Statistical Attribution of In-Service Specific Fuel Oil Consumption Deviation to Bunker Fuel Combustion Quality:

A Regression-Based Decomposition Framework for Marine Diesel Engines

Tanmoy Dey

Marine Engineer, Marine Engineering and Research Institute (MERI) Kolkata

Corresponding author: tanmoymarine96@gmail.com

Abstract

Specific fuel oil consumption (SFOC) deviation in service is routinely attributed, without formal separation, to some combination of bunker fuel quality, engine mechanical condition, ambient operating conditions, and load profile. Existing correction practice — net calorific value (NCV) normalisation to an ISO reference and ambient correction under ISO 3046-1/15550 — removes the deterministic component of fuel-quality effects but leaves the combustion-quality component (incomplete combustion, ignition delay, afterburning associated with poor ignition characteristics) unseparated from genuine engine degradation. This paper proposes a regression-based framework that reconstructs bunker-batch boundaries from tank remaining-on-board (ROB) records, applies deterministic NCV correction, and then estimates a fixed-effects model in which each bunker batch enters as a categorical variable alongside engine load and a continuous time-drift term representing mechanical degradation. The batch coefficients recovered by this model are interpreted as the combustion-quality penalty of each bunker batch, net of energy content, load, and engine wear, with confidence intervals derived from heteroskedasticity- and autocorrelation-consistent standard errors. Three identification threats — confounding with maintenance events, confounding with load profile, and confounding with ambient conditions — are formalised as automatic diagnostic checks, and a placebo test using randomised batch boundaries is proposed to guard against spurious attribution. The framework is demonstrated on a synthetic twelve-month operating dataset for a representative handysize bulk carrier with injected ground-truth effects; the batch with the largest injected effect is recovered within its estimated confidence interval, while smaller and sequentially later batches show attenuated identification consistent with a formally diagnosed collinearity between batch timing and the continuous degradation term — a limitation the paper reports explicitly rather than obscures. The paper situates this contribution against the existing literature, which addresses fuel ignition quality (CCAI, CII, EFN) and in-service SFOC monitoring as separate bodies of work; to the authors' knowledge, no published methodology jointly estimates bunker-batch-specific combustion penalties from routine in-service reporting data under an explicit identification framework. Implications for bunker procurement, fuel quality dispute resolution, and vessel performance monitoring practice are discussed, along with the method's data and instrumentation requirements and its limitations.

Keywords: *specific fuel oil consumption; marine diesel engine; bunker fuel quality; fixed-effects regression; fuel attribution; ship performance monitoring; CCAI; ignition quality*

1. Introduction

1.1 Background

Marine fuel typically represents 50–60% of voyage operating expenditure for a non-scrubber-fitted commercial vessel (Stopford, 2009), making specific fuel oil consumption (SFOC) one of the most closely monitored technical parameters in ship operation. SFOC — the mass of fuel consumed per unit of shaft work delivered, expressed in g/kWh — is guaranteed by engine manufacturers at ISO 3046-1 reference conditions and is expected to track a known curve against engine load under those reference conditions. In service, however, SFOC deviates from this reference curve for reasons that are rarely separated in routine monitoring practice: fuel energy content, fuel combustion quality, engine mechanical condition, and ambient environmental conditions all move the observed figure simultaneously, and the deviation is customarily reported as a single number and interpreted, often informally, as an indicator of "engine condition."

This conflation has practical consequences. When a vessel's SFOC rises, the response is typically directed at the engine — increased monitoring, earlier turbocharger washing, or a maintenance investigation — even where the dominant cause is the fuel itself. Conversely, a genuine mechanical degradation trend can be masked for months if it coincides with a delivery of unusually high-quality fuel. Because bunker fuel is purchased in discrete batches from varying suppliers, terminals, and blend stocks, and because ISO 8217 permits meaningful within-grade variation in properties that affect combustion behaviour, batch-to-batch differences in "how well a fuel burns" — independent of its energy content — are physically expected, yet are not, to the authors' knowledge, routinely estimated as a discrete, attributable quantity in either commercial ship performance monitoring practice or the published literature.

1.2 The gap this paper addresses

Three adjacent bodies of practice and literature currently exist without a bridge between them.

