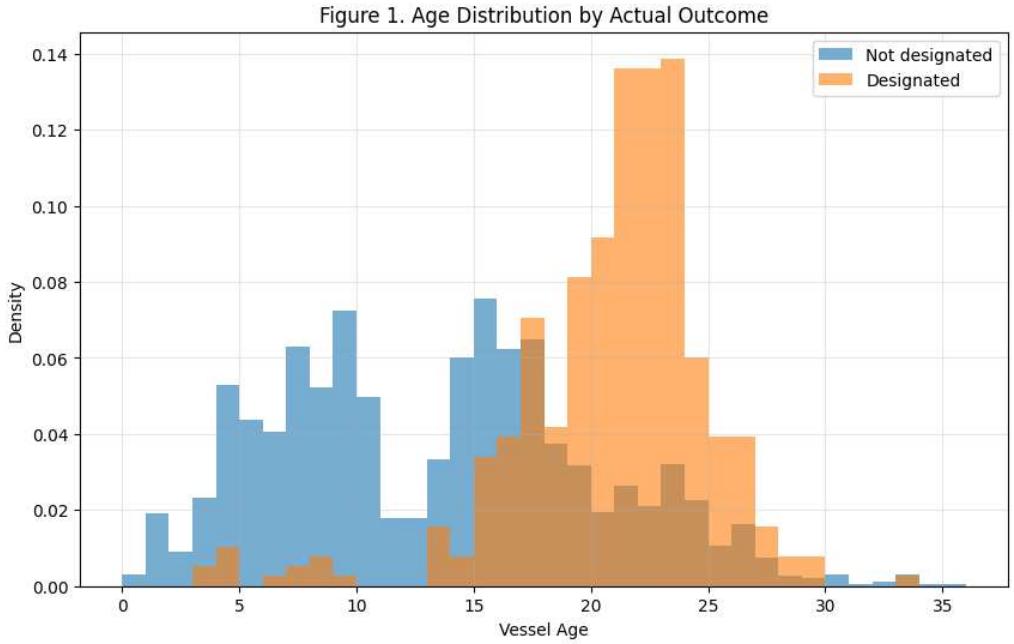


Figure 1: Age distribution by actual outcome

Figure 1 shows that the designated subset is concentrated much more heavily in the late-teen to mid-twenties age range than the non-designated fleet. This is consistent with a broad strand of the literature portraying the shadow fleet as disproportionately composed of ageing tonnage and vessels operating under weaker maintenance and oversight conditions (Knapp et al., 2025; Gözübenli, 2025; Polayeva & Pustovoit-Dnikov, 2024; Rodriguez-Diaz et al.,

2025). What the figure adds is a more precise empirical hint that the relevant age profile may be curved rather than monotonic.

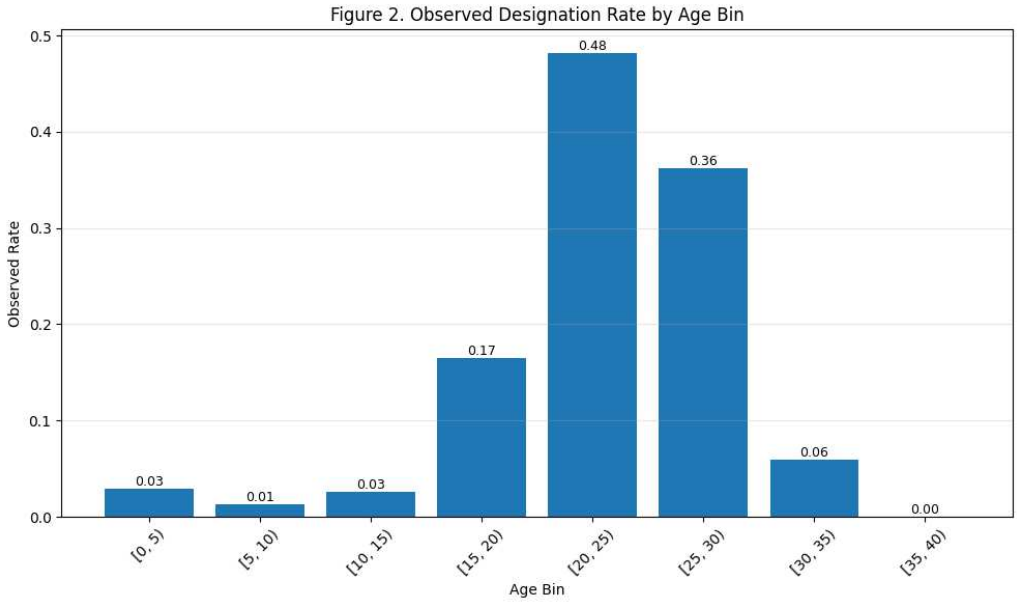


Figure 2: Observed designation rate by age bin

Figure 2 confirms that intuition. The raw designation rate remains very low through the 0–15 year range, rises sharply in the 15–20 year bin, peaks at 0.48 in the 20–25 year bin, and remains elevated at 0.36 in the 25–30 year bin before declining in the sparse upper tail. Existing policy and risk studies often note that older vessels matter, but they rarely specify *where* within the age distribution the risk is concentrated (Knapp et al., 2025; Gözübenli, 2025; Rodriguez-Diaz et al., 2025). The descriptive evidence here suggests that the relevant pattern is not simply ageing in general, but a distinct intermediate-to-old segment.

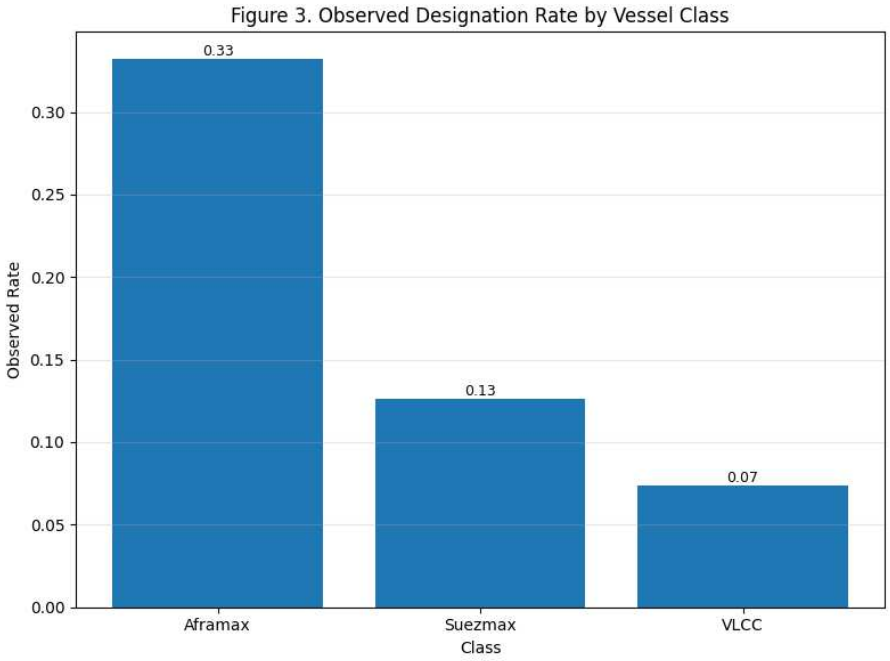


Figure 3: Observed designation rate by vessel class

Figure 3 reveals a pronounced class gradient. Aframax vessels have an observed positive rate of 0.3314, compared with 0.1259 for Suezmax and 0.0738 for VLCC. This is broadly consistent with qualitative discussions in the sanctions literature that emphasise the flexibility, second-hand market depth, and operational adaptability of smaller tanker segments within shadow-fleet development (Hilgenstock et al., 2024; Domballe et al., 2024; Cardoso et al., 2025). The descriptive evidence therefore suggests that the designated subset is not merely older, but compositionally concentrated in a particular class segment.

