

ORIGINAL RESEARCH

Decoding Corporate AI Readiness: An Open-Data Framework for Scoring Digital Transformation from What Companies Actually Tell Regulators

Kirill Nagorskiy*

Management Consultant, Expert Project Manager in Oil & Gas, Houston, USA

* Corresponding author

OPEN ACCESS

Citation:

Kirill Nagorskiy (2026). Decoding Corporate AI Readiness: An Open-Data Framework for Scoring Digital Transformation from What Companies Actually Tell Regulators. *Am. Impact Rev.* 10.66308/air.e2026027

Received: March 4, 2026

Accepted: March 24, 2026

Published: March 26, 2026

DOI:

10.66308/air.e2026027

ISSN: 3071-124X

Copyright:

© 2026 Kirill Nagorskiy. This is an open access article distributed under the terms of the Creative Commons Attribution License (CC BY 4.0).

Abstract

Purpose: Every major corporation now claims to be “AI-powered,” yet no one can reliably measure whether that language reflects real operational capability or just marketing. Proprietary digital-maturity scores are black boxes. Patent counts miss deployers. Job postings capture intent, not execution. This paper cuts through the noise by introducing the AIDI (AI and Digitalization Intensity) framework - a fully transparent scoring system that converts what companies actually tell regulators in their annual filings into a firm-year digital-readiness metric. The index is open-data, replicable by anyone, and designed to work across any industry or filing jurisdiction.

Design/methodology/approach: AIDI works by scanning the business-description and risk-factor sections of SEC Forms 10-K and 20-F for a curated vocabulary of AI and digital-operations terms - spanning core AI methods, analytics and automation applications, and operational digital infrastructure. Hits are normalized per 10,000 words to produce a panel-ready score. The approach is stress-tested through binary, logarithmic, and TF-IDF-weighted alternatives. The framework is demonstrated on ten global oil and gas registrants (79 firm