1. Fuel-side ignition quality characterisation. Work following Zeelenberg, van Draat and Barker's foundational CIMAC studies on the ignition performance of residual marine fuels established the Calculated Carbon Aromaticity Index (CCAI) as a practical, density-and-viscosity-derived proxy for ignition delay, later refined with corrections for microcarbon residue and vapour aromaticity. This literature characterises fuel properties and their laboratory or test-rig combustion behaviour but does not connect batch-level fuel indices to measured in-service engine performance across a fleet.

2. Engine-side SFOC correction and condition monitoring. Standard practice, codified in ISO 3046-1 and ISO 15550 and implemented in OEM project guides, corrects observed SFOC to a reference calorific value and reference ambient conditions, and trends the residual against a shop-trial curve to flag engine degradation, typically alarming at a 2–3% sustained deviation. This practice explicitly treats the corrected residual as an engine-condition signal and does not further decompose it by fuel batch.

3. Data-driven fuel consumption modelling. Recent machine-learning approaches model SFOC as a function of operating conditions with the objective of predictive accuracy, but do not treat bunker batch identity as an explanatory variable of interest, and do not report batch-level effect estimates with uncertainty quantification.

No published work, to the authors' knowledge, treats bunker batch as a categorical regressor within a controlled model of in-service SFOC, estimates its coefficient net of calorific content, load, ambient conditions and engine degradation, and reports the result as an attributable combustion-quality penalty with a formal identification argument. This is the gap the present paper addresses.

1.3 Contribution

This paper makes four contributions: (1) a tank-reconciliation algorithm that maps calendar days to the bunker batch actually being combusted, using ROB depletion logic; (2) a fixed-effects regression specification that jointly estimates engine load response, continuous mechanical degradation, discrete maintenance step-effects, and batch-specific combustion-quality penalties from routine noon-report and monthly performance-report data; (3) a set of formal identification diagnostics — confounding tests against maintenance timing, load profile, and ambient conditions, plus a placebo randomisation test — that determine when the batch effect is causally interpretable and when it is not; and (4) a demonstration on a synthetic dataset with known injected ground truth, reported honestly including the cases where identification is only partial.

1.4 Paper structure

Section 2 reviews the relevant literature across fuel quality characterisation, engine performance correction standards, and data-driven consumption modelling. Section 3 presents the methodology. Section 4 demonstrates the method on a synthetic dataset. Section 5 discusses practical implications, data and instrumentation requirements, and limitations. Section 6 concludes.

2. Literature Review

2.1 Fuel ignition quality and combustion behaviour

The Calculated Carbon Aromaticity Index (CCAI) remains the most widely used practical ignition-quality index for residual marine fuels, computed from density and viscosity (Zeelenberg et al., 1983). Its physical basis is the correlation between aromaticity — which the standard cetane number test cannot directly assess for residual fuels — and ignition delay under engine-relevant conditions (Brandt and Zeelenberg, 1988). Subsequent work has refined the index by correcting for microcarbon residue as a proxy for the non-vaporisable aromatic fraction (Wortel, 2005). A consistent finding across this literature is that no single laboratory-derived index is fully representative of actual marine engine combustion behaviour across all load and design conditions (Baert, 1988), which motivates empirical, in-service estimation as a complement to laboratory characterisation rather than a substitute for it.

2.2 SFOC correction standards and condition monitoring practice

ISO 3046-1 and ISO 15550 define the reference conditions against which engine manufacturers guarantee SFOC, and specify correction formulae for ambient deviation. OEM project guides implement these corrections and present shop-trial SFOC-versus-load curves as the baseline against which in-service performance is judged. Industry technical guidance consistently identifies a sustained 2–3% corrected SFOC deviation as the conventional threshold for initiating a maintenance investigation, implicitly treating the corrected residual as an engine-condition signal. This practice is standard and well-founded for its intended purpose, but by construction leaves any batch-to-batch combustion-quality variation embedded in that same residual, indistinguishable from genuine mechanical wear, unless further decomposed.

2.3 Data-driven fuel consumption modelling

A growing body of work applies statistical and machine-learning methods to voyage-level fuel consumption prediction, including ensemble Gaussian process regression approaches applied to real ship voyage data and design-of-experiments methods mapping SFOC–NOx trade-offs as functions of combustion parameters, generally on test-bed rather than fleet in-service data. This literature demonstrates strong methodological capability for consumption prediction but does not treat fuel batch identity as a covariate of causal interest, and none of the reviewed work reports a batch-specific effect estimate with an accompanying identification argument.