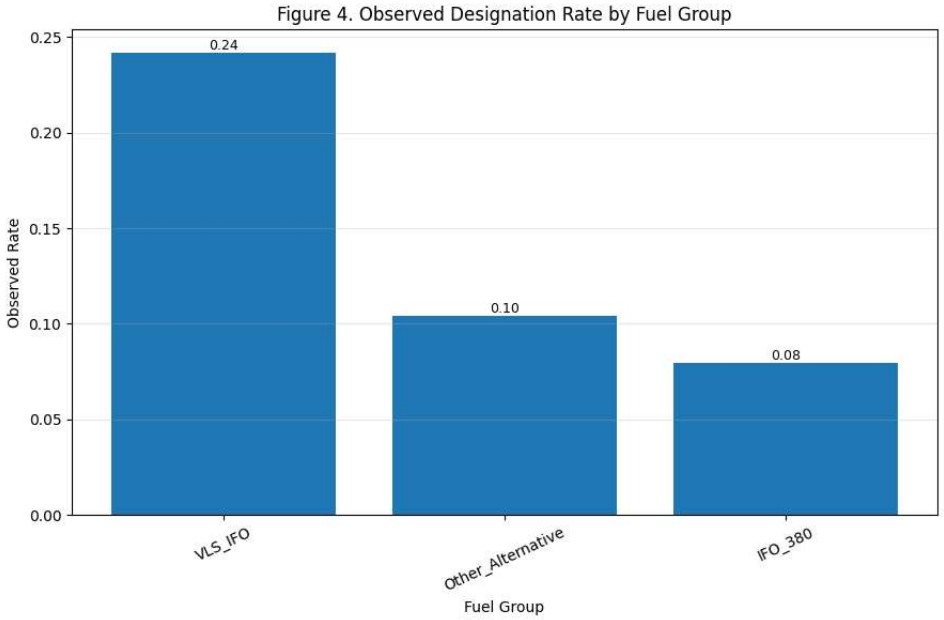


Figure 4: Observed designation rate by fuel group

Figure 4 shows that VLS IFO vessels have the highest raw positive rate, with IFO 380 substantially lower. This should be interpreted cautiously. Fuel type may partly proxy for vessel vintage, retrofit history, or class composition, and the literature does not yet provide a settled theoretical expectation on fuel as a primary determinant of shadow-fleet participation (Knapp et al., 2025; Rodriguez-Diaz et al., 2025). The figure is therefore suggestive, but not conclusive.

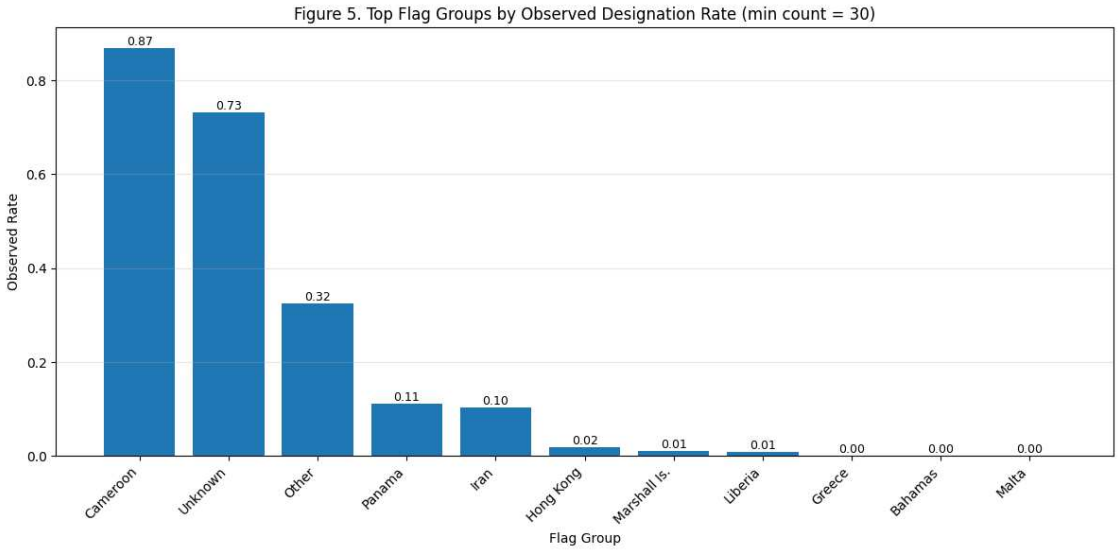


Figure 5: Top flag groups by observed designation rate (minimum count = 30)

Figure 5 highlights large observed differences across pooled registry groups. Categories such as Cameroon and Unknown display very high raw positive rates, whereas major registries such as Liberia, Malta, Greece, and the Bahamas sit near zero. This is consistent with the literature on permissive registries, opportunistic flag use, and jurisdictional arbitrage in the shadow-fleet ecosystem (Domballe et al., 2024; Eun-Ju Lee, 2025; Nainani, 2025). At the same time, it reinforces the methodological decision to treat flag state as an operational screening variable rather than a core structural regressor.

3.4 Econometric Strategy

The econometric strategy separates explanatory inference from predictive screening. This distinction is central both methodologically and substantively. Much of the current literature mixes descriptive identification, behavioural monitoring, and sanctions-risk assessment within a single narrative (Domballe et al., 2024; Hilgenstock et al., 2024; Rodriguez-Diaz et al., 2025). The present design instead distinguishes between a model intended to identify structural vessel correlates of designation and a model intended to maximise out-of-sample screening performance.

The baseline explanatory specification is a binary logit:

$$P(Y_i = 1 \mid X_i) = \Lambda(X_i\beta) = \frac{e^{X_i\beta}}{1 + e^{X_i\beta}},$$

where Y_i is the composite designation indicator and X_i contains age, class, simplified fuel, and shuttle-tanker status. Age is entered as a four-degree-of-freedom B-spline rather than as a linear term because the descriptive profile clearly shows curvature. Aframax and IFO 380 are used as reference categories. The baseline model reports HC3-robust standard errors, and a builder-clustered specification is added as a robustness check to allow for unobserved correlation across shipyards. That clustering choice is reasonable given the possibility that common build practices, yard-specific characteristics, or unobserved technical features may induce within-builder dependence.

A probit model with the same covariates is also estimated. This is not because the probit link is expected to dominate substantively, but because sign consistency across links strengthens confidence that the main results do not depend upon a particular distributional assumption. In a literature where much evidence remains descriptive or case-based, such link-function robustness is worth establishing explicitly.

The predictive screening model retains age, class, fuel, and shuttle-tanker status, but adds

pooled flag state. It is estimated using five-fold stratified cross-validation and evaluated through out-of-fold predictions. Performance is summarised with ROC-AUC, PR-AUC, Brier score, threshold trade-offs, calibration, and permutation importance. This design is appropriate for a moderately imbalanced classification problem and aligns with the operational realities of sanctions enforcement, where authorities and private actors face trade-offs between broad surveillance and narrow high-confidence escalation (Eurasia Institute, 2025; Jones, 2025).

Finally, regime-specific technical models are estimated separately for the EU, UK, and Ukraine. The US-specific model is omitted because only 13 positives are available, which is too sparse for stable maximum-likelihood estimation in the present setting. This choice is analytically preferable to forcing an unstable rare-events specification into the main paper, particularly given that the purpose of the regime models is comparison rather than exhaustive estimation.

3.5 Identification and Scope Conditions

The empirical strategy is designed to identify the vessel-level correlates of *observed designation risk*, not to establish a causal model of shadow-fleet participation in the strict sense. The dependent variable is constructed from formal sanction or designation outcomes across the European Union, the United Kingdom, Ukraine, and the United States. As a consequence, the positive class reflects the intersection between underlying vessel characteristics and the legal, institutional, and strategic logic of sanctioning authorities. The models therefore estimate which types of vessels are more likely to appear in the observed designation universe, rather than proving that those vessels constitute the full latent shadow fleet (Hilgenstock et al., 2024; Domballe et al., 2024; Eun-Ju Lee, 2025; Yusuf, 2024).