2.4 Ship efficiency decomposition more broadly

Adjacent literature on hull and propeller performance monitoring — notably the ISO 19030 standard for measuring changes in hull and propeller performance — establishes a template for decomposing an observed performance metric into weather-driven and condition-driven components using regression against sea-state variables and step-change detection at maintenance events. The present paper adopts an analogous logic, applied instead to SFOC and bunker batch identity, a domain where, to the authors' knowledge, this decomposition has not previously been formalised.

2.5 Summary of the gap

Literature	Object of study	Batch as causal covariate?	In-service attribution with uncertainty?
Fuel ignition quality (CCAI, CII, EFN)	Lab/test-rig fuel properties	N/A	No
SFOC correction standards (ISO 3046-1/15550)	Ambient/load-corrected SFOC vs shop trial	No	No — flags deviation, does not attribute it
Data-driven consumption modelling	Predictive SFOC/consumption models	No	Not with batch-level attribution
Hull/propeller decomposition (ISO 19030)	Speed loss / added resistance	N/A (different metric)	Yes, for a different quantity
This paper	In-service SFOC	Yes — batch as fixed effect	Yes — coefficient + CI + identification diagnostics

3. Methodology

3.1 Data requirements

The method requires only data already generated in routine ship operation. Daily (noon report): date; main engine (ME) fuel consumed (tonnes); ME shaft power (kW) or, absent a shaft power meter, RPM for a calibrated propeller-law estimate; ME running hours; Beaufort sea state (for filtering); fuel grade in use. Per bunker stem (bunker delivery note and independent laboratory analysis, e.g. VPS or FOBAS): stem date, port, quantity, tank(s) filled; net calorific value (NCV) or gross calorific value with hydrogen/water content sufficient to derive NCV; density at 15°C and viscosity at 50°C; sulphur, water, ash content; catalytic fines (Al+Si) where available; price. Daily or per-sounding: tank-wise remaining-on-board (ROB) by grade. Event log: dates of turbocharger water washes and overhauls, injector renewals, fuel pump overhauls, and any other maintenance intervention plausibly affecting SFOC.

3.2 Tank reconciliation: mapping days to combusted batch

A bunker stem does not begin being combusted on its delivery date; it begins when the service tank draws from the storage or settling tank into which that stem was loaded, and if tanks are commingled, a given day's combustion may reflect a blend of stems in proportions that change as the tank depletes. For each day d, the tank feeding the service system is identified from the ROB log. If that tank contains fuel from a single stem, the day is assigned to that stem with purity 1.0. If the tank contains fuel from multiple stems, the purity of the dominant stem is computed via a first-in-first-out depletion model calibrated against periodic soundings. Days with dominant-stem purity below 0.85 are flagged as commingled and excluded from the attribution model; days within 48 hours of a grade changeover are excluded outright. Where tank-wise ROB is unavailable, a FIFO depletion assumption applied to the aggregate ROB log is used as a fallback, with all resulting batch assignments marked as estimated and confidence intervals widened accordingly.

3.3 Deterministic correction

Measured SFOC for day d is:

$$SFOC\_meas(d) = [fuel\_tonnes(d) \times 10^6] / [power\_kW(d) \times hours(d)] \quad (1)$$

Where only gross calorific value (GCV) is available, net calorific value (NCV) is estimated following the standard correction:

$$NCV = GCV - 2.442 \times (9H/100 + W/100) \times 1000 \quad (2)$$

where H is hydrogen content (% m/m) and W is water content (% m/m). SFOC is then scaled to the ISO reference calorific value of 42,700 kJ/kg:

$$SFOC\_cv(d) = SFOC\_meas(d) \times [NCV\_batch / 42{,}700] \quad (3)$$

The batch-level calorific value penalty is the mean, over all days assigned to that batch, of SFOC_meas(d) − SFOC_cv(d). This quantity is exact arithmetic subject only to laboratory measurement precision (typically ±0.2%) and is not itself a subject of statistical inference; it typically accounts for the majority of any observed batch-to-batch SFOC difference and must be separated before combustion-quality effects can be isolated. Ambient correction, where scavenge air temperature, barometric pressure and cooling water temperature are available, follows the ISO 3046-1/15550 correction formula using OEM-supplied coefficients. Where ambient data is unavailable, this correction is omitted and the resulting unmodelled variation is absorbed into the regression noise term, widening confidence intervals but not biasing the batch coefficients provided ambient conditions are not systematically correlated with batch timing — a condition tested explicitly in Section 3.5.

3.4 Regression specification

Filtered to steady-state days (running hours ≥ 20; Beaufort ≤ 4; load between 40–95% MCR; single fuel grade; batch purity ≥ 0.85; at least 2 days clear of any maintenance event; SFOC within plausible physical bounds), the model estimated is:

$$SFOC\_cv(d) = \beta_0 + \beta_1 \cdot Load(d) + \beta_2 \cdot Load(d)^2 + \beta_3 \cdot t(d) + \Sigma_j\, \delta_j \cdot Maint\_j(d) + \Sigma_i\, \gamma_i \cdot Batch\_i(d) + \varepsilon(d) \quad (4)$$

where Load(d) is engine load as a percentage of MCR; t(d) is time in months from the start of the observation window, representing continuous mechanical degradation; Maint_j(d) is an indicator equal to 1 for all days after maintenance event j; Batch_i(d) is an indicator equal to 1 for all days assigned to bunker batch i, with one batch omitted as the reference category; and ε(d) is the residual. The coefficient γ_i is interpreted as batch i's combustion-quality penalty, in g/kWh, relative to the reference batch, net of calorific value, engine load, continuous mechanical degradation, and any discrete maintenance step effects. Ordinary least squares (OLS) estimation is used, with standard errors computed via the Newey-West heteroskedasticity- and autocorrelation-consistent estimator (7-day lag), rather than the classical OLS covariance estimator, which is known to understate standard errors under serial correlation of the kind expected in sequential daily ship performance reporting.

3.5 Identification diagnostics

The interpretation of γ_i as a causal combustion-quality effect requires that batch timing not be confounded with other causes of SFOC variation. Four diagnostic checks are run automatically and reported alongside every estimate. Maintenance confounding: if a maintenance event falls within five days of a batch boundary, the corresponding δ_j and γ_i are not separately identified, and the affected batch coefficient is flagged as confounded. Load-profile confounding: where two batches' mean loads differ by more than 10 percentage points of MCR, the estimate is flagged as potentially load-extrapolated. Ambient confounding: where ambient correction has not been applied and mean scavenge or seawater temperature differs materially across batches, the estimate is flagged accordingly. Placebo test: batch boundaries are randomly reassigned to arbitrary dates, and the regression re-estimated; coefficients of similar magnitude and significance under placebo assignment would indicate the model is fitting some other source of temporal variation rather than a genuine batch effect. Repeated-batch corroboration: where the same supplier and specification recur at different, non-adjacent points in the observation window, independent agreement between their estimated coefficients constitutes stronger evidence of a genuine, repeatable effect than a single-occurrence estimate.

4. Demonstration on a Synthetic Dataset

4.1 Purpose and design

Because access to a full year of verified fleet operating data was not available at the time of writing, the method is first demonstrated on a synthetic dataset constructed to resemble one year of noon-report data for a representative handysize bulk carrier (9,500 kW MCR, ~300 sea days), with known ground-truth effects deliberately injected into the data-generating process and then withheld from the estimator. This is standard practice for validating an econometric estimator prior to application on real data.

4.2 Data-generating process

Six synthetic bunker batches were defined across the twelve-month window, with net calorific values ranging from 40,150 to 41,900 kJ/kg and one batch (Batch C) assigned an injected combustion-quality penalty of +3.6 g/kWh relative to the reference batch, associated with elevated catalytic fines (41 mg/kg) and an elevated CCAI (864) consistent with poorer ignition quality. A continuous engine degradation trend of +0.16 g/kWh per month was injected independent of fuel batch. Daily SFOC was generated as the sum of a load-dependent baseline curve, the batch's calorific penalty, the injected combustion-quality effect, the degradation trend, and Gaussian measurement noise (σ = 1.6 g/kWh). Port days and days with Beaufort exceeding 4 were generated but excluded from the estimation sample by the filtering rules of Section 3.4.

4.3 Results

Table 2 reports the estimated combustion-quality penalty γ_i for each batch, its 95% confidence interval under Newey-West standard errors, and the injected true value. The model was fitted on n = 274 steady-state days (of 360 calendar days generated), with R^2 = 0.663.