This distinction is important because shadow-fleet activity almost certainly extends beyond formally designated vessels. Existing studies suggest that sanctions circumvention relies on adaptive networks of ownership opacity, permissive registries, insurance substitution, and deceptive operational practices, many of which may remain undetected or undesignated for substantial periods (Androjna et al., 2023; Dutt et al., 2024; Nainani, 2025; Cardoso et al., 2025). The dependent variable used here should therefore be interpreted as a legally observed and institutionally filtered proxy, not as a complete census of illicit maritime activity.

The paper consequently adopts an associational rather than a causal interpretive framework. Coefficients from the technical models indicate how age, class, fuel grouping, and shuttle-tanker status are related to designation risk conditional on the included covariates, but

they should not be read as demonstrating that these characteristics *cause* designation or illicit behaviour. In particular, age may proxy for acquisition cost, market marginality, maintenance quality, or the commercial attractiveness of second-hand tonnage; class may proxy for route flexibility, transaction size, or market visibility; and flag state, where included in the predictive model, may partly capture the geography of enforcement itself (Knapp et al., 2025; Rodriguez-Diaz et al., 2025; Domballe et al., 2024).

A related limitation concerns timing. Age is measured using the snapshot year embodied in the estimation file rather than the vessel's age at first designation event. The resulting estimates should therefore be read as cross-sectional relationships observed at the time of the compiled dataset, not as event-timed causal effects. Likewise, the regime-specific models illuminate differences in observed designation profiles across sanctioning authorities, but they do not identify the internal decision rules or evidentiary processes that generated those listings. This is especially relevant because legal thresholds, intelligence inputs, and strategic priorities vary across regimes (Eun-Ju Lee, 2025; Jones, 2025; Ibadoghlu, 2025).

For these reasons, the paper deliberately separates technical inference from operational screening. The technical models are intended to recover stable vessel-level patterns associated with the observed designation universe. The predictive model is intended to support triage and prioritisation under real-world conditions of incomplete information. Its stronger performance should not be taken to imply a deeper structural explanation; rather, it reflects the fact that pooled flag state and similar variables encode both vessel risk and observed sanctioning geography. The most appropriate interpretation of the model, therefore, is as a designation-risk classifier with practical screening value, not as an automated mechanism for inferring latent shadow-fleet membership.

4 Model

4.1 Baseline Technical Inference

Table 2 reports the principal odds-ratio results from the baseline logit model and the builder-clustered robustness specification.

Table 2: Baseline technical designation-risk models

	Logit (HC3)			Logit (Builder-clustered)		
Variable	**OR**	**95% CI**	***p*-value**	**OR**	**95% CI**	***p*-value**
Suezmax (vs Aframax)	0.340	0.245 to 0.471	< 0.001	0.340	0.215 to 0.537	< 0.001
VLCC (vs Aframax)	0.143	0.100 to 0.206	< 0.001	0.143	0.090 to 0.228	< 0.001
Other/Alternative fuel (vs IFO 380)	4.850	1.723 to 13.652	0.0028	4.850	0.906 to 25.952	0.0650
VLS IFO (vs IFO 380)	0.761	0.541 to 1.072	0.1184	0.761	0.484 to 1.198	0.2385
Shuttle tanker	0.644	0.279 to 1.487	0.3024	0.644	0.327 to 1.269	0.2032

Notes: Age is modelled using a four-degree-of-freedom spline and is interpreted through Figure 6 rather than through the sp basis terms themselves. Full odds-ratio tables are reported in the Appendix.

The class effect is the strongest and most stable result. Relative to Aframax vessels, Suezmaxes have approximately 66% lower odds of designation, while VLCCs have approximately 86% lower odds. These magnitudes are virtually unchanged once errors are clustered by builder. This finding is important because much of the existing literature treats the shadow fleet as an ageing and opaque aggregate, without estimating class-conditioned designation differentials explicitly (Domballe et al., 2024; Hilgenstock et al., 2024; Rodriguez-Diaz et al., 2025). The present results extend that literature by showing that the observed designation universe is not simply older on average; it is also compositionally concentrated in smaller crude tanker segments.

That pattern sits comfortably alongside broader arguments about sanctions adaptation. Studies of shadow-fleet formation and oil-cap evasion suggest that circumvention is facilitated by flexible, transferable, second-hand tonnage capable of operating within fragmented logistics chains (Cardoso et al., 2025; Domballe et al., 2024; Hilgenstock et al., 2024). The Aframax concentration found here is therefore consistent with a market logic of modularity and operational adaptability. By contrast, the much lower odds associated with VLCCs suggest that very large tonnage is less central to the observed designation profile, whether because of cost, visibility, operational constraints, or enforcement salience.

Fuel effects are less robust. In the HC3 model, the Other/Alternative category is associated with much higher odds of designation, but that estimate weakens materially and becomes only marginal once builder clustering is introduced. VLS IFO is not statistically significant in either baseline logit. The broader literature does not offer strong prior expectations that fuel configuration alone should dominate shadow-fleet participation, and the present results support that caution (Knapp et al., 2025; Rodriguez-Diaz et al., 2025). Fuel appears to matter, but not with the same consistency as age and class. Shuttle-tanker status is

consistently negative and insignificant, which is also instructive. It suggests that specialised offshore service functionality is not a principal marker of the designated subset once basic fleet structure is taken into account.

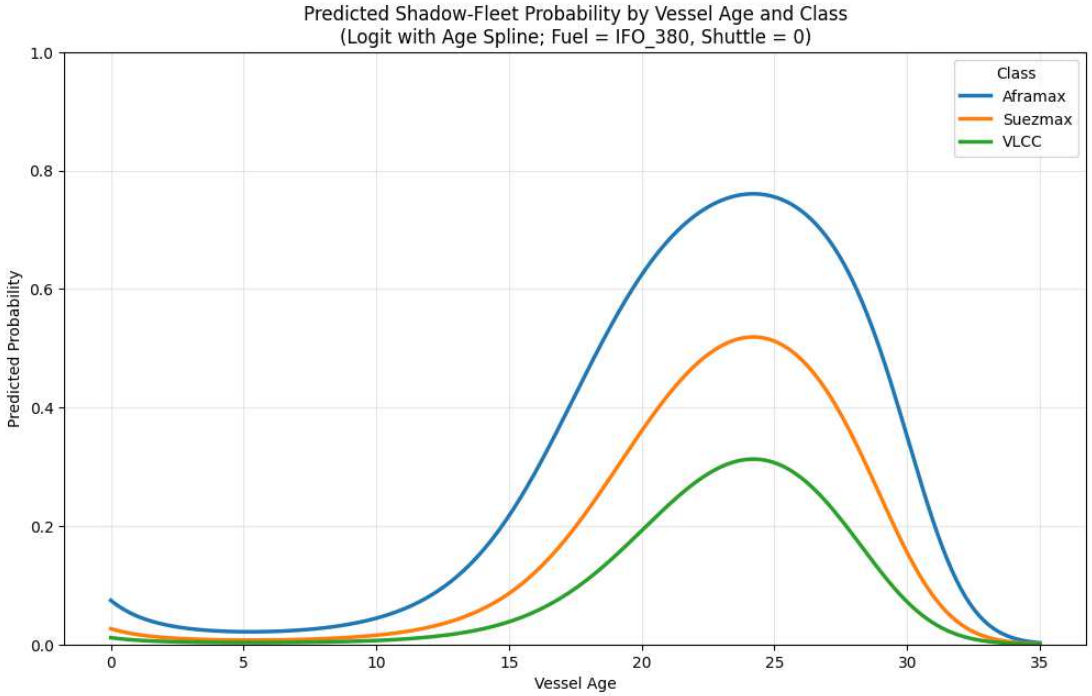


Figure 6: Predicted shadow-fleet probability by vessel age and class

Figure 6 provides the substantive interpretation of the age spline. Holding fuel at IFO 380 and shuttle status at zero, predicted designation probability is low for very young vessels, rises sharply from the mid-teens onwards, peaks in the early-to-mid twenties, and then declines in the upper tail. Aframax remains above Suezmax and VLCC across the full age profile. This is a more precise finding than the broad claim, common in the policy literature, that "older vessels are riskier" (Knapp et al., 2025; Gözübenli, 2025; Polayeva & Pustovoit-Dnikov, 2024). What the model shows is that designation risk is concentrated in a distinct intermediate-to-old age window rather than increasing monotonically across the entire age distribution.