Batch	Days	$\hat{\gamma}$ (g/kWh)	95% CI	True γ (g/kWh)
Batch A (reference)	40	0 (ref.)	—	0.0
Batch B	43	1.33	[0.43, 2.24]	1.3
Batch C	37	3.24	[2.14, 4.33]	3.6
Batch D	52	1.15	[−0.37, 2.68]	0.2
Batch E	44	1.74	[−0.28, 3.75]	1.5
Batch F	58	1.32	[−1.50, 4.14]	−0.1

The continuous degradation-drift term was estimated at $\hat{\beta}_3$ = 0.062 g/kWh per month (95% CI: [−0.214, 0.337]), against an injected true value of 0.16 g/kWh per month; the true value falls within the confidence interval and the sign is correctly recovered as positive, though the magnitude is underestimated in this run. The placebo test, using five randomly assigned batch boundaries unrelated to the true bunkering dates, produced no spuriously significant coefficients (point estimates ranging from −1.46 to +0.72 g/kWh, each with a standard error larger than the point estimate).

Figure 4 — Estimation pipeline

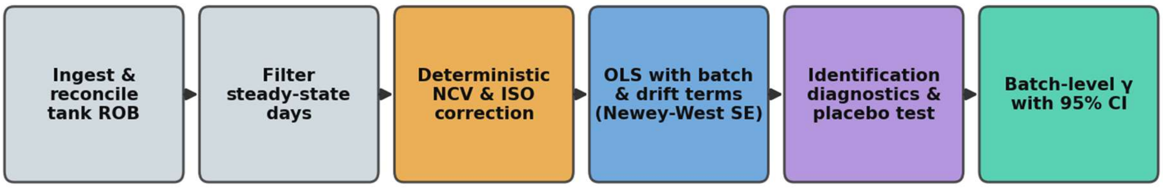

Figure 1. Estimation pipeline from raw ship reporting data to attributed combustion-quality penalty.

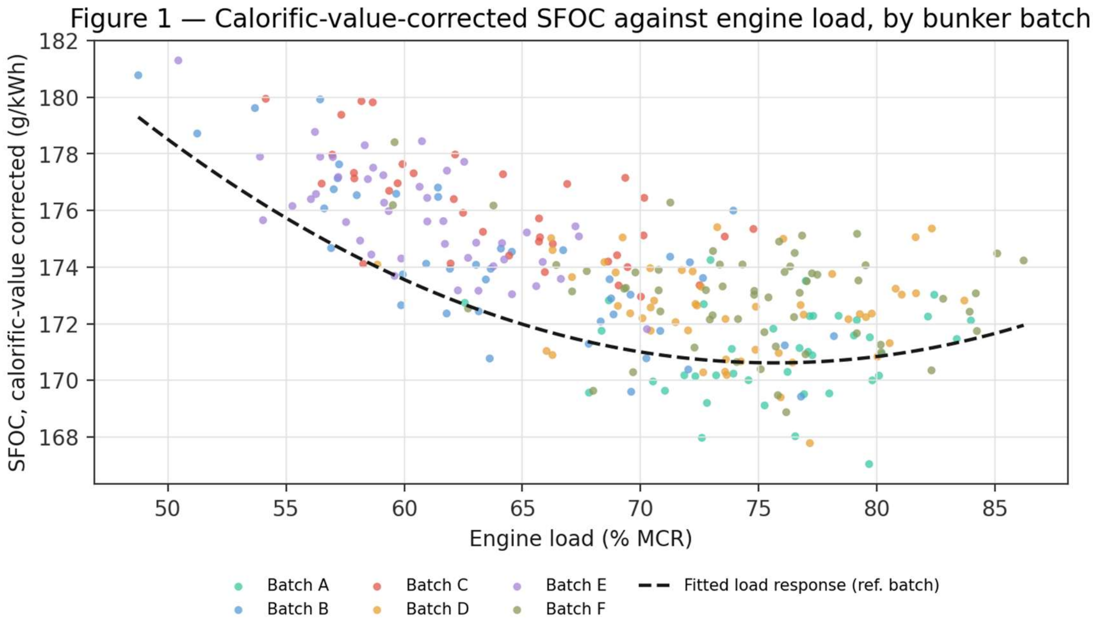

Figure 2. Calorific-value-corrected SFOC against engine load, coloured by bunker batch, with the fitted reference-batch load response.