That non-linearity matters theoretically. The literature on shadow fleets and sanctions circumvention often links ageing tonnage to lower acquisition cost, weaker resale prospects in mainstream markets, and greater willingness to operate outside conventional compliance systems (Hilgenstock et al., 2024; Domballe et al., 2024; Rodriguez-Diaz et al., 2025). The spline profile fits that logic more closely than a linear effect would. Extremely young vessels are unlikely to enter the designated subset because they remain too valuable and too visible

within standard markets; extremely old vessels may be too sparse, too technically degraded, or too commercially marginal to dominate the observed designation universe. The risk peak in the early-to-mid twenties is therefore economically and operationally plausible.

The probit robustness model preserves the same sign pattern and does not materially alter inference. Relative to Aframax, both Suezmax and VLCC remain strongly negative. The positive coefficient attached to the Other/Alternative fuel category also survives in the HC3 probit. This consistency is valuable in a research area where many claims are still based on descriptive reports or case studies rather than formal cross-sectional estimation (Domballe et al., 2024; Hilgenstock et al., 2024; Eurasia Institute, 2025). The central structural conclusions do not depend upon the choice of binary link function.

4.2 Operational Screening Performance

The predictive model adds pooled flag information to the technical feature set and evaluates performance out of sample. Table 3 summarises the principal results.

Table 3: Out-of-fold screening performance and threshold trade-offs

Panel A. Overall performance

Metric	ROC-AUC	PR-AUC	Brier score	Positive rate	Best Youden threshold
Value	0.9306	0.7288	0.0750	0.1678	0.1817

Panel B. Threshold trade-offs

Threshold	Precision	Recall	Specificity	F1	Flagged share
0.1000	0.4496	0.9450	0.7666	0.6093	0.3528
0.1800	0.5210	0.9110	0.8310	0.6629	0.2935
0.1817	0.5226	0.9084	0.8326	0.6635	0.2917
0.2500	0.5614	0.8377	0.8680	0.6723	0.2504
0.5000	0.7104	0.6099	0.9498	0.6563	0.1441

Notes: Metrics are based on five-fold stratified cross-validation and out-of-fold predictions. At the Youden threshold of 0.1817, the confusion matrix is TN = 1,577, FP = 317, FN = 35, TP = 347. At a threshold of 0.50, the confusion matrix is TN = 1,799, FP = 95, FN = 149, TP = 233.

The screening model performs strongly. ROC-AUC reaches 0.9306, PR-AUC 0.7288, and the Brier score is 0.0750. Because the positive class accounts for only 16.8% of the sample, the PR-AUC is especially informative: the classifier is not merely ordering vessels well,

but recovering substantial positive-class signal in a genuinely useful way. This strengthens the more impressionistic claims in practitioner and policy work that certain vessel profiles, registry patterns, and operational signals can materially narrow the enforcement search space (Domballe et al., 2024; Eurasia Institute, 2025; Jones, 2025). The contribution here is to show that those intuitions survive out-of-fold validation.

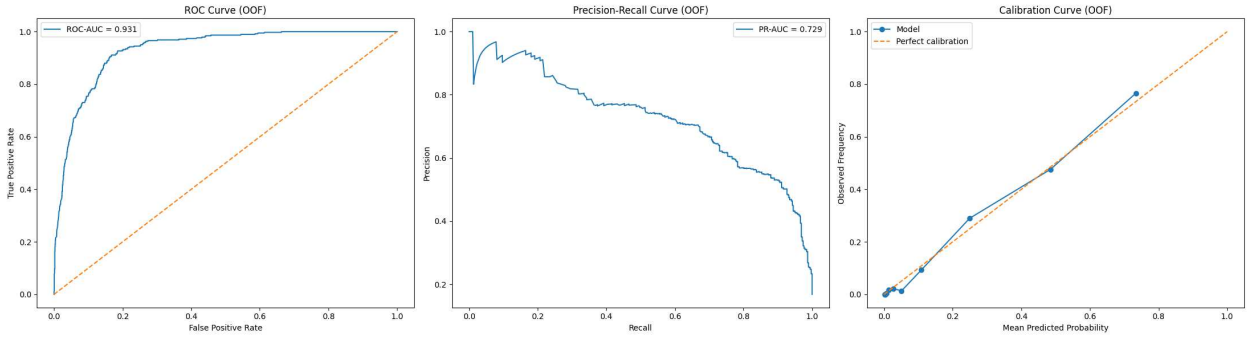


Figure 7: Out-of-fold diagnostic plots: ROC curve, precision–recall curve, and calibration

Figure 7 confirms this interpretation. The ROC curve is steep, the precision–recall curve remains strong across a wide range of recall values, and the calibration profile stays close to the 45-degree line over most of the support. That matters in a literature where many risk claims are difficult to evaluate because their predictive discipline is unclear. The model here is not simply descriptive; it discriminates well and produces reasonably calibrated probabilities.

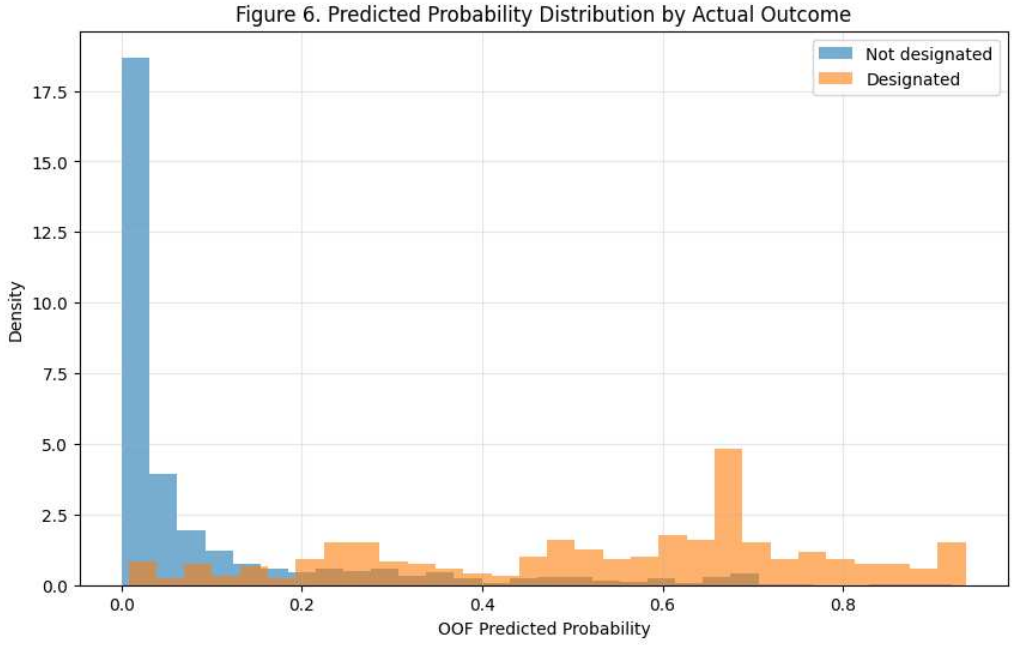


Figure 8: Predicted probability distribution by actual outcome

Figure 8 shows marked separation between designated and non-designated vessels, although overlap remains. That overlap is analytically important because it mirrors the ambiguity highlighted throughout the sanctions literature: shadow-fleet participation is not reducible to one signal alone, and no single model can collapse uncertainty into certainty (Domballe et al., 2024; Hilgenstock et al., 2024; Nainani, 2025). The classifier is therefore best understood as a triage device rather than a deterministic designation rule.

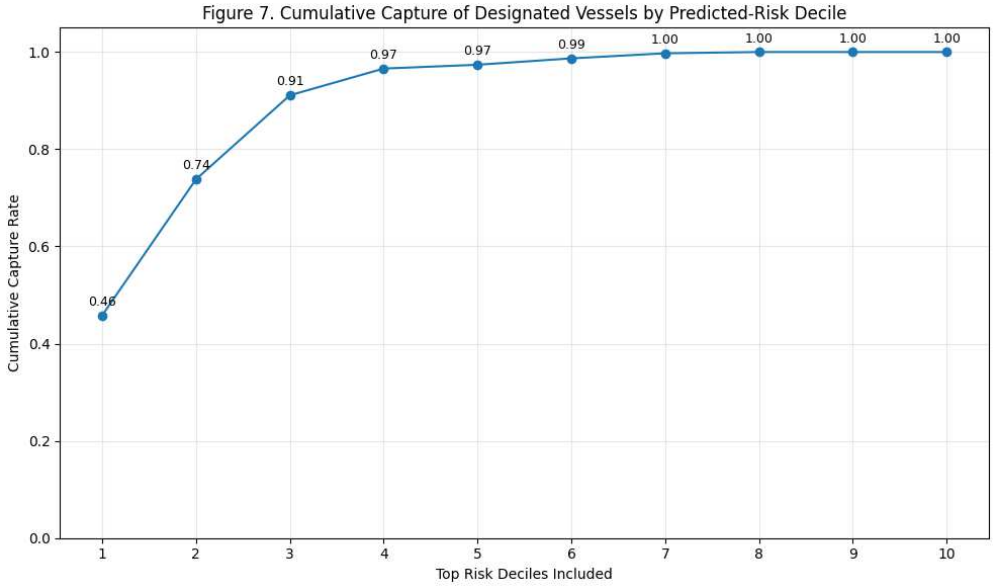


Figure 9: Cumulative capture of designated vessels by predicted-risk decile

Figure 9 translates predictive power into a decision-relevant form. The top risk decile captures 46% of all positives, the top two deciles capture 74%, the top three capture 91%, and the top four capture 97%. This indicates that a relatively narrow watchlist can recover most of the designated subset, which is precisely the type of leverage implied, though not always quantified, in sanctions and maritime-risk analyses (Hilgenstock et al., 2024; Eurasia Institute, 2025; Jones, 2025).

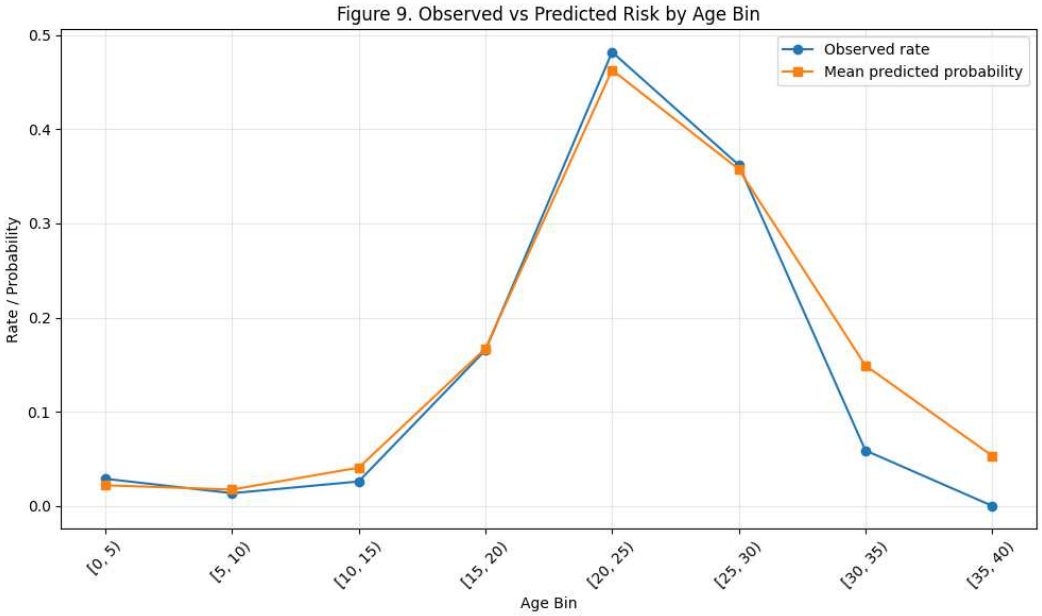


Figure 10: Observed versus predicted risk by age bin

Figure 10 shows that the screening model tracks the age profile of risk closely in the central range where most of the empirical mass lies. In the oldest bins, predicted risk modestly exceeds observed risk, which is unsurprising given thin support in the upper tail. This again reinforces the substantive point that the model is capturing the same age-structured risk pattern suggested by the safety and environmental literature, but in a more formalised way (Knapp et al., 2025; Gözübenli, 2025; Polayeva & Pustovoit-Dnikov, 2024).



Figure 11: Observed designation-rate heatmap by class and fuel

Figure 11 shows that raw risk is most concentrated in the Aframax–VLS IFO cell. This is descriptively informative, but the earlier regression results caution against treating it as a standalone structural conclusion. Much of the apparent fuel pattern overlaps with class and age composition, and the literature offers stronger support for class-age risk than for fuel-specific shadow-fleet mechanisms (Rodriguez-Diaz et al., 2025; Knapp et al., 2025).

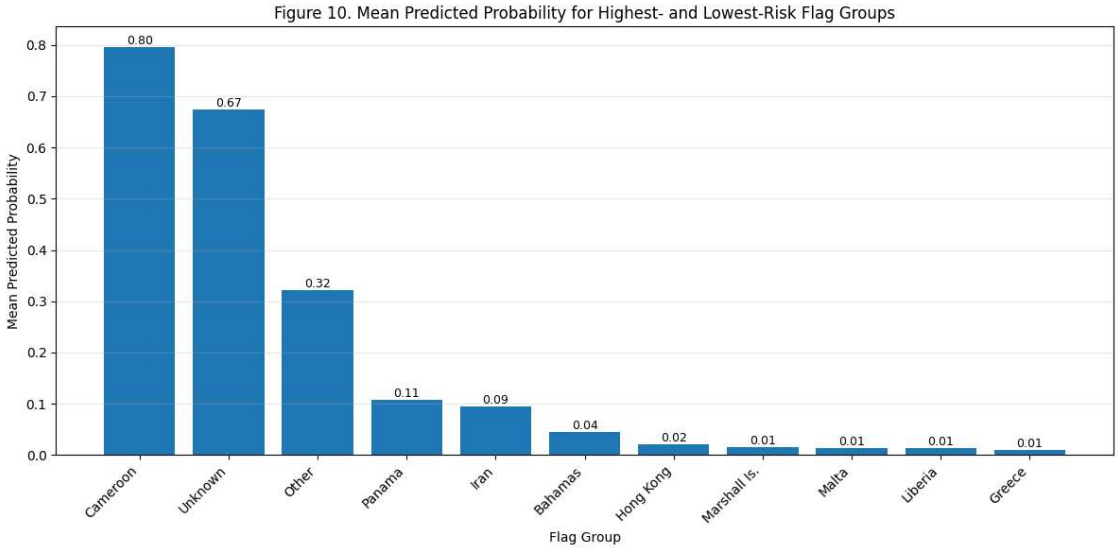


Figure 12: Mean predicted probability for highest- and lowest-risk flag groups

Figure 12 shows that pooled flag state contributes heavily to classification performance. Mean predicted probability is especially high for Cameroon and Unknown and especially low for major registries such as Greece, Liberia, Malta, and the Marshall Islands. This fits the literature on permissive registries and administrative arbitrage, which treats flag selection as part of the shadow-fleet toolkit rather than as a neutral descriptor (Domballe et al., 2024; Eun-Ju Lee, 2025; Nainani, 2025). It also explains why the predictive model outperforms the technical-only inference model: it combines durable vessel characteristics with administrative geography.