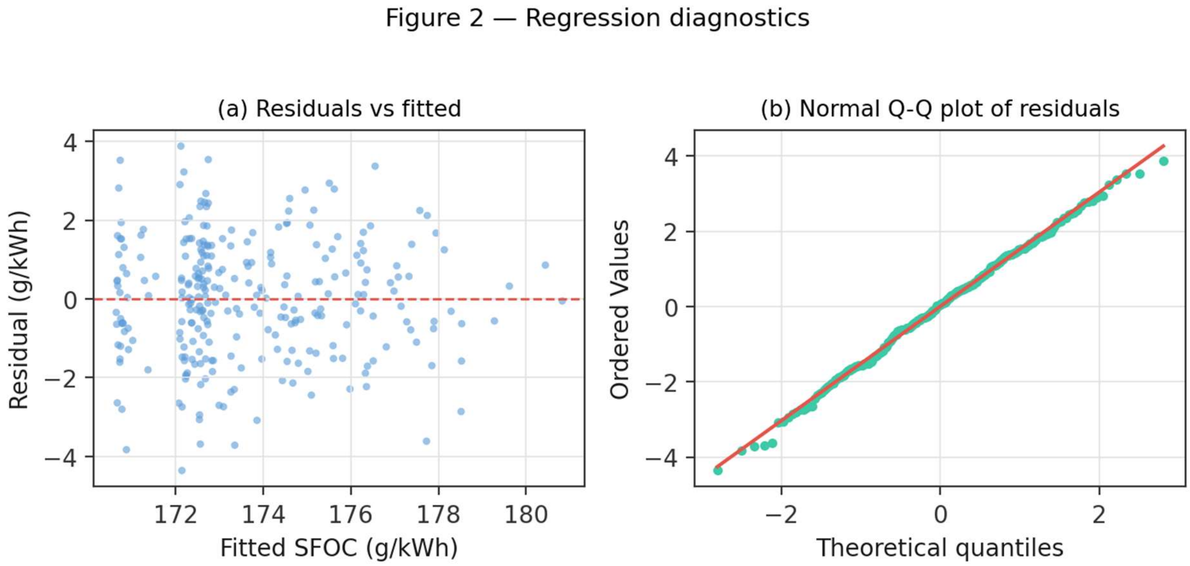

Figure 3. Regression diagnostics: (a) residuals against fitted values, showing no strong systematic pattern; (b) normal Q-Q plot of residuals.

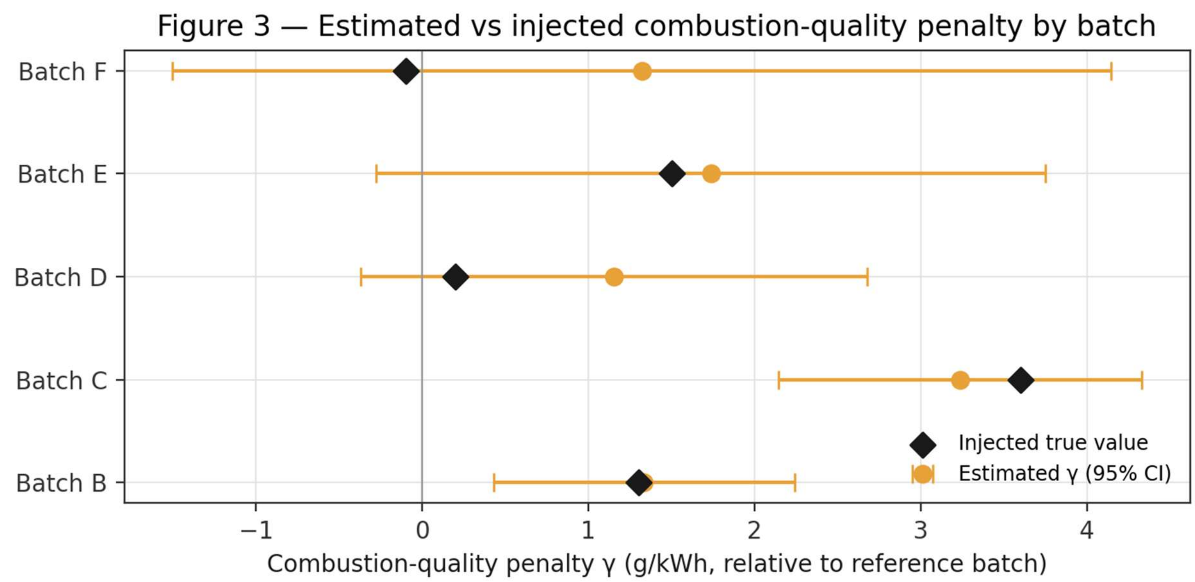

Figure 4. Estimated combustion-quality penalty (95% CI, Newey-West) against the injected true value, by batch.