Permutation importance supports that reading. Pooled flag state has the largest average importance score (0.3660), followed by age (0.2294) and class (0.1352), while simplified fuel (0.0168) and shuttle-tanker status (0.0020) contribute much less. This ranking is highly informative when placed against existing literature. Studies centred on AIS anomalies, dark activity, and legal opacity emphasise mutable operational markers (Androjna et al., 2023; Dutt et al., 2024; Nainani, 2025); the present model shows that once one moves to prediction, the most useful features are a combination of administrative context and durable vessel structure. That is a more disciplined way of expressing a point long present in policy reports but rarely quantified.

Threshold choice should consequently be understood as a policy trade-off rather than a purely statistical fact. Around 0.18, recall exceeds 0.90, but precision is only moderate and the watchlist remains broad. At 0.25, precision improves and F1 is slightly stronger, making it a plausible triage threshold. At 0.50, false positives fall sharply, but many positives are

missed. This is entirely consistent with the enforcement dilemmas described in the sanctions literature: broad monitoring captures more suspicious activity but imposes higher review costs, whereas tighter escalation thresholds conserve capacity at the risk of omission (Eurasia Institute, 2025; Jones, 2025; Domballe et al., 2024).

4.3 Regime-Specific Divergence

The composite label was then disaggregated into EU, UK, and Ukrainian outcomes. Table 4 summarises the main odds-ratio terms.

Table 4: Regime-specific technical designation models

Variable	EU OR (95% CI)	Ukraine OR (95% CI)	UK OR (95% CI)
Suezmax (vs Aframax)	0.422 (0.293 to 0.608)	0.352 (0.251 to 0.494)	0.418 (0.294 to 0.593)
VLCC (vs Aframax)	0.016 (0.006 to 0.043)	0.124 (0.087 to 0.179)	0.022 (0.010 to 0.047)
Other/Alternative fuel (vs IFO 380)	8.039 (2.582 to 25.036)	5.081 (1.755 to 14.708)	9.684 (2.919 to 32.128)
VLS IFO (vs IFO 380)	0.861 (0.549 to 1.352)	0.659 (0.466 to 0.934)	0.786 (0.513 to 1.205)
Shuttle tanker	0.510 (0.163 to 1.591)	0.728 (0.313 to 1.697)	0.471 (0.154 to 1.442)

Notes: Each model includes the four-degree-of-freedom age spline used in the baseline technical specification. Full odds
including spline basis terms, are reported in the Appendix. The US-specific model is not estimated because only 13 positives

The class pattern is again highly consistent. Relative to Aframax vessels, both Suezmaxes and VLCCs are much less likely to appear on the EU, UK, and Ukrainian lists. The VLCC penalty is especially severe in the EU and UK models. The Ukrainian model also disfavors VLCCs strongly, though less dramatically than the EU and UK. These findings suggest a broad convergence across sanctioning regimes on smaller crude tanker classes, even if full harmonisation is absent. That inference is broadly compatible with the literature describing a shared but incomplete international response to Russian shadow shipping (Hilgenstock et al., 2024; Domballe et al., 2024; Eun-Ju Lee, 2025; Jones, 2025).

Fuel effects are more heterogeneous. The Other/Alternative category is positive across all three regimes, while VLS IFO is significantly negative only in the Ukrainian model. Shuttle-tanker status remains negative and insignificant throughout. The age spline terms remain important in all three regime-specific models, indicating that the non-linear age effect is not simply an artefact of pooling all regimes together. In substantive terms, the regimes appear to share a broad targeting profile centred on older, smaller crude tankers, but with some divergence in how strongly they weight fuel-related characteristics and ageing structure.

That regime-level divergence is consistent with the institutional and legal heterogeneity

described in the literature. The shadow-fleet problem is transnational, but enforcement is mediated through different legal frameworks, intelligence pipelines, political incentives, and evidentiary thresholds (Eun-Ju Lee, 2025; Ibadoghlu, 2025; Jones, 2025). The present results do not imply that the EU, UK, and Ukraine operate identical sanction logics. They do, however, suggest that a common vessel-centred risk template underlies much of the observed designation practice. This is an important contribution because the literature has so far described that convergence more often than it has measured it.

5 Policy Implications

The empirical results suggest that sanctions enforcement in maritime oil transport would benefit from moving away from a largely reactive, list-driven posture towards a more structured risk architecture grounded in vessel characteristics. Much of the existing policy debate has concentrated on opacity in ownership, insurance, and routing, and with good reason: shell-company layering, permissive flags, AIS interference, and deceptive ship-to-ship activity are now well-established features of sanctions circumvention (Androjna et al., 2023; Dutt et al., 2024; Domballe et al., 2024; Nainani, 2025). Yet the present model shows that a substantial part of the observed designation universe is not random. It clusters within a recognisable material subset of the crude tanker fleet, above all in older Aframax tonnage. That finding matters because vessel hardware is far more durable than many of the administrative and digital markers upon which compliance systems currently rely.

For regulators and port-state authorities, this creates a practical basis for a layered screening regime. The technical model demonstrates that class and age are doing most of the explanatory work, while the predictive model shows that these structural characteristics can be translated into operational triage with meaningful out-of-fold performance. A screening architecture that begins with immutable or slow-moving vessel attributes can therefore reduce the search space before investigators turn to more resource-intensive inquiries into beneficial ownership, insurance provenance, voyage anomalies, or transactional counterparties. In institutional terms, this is attractive because it economises on scarce investigative capacity while remaining legally more defensible than acting on opaque intelligence alone. In contexts where data are incomplete, deliberately manipulated, or fragmented across jurisdictions, a vessel-centred first filter is likely to be more resilient than any system built principally on mutable metadata (Hilgenstock et al., 2024; Eurasia Institute, 2025; Rodriguez-Diaz et al., 2025).

The threshold analysis makes clear that the model is especially useful as a triage instrument.

At an out-of-fold threshold around 0.18, recall exceeds 0.90, which means the screening tool captures most vessels in the observed designation universe. The cost is a fairly broad watchlist and only moderate precision. That trade-off is not a defect so much as a policy choice. Where the objective is surveillance breadth, secondary documentation review, or prioritisation for enhanced due diligence, a lower threshold is entirely defensible. A threshold around 0.25 produces a narrower and more operationally manageable list, with stronger precision and the best balance between false negatives and false positives in the reported results. A threshold around 0.50 is better suited to circumstances in which enforcement capacity is highly constrained or the downstream consequences of escalation are particularly severe. In other words, the model supports differentiated use rather than a single cut-off. Maritime compliance teams, port-state inspectors, insurers, classification societies, and financial institutions do not face identical optimisation problems, and they should not be expected to use identical thresholds.

This has immediate implications for sanctions implementation in the financial and commercial sectors. Banks, commodity traders, P&I clubs, and shipbrokers increasingly operate in an environment where traditional know-your-customer and sanctions-screening procedures struggle to keep pace with shipping-specific forms of obfuscation (Domballe et al., 2024; Jones, 2025; Ibadoghlu, 2025). The model developed here provides a way of supplementing entity-based controls with asset-based risk scoring. A vessel that is not yet listed but sits squarely in the high-risk technical profile identified by the model may warrant enhanced due diligence, deeper ownership verification, closer scrutiny of bills of lading and chartering arrangements, or the imposition of additional internal approvals before financing or cover is extended. This is particularly important because the literature increasingly suggests that sanctions circumvention works not merely through hidden actors, but through the gradual assembly of intermediary service ecosystems capable of sustaining trade under pressure (Cardoso et al., 2025; Wang et al., 2025; Omokaro et al., 2025). A vessel-centred risk score can therefore function as an early-warning device within a wider compliance workflow.

The predictive model's strong dependence on pooled flag state, however, should encourage caution as much as confidence. Flag state is the most important predictor in the operational specification, exceeding the contribution of age and class in permutation importance. That result is highly informative from an enforcement standpoint, but it should not be mistaken for proof that flag state is a deep structural cause of illicit behaviour. The literature on shadow fleets repeatedly shows that registries are embedded in opportunistic governance, weak oversight, and legal arbitrage (Eun-Ju Lee, 2025; Nainani, 2025; Domballe et al., 2024). In this sense, the model is learning a mixture of vessel structure and administrative geography. That

is valuable for screening, but it carries legal and policy implications. A regulator or private actor relying on the predictive model should be careful not to treat registry-enhanced probabilities as conclusive evidence of wrongdoing. They are better understood as indicators of elevated designation risk, to be combined with corroborative evidence concerning ownership opacity, transactional behaviour, or voyage patterns. Used properly, the model strengthens proportionality by directing attention; used carelessly, it could encourage overreach against vessels whose principal risk marker is jurisdictional rather than operational.

The regime-specific results also speak to the coordination problem that has shaped sanctions enforcement since 2022. Despite inevitable legal and political differences, the EU, UK, and Ukraine appear to converge on a broadly similar technical profile. In all three models, designation risk is concentrated in older, smaller crude tanker segments, and Aframax tonnage is consistently more exposed than Suezmax or VLCC tonnage. This convergence is important. It suggests that a common vessel-risk grammar may already exist beneath the surface of ostensibly separate sanctioning systems. That creates an opening for more systematic intelligence-sharing, for common risk indicators, and for partial harmonisation of technical screening criteria across authorities. At the same time, the models also reveal that the regimes do not weight every characteristic identically. The Ukrainian model, for instance, appears somewhat more sensitive to certain age and fuel patterns than the EU or UK models. That divergence is unsurprising in light of the different institutional mandates, evidentiary thresholds, and strategic priorities at work across sanctioning bodies (Hilgenstock et al., 2024; Jones, 2025; Eun-Ju Lee, 2025). It does, however, imply that multilateral coordination is still incomplete. A shared vessel-centred typology could reduce informational fragmentation without requiring absolute legal uniformity.

Beyond sanctions compliance, the findings reinforce the close relationship between shadow-fleet enforcement and maritime safety governance. The concentration of observed designation risk in ageing tanker tonnage is directly relevant to the environmental and casualty concerns raised in the emerging literature on shadow fleets (Knapp et al., 2025; Gözübenli, 2025; Polayeva & Pustovoit-Dnikov, 2024; Rodriguez-Diaz et al., 2025). Older vessels operating in opaque commercial conditions, often under fragmented oversight, create overlapping risks: sanctions evasion, insurance uncertainty, pollution exposure, and accident vulnerability. This means that the public value of risk-based screening is wider than sanctions policy alone. A port-state authority that targets vessels because they fit a high-risk technical profile may simultaneously improve compliance enforcement and reduce exposure to marine incidents. The same logic applies to coastal states and regional monitoring frameworks in ecologically sensitive waters. In that respect, the model's strongest contribution may lie in helping

policymakers see that the shadow fleet is not only a problem of hidden ownership and evasion, but also a problem of ageing maritime infrastructure operating at the margins of effective governance.

The broader strategic literature points in the same direction. Work on hybrid warfare, seabed security, and maritime grey zones has shown that shipping can no longer be understood as a neutral commercial backdrop to geopolitical conflict (Kozmenko, 2024; Antonova & Karatash, 2025; Meiertöns, 2025). The shadow fleet exemplifies this reality. It is a commercial transport system, but also an instrument of fiscal resilience, strategic ambiguity, and coercive adaptation. A vessel-centred model cannot solve that problem by itself. What it can do is impose more structure on a policy space that is often dominated by anecdote, fragmented intelligence, or post hoc designation practice. The findings therefore support a more disciplined enforcement framework: one in which vessel characteristics provide the initial screen, administrative and behavioural indicators provide the refinement, and formal sanctions action follows only where the evidentiary picture becomes sufficiently strong.

Taken together, the policy message is straightforward. The model should not be used as an automated designation engine, nor should it be read as direct proof of latent shadow-fleet membership. Its proper function is narrower, but still powerful. It helps identify where scrutiny is most likely to pay off, where compliance resources should be concentrated, and where maritime governance and sanctions enforcement overlap most closely. In a domain defined by adaptation, opacity, and legal fragmentation, that is already a substantial gain.

6 Conclusion

This article has examined whether sanction-designated crude tankers can be distinguished from the wider global tanker fleet through vessel-level technical characteristics and an operational screening model. Using a dataset of 2,277 crude tankers and a composite designation indicator across the EU, UK, Ukraine, and the United States, the paper shows that observed designation risk is far from random.

Three conclusions are central. First, designation risk is concentrated in Aframax tonnage. Relative to Aframax vessels, both Suezmaxes and VLCCs have markedly lower odds of appearing in the designated subset across baseline, clustered, probit, and regime-specific models. Secondly, the age effect is strongly non-linear: risk is concentrated in an intermediate-to-old ageing window, especially in the late teens through the mid-twenties, rather than increasing monotonically across the entire age distribution. Thirdly, fuel effects are secondary and

somewhat unstable, while shuttle-tanker status contributes little once the broader technical profile of the vessel is controlled for.

The operational screening model performs strongly. With pooled flag information added, it achieves ROC-AUC 0.9306, PR-AUC 0.7288, and a Brier score of 0.0750. A threshold around 0.18 captures just over 90% of positives, although at the cost of moderate precision. This makes the model useful for triage, prioritisation, and risk-based escalation. Yet it also clarifies an important conceptual limit. The strongest predictive model partly learns the administrative geography of observed sanctioning practice. It is therefore more defensibly interpreted as a designation-risk classifier than as direct proof of latent shadow-fleet membership.

More broadly, the article supports a shift in emphasis within the sanctions-evasion literature. Administrative opacity and digital deception remain vital, but the physical fleet also matters. Vessel structure, ageing, and class composition provide durable empirical leverage in a domain characterised by fluid identities and evasive behaviour (Domballe et al., 2024; Hilgenstock et al., 2024; Nainani, 2025). For maritime governance, that insight is analytically valuable and practically actionable.

AI Disclosure

Portions of the writing process were assisted by ChatGPT (OpenAI) for language editing, formatting suggestions, and stylistic improvements. All arguments, analysis, interpretations, and conclusions are solely those of the authors, who reviewed and verified the final text.

Appendix: Full Regression Tables

Table 5: Full baseline odds-ratio table: logit with HC3-robust standard errors

Variable	**Odds ratio**	**95% CI**	p**-value**
Constant	0.080	0.022 to 0.292	0.0001
Suezmax (vs Aframax)	0.340	0.245 to 0.471	< 0.001
VLCC (vs Aframax)	0.143	0.100 to 0.206	< 0.001
Other/Alternative fuel (vs IFO 380)	4.850	1.723 to 13.652	0.0028
VLS IFO (vs IFO 380)	0.761	0.541 to 1.072	0.1184
Age spline basis 1	0.070	0.006 to 0.825	0.0346
Age spline basis 2	6.290	0.990 to 39.967	0.0513
Age spline basis 3	720.233	50.354 to 10301.831	< 0.001
Age spline basis 4	0.034	0.000 to 3.919	0.1622
Shuttle tanker	0.644	0.279 to 1.487	0.3024

Table 6: Full baseline odds-ratio table: logit with builder-clustered standard errors

Variable	**Odds ratio**	**95% CI**	p**-value**
Constant	0.080	0.010 to 0.635	0.0168
Suezmax (vs Aframax)	0.340	0.215 to 0.537	< 0.001
VLCC (vs Aframax)	0.143	0.090 to 0.228	< 0.001
Other/Alternative fuel (vs IFO 380)	4.850	0.906 to 25.952	0.0650
VLS IFO (vs IFO 380)	0.761	0.484 to 1.198	0.2385
Age spline basis 1	0.070	0.002 to 2.256	0.1332
Age spline basis 2	6.290	0.617 to 64.162	0.1207
Age spline basis 3	720.233	60.131 to 8626.742	< 0.001
Age spline basis 4	0.034	0.000 to 4.806	0.1801
Shuttle tanker	0.644	0.327 to 1.269	0.2032