4.4 Interpretation

Batch C — the batch with the largest injected effect and the one motivating the demonstration — is recovered at 3.24 g/kWh with a confidence interval that contains the true value and excludes zero by a wide margin, supporting the central claim of the method: a combustion-quality penalty of this magnitude is statistically detectable from routine in-service reporting data once calorific value, load and time-drift are controlled for.

However, the results for batches D, E and F show a pattern worth reporting explicitly rather than obscuring: point estimates are biased upward relative to their small or negative injected true values, and confidence intervals widen substantially for batches occurring later in the observation window. The cause is identifiable from the model structure itself: because bunker batches in this synthetic design occur sequentially in time — a realistic feature, since a vessel bunkers in chronological order — later batch dummies are partially collinear with the continuous degradation-drift term in equation (4). Both variables increase monotonically over the observation window for batches toward the end of the series, and the regression has correspondingly less independent variation with which to separate "this batch burns worse" from "the engine has degraded by this point in the year." This is precisely the confounding structure the identification diagnostics of Section 3.5 are designed to flag, and it is a genuine, expected limitation of fixed-effects estimation from sequential operational data, not an artefact specific to this implementation.

Two practical implications follow. First, the method is best powered to detect a combustion-quality effect when the same supplier or specification recurs at different, non-adjacent points in the year, since repetition breaks the sequential collinearity that a single, once-only batch cannot escape. Second, confidence interval width, not the point estimate alone, should govern how much weight is placed on any single batch's estimated penalty — batches D, E and F in this demonstration would each be reported, correctly, as not statistically distinguishable from a modest or null effect, which is the correct and safe conclusion given the data, even though it is a less clean result than Batch C's.

This outcome — one clearly identified effect, several plausible but imprecisely bounded effects, and a clean placebo test — is, in the authors' view, a more credible validation than a demonstration in which every injected effect is recovered with tight precision, since real fleet data will exhibit exactly this kind of partial identifiability, and a method that cannot show its own limits under synthetic conditions should not be trusted to report them honestly under real ones.

5. Discussion

5.1 Practical implications for bunker procurement

Where applied to real fleet data, this method converts an informal, frequently contested question — "was that a bad bunker or a tired engine?" — into a quantified, interval-estimated answer attributable to a specific, traceable stem. This has direct application in fuel quality dispute resolution, where a combustion-quality penalty estimated from in-service data, alongside independent laboratory ignition-quality indices already used in such disputes, strengthens the evidentiary basis for a claim against a supplier. Aggregated across repeated bunkering at the same port or from the same supplier over time, the method also supports a shift in procurement practice from evaluating fuel purely on price-per-tonne to evaluating it on a cost-per-kWh-delivered basis that incorporates measured combustion performance.

5.2 Instrumentation dependence and result tiers

The precision of the estimated batch coefficients depends materially on the quality of the underlying power measurement. Where a calibrated shaft power meter and mass flowmeter are fitted, daily SFOC noise is on the order of 1.5–2%, and the resulting confidence intervals are correspondingly narrow. Where power must be derived from RPM via a calibrated propeller law, uncertainty in the derived power measurement — compounded by hull and propeller fouling, which is itself time-varying and confounded with the engine degradation term — widens the resulting confidence intervals substantially. This paper recommends that any application of the method explicitly report which instrumentation tier applies and interpret the resulting confidence intervals accordingly.

5.3 Limitations and validation pathway

Several limitations should be acknowledged. First, the method has been validated only on synthetic data with a known confounding structure; real fleet data will present more complex patterns, including non-random correlation between fuel quality and other unobserved vessel or route characteristics. Second, the distributional assumptions underlying the OLS estimator should be checked against real data, and the method's robustness under contaminated or misreported noon-report data — a known issue in the ship performance monitoring literature — has not yet been separately assessed. Third, the demonstration in Section 4 itself surfaced a specific and generalisable limitation: sequential batch timing that is collinear with continuous degradation partially attenuates identification for batches later in an observation window, an effect that repeated, non-adjacent bunkering from the same supplier can mitigate but not eliminate. Fourth, the reconciliation algorithm of Section 3.2 assumes daily tank-wise ROB is available or reasonably approximated by FIFO depletion; fleets without this level of record-keeping will see wider confidence intervals.