Table 7: Full regime-specific odds-ratio table: European Union

Variable	**Odds ratio**	**95% CI**	p**-value**
Constant	0.044	0.010 to 0.193	< 0.001
Suezmax (vs Aframax)	0.422	0.293 to 0.608	< 0.001
VLCC (vs Aframax)	0.016	0.006 to 0.043	< 0.001
Other/Alternative fuel (vs IFO 380)	8.039	2.582 to 25.036	0.0003
VLS IFO (vs IFO 380)	0.861	0.549 to 1.352	0.5164
Age spline basis 1	0.253	0.013 to 4.872	0.3622
Age spline basis 2	1.119	0.082 to 15.192	0.9328
Age spline basis 3	3303.390	68.340 to 159678.616	< 0.001
Age spline basis 4	0.000	0.000 to 0.112	0.0180
Shuttle tanker	0.510	0.163 to 1.591	0.2461

Table 8: Full regime-specific odds-ratio table: Ukraine

Variable	**Odds ratio**	**95% CI**	p**-value**
Constant	0.081	0.024 to 0.277	0.0001
Suezmax (vs Aframax)	0.352	0.251 to 0.494	< 0.001
VLCC (vs Aframax)	0.124	0.087 to 0.179	< 0.001
Other/Alternative fuel (vs IFO 380)	5.081	1.755 to 14.708	0.0027
VLS IFO (vs IFO 380)	0.659	0.466 to 0.934	0.0190
Age spline basis 1	0.081	0.007 to 0.983	0.0484
Age spline basis 2	2.228	0.320 to 15.508	0.4183
Age spline basis 3	3610.209	251.238 to 51877.620	< 0.001
Age spline basis 4	0.001	0.000 to 0.152	0.0080
Shuttle tanker	0.728	0.313 to 1.697	0.4628

Table 9: Full regime-specific odds-ratio table: United Kingdom

Variable	**Odds ratio**	**95% CI**	p**-value**
Constant	0.073	0.019 to 0.272	0.0001
Suezmax (vs Aframax)	0.418	0.294 to 0.593	< 0.001
VLCC (vs Aframax)	0.022	0.010 to 0.047	< 0.001
Other/Alternative fuel (vs IFO 380)	9.684	2.919 to 32.128	0.0002
VLS IFO (vs IFO 380)	0.786	0.513 to 1.205	0.2694
Age spline basis 1	0.047	0.003 to 0.743	0.0299
Age spline basis 2	3.184	0.230 to 44.152	0.3880
Age spline basis 3	1005.524	33.311 to 30352.450	0.0001
Age spline basis 4	0.000	0.000 to 0.429	0.0318
Shuttle tanker	0.471	0.154 to 1.442	0.1876

References

Androjna, A., Pavić, I., Gucma, L., Vidmar, P., & Perkovič, M. (2023). AIS data manipulation in the illicit global oil trade. *Journal of Marine Science and Engineering*, 12(1), 6.

Antonova, M. L., & Karatash, S. (2025). The strategic significance of the maritime theatre in the Russia–Ukraine war.

Astrov, V., Brockhaus, C., Hinz, J., Jessen-Thiesen, L., Mahlkow, H., & Svab, P. (2024). *Russia Monitor 3, Navigating Trade Restrictions*. The Vienna Institute for International Economic Studies (wiiw).

Cardoso, D. S., Salant, S. W., Daubanes, J., & Daubanes, J. X. (2025). The dynamics of evasion: The price cap on Russian oil exports and the amassing of the shadow fleet. *CESifo Working Paper*.

Domballe, J., Esposito, M., Lin, M., & McKinney, B. (2024). *Russia's shadow fleet – Formation, operation and continued risks for sanctions compliance teams*. S&P Global Commodity Insights.

Dutt, A., Mukherjee, A., Panayotov, G., Roy, D., & Wen, X. (2024). *Sanctions, Satellites, Suspect Ships*. SSRN.

Eun-Ju Lee. (2025). The shadow fleet tankers: Threats and international legal responses. *Wonkwang Law Review*, 41(1), 111–141.

Eurasia Institute. (2025). *Risk Exposure and Shadow Fleets: 2025 Report*.

Götz, R. (2023). War economy: Russia's economic development in 2022. *Open Access Journal of the Center for Governance and Culture in Europe*, University of St. Gallen, 13(35).

Gözübenli, A. (2025). *The Shadow Fleet: Navigating the Nexus of Ecological Peril and Geopolitical Strife*.

Grebennikov, I. V. (2024). Impact of sanctions on Russia's oil trade. *Russian Foreign Economic Journal*, 4, 121–130.

Hilgenstock, B., Hrybanovskii, O., & Kravtsev, A. (2024). *Assessing Russia's Shadow Fleet: Initial Build-Up, Links to the Global Shadow Fleet, and Future Prospects*. KSE Institute.

Ibadoghlu, G. (2025). *Azerbaijan in the Gray Zone: Shadow Energy Trade and Russian Sanctions Circumvention*.

Johnson, S., Rachel, L., & Wolfram, C. (2023). *A Theory of Price Caps on Non-Renewable Resources*. NBER Working Paper No. 31347.

Jones, S. (2025). *Russia's Shadow War Against the West*. Center for Strategic and International Studies.

Kilian, L., Rapson, D., & Schipper, B. C. (2024). *The Impact of the 2022 Oil Embargo and Price Cap on Russian Oil Prices*. FRB of Dallas Working Paper No. 2401.

Knapp, S., Rydbergh, T., Franses, P. H., & Berneblad, B. (2025). *Maritime Risk Assessment for Portugal, Sweden and the Baltic Sea with an emphasis on shadow fleets*. Econometric Institute – ERIM Report Series.

Kozmenko, A. S. (2024/2025). Shadow oil fleet of Russia: A new chapter of hybrid warfare. *Hybrid War Studies Journal*.

Madaki, A. G. (2025). *Economic sanctions: Impact, pressure and geopolitical consequences*.

Malykh, A. (2025). Rossiya na mirovom energeticheskom rynke [Russia in the global energy market]. *Russian Journal of Energy Policy*.

Meiertöns, H. (2025). *Seabed Warfare: Protection of Submarine Cables, Grey Zone Aggression and Sea-Blindness*.

Nainani, D. (2025). The unfree sea: a legal geography of the dark fleet. *Journal of Property, Planning and Environmental Law*.

Omokaro, G. O., Nafula, Z. S., Iloabuchi, N. E., Efeni, O. S., Adeyanju, O. I., Janet, O. O., & Idiong, O. U. (2025). Energy sanctions in the global economy: Geopolitical disruptions, market fragmentation, innovation and green transition. *International Journal of Innovation Studies*.

Polayeva, B., & Pustovoit-Dnikov, G. (2024). Risks of maritime hydrocarbon transport under intensifying anti-Russian sanctions. *Ekonomika Stroitelstva*, Russian Academy of Sciences.

Poussenkova, N. (2024). Survival of the fleetest: Russian oil companies under Western sanctions. *Russian Analytical Digest*, No. 308.

Rodriguez-Diaz, E., Alcaide, J. I., & Endrina, N. (2025). Shadow fleets: A growing challenge in global maritime commerce. *Applied Sciences*, 15(12), 6424.

Wang, X. et al. (2025). Sanctions evasion in global oil flows: Evidence from shipping activity. *Journal of Asian Economics*.

Yusuf, T. (2024). *US Sanctions Are Ineffective: Russia's Dark Fleet and Gray Fleet and its Circumvention of Sanctions*. MPRA Paper No. 121829.