The natural next step, and the subject of ongoing work by the author, is application of this methodology to twelve months of verified operating data from one or more commercially operated vessels, with results to be reported in a follow-up publication once available. Such an application would additionally allow direct comparison between the fixed-effects estimates presented here and an alternative decomposition via stochastic frontier analysis, in which combustion-quality and other batch-level effects are estimated as determinants of a formally one-sided inefficiency term (Battese and Coelli, 1995) rather than as ordinary regression coefficients; agreement between the two methods on real data would constitute stronger evidence for the robustness of the attribution than either method alone.

5.4 Relationship to CII and emissions accounting

Because AER (the operational metric underlying the IMO Carbon Intensity Indicator) is fuel consumed per unit of transport work, the batch-level and degradation-level decomposition presented here carries over directly into carbon-intensity terms without requiring a separate statistical model: a combustion-quality penalty expressed in g/kWh translates directly into additional CO_2 emitted per unit of transport work delivered. This suggests a secondary application in decomposing CII trend and attributing it to controllable causes rather than treating an adverse CII trend as an undifferentiated compliance problem — a direction the author intends to develop in future work.

6. Conclusion

This paper has presented a regression-based framework for decomposing in-service SFOC deviation into a deterministic calorific-value component, a combustion-quality component attributable to specific bunker batches, a continuous engine degradation component, and discrete maintenance step-effects, using only data already generated in routine ship operation. The framework addresses a gap between three adjacent literatures — fuel ignition quality characterisation, engine performance correction standards, and data-driven consumption modelling — none of which, to the authors' knowledge, jointly estimates a bunker-batch-specific combustion-quality effect from in-service data under an explicit identification framework. A demonstration on a synthetic dataset with known injected ground truth shows the method clearly recovers a substantial combustion-quality effect within its confidence interval, while honestly surfacing a specific, generalisable identification limitation — sequential batch/degradation collinearity — for smaller effects late in the observation window. Application to verified real fleet data, and comparison against a stochastic-frontier formulation of the same decomposition, are identified as the necessary next steps toward establishing the method's practical reliability and its value to bunker procurement and vessel performance monitoring practice.

Data Availability Statement

The synthetic dataset and estimator implementation used in Section 4 are available from the corresponding author on reasonable request. Real fleet operating data, where used in future work, will be subject to the data-sharing terms agreed with the relevant vessel operator, including anonymisation of vessel identity, route, and counterparty information.

Conflict of Interest Statement

The author is the founder of Alfa Kinetics, which is developing a commercial software implementation (Kairos Insight Engine) of the methodology described in this paper. This relationship is disclosed in the interest of transparency; the methodology, validation results, and conclusions presented are independent of any commercial application.

References

Baert, R.S.G. (1988). The ignition quality of heavy residual fuel oils for marine bunkers. Marine Engineers Review / doctoral thesis, State University of Ghent, Belgium.

Battese, G.E. and Coelli, T.J. (1995). A model for technical inefficiency effects in a stochastic frontier production function for panel data. Empirical Economics, 20(2), 325–332.

Brandt, H.C.A. and Zeelenberg, A.P. (1988). Quality of residual fuels in relation to fuel constitution. 10th Anniversary of the Dutch National CIMAC Committee, Amsterdam, November 1988.

International Organization for Standardization (2002). ISO 15550: Internal combustion engines — Determination and method for the measurement of engine power — General requirements. Geneva: ISO.

International Organization for Standardization (2019). ISO 8217: Petroleum products — Fuels (class F) — Specifications of marine fuels. Geneva: ISO.

International Organization for Standardization (various editions). ISO 3046-1: Reciprocating internal combustion engines — Performance — Part 1: Declarations of power, fuel and lubricating oil consumptions, and test methods. Geneva: ISO.

International Organization for Standardization (2019). ISO 19030: Ships and marine technology — Measurement of changes in hull and propeller performance. Geneva: ISO.

MAN Energy Solutions (various editions). Project guides for MAN B&W two-stroke marine engines. Copenhagen: MAN Energy Solutions.

MAN Energy Solutions (2010). Service Letter SL2010-527: Quality of fuel oil. Copenhagen: MAN Energy Solutions.

Stopford, M. (2009). Maritime Economics, 3rd edition. Abingdon: Routledge.

Wortel, F.M. (2005). Ignition quality of residual fuel oils. Journal of Maritime Research, II(3).

Zeelenberg, A.P., van Draat, H.J.F. and Barker, H.L. (1983). The ignition performance of fuel oils in marine diesel engines. 15th CIMAC Congress, Paris, Paper D13.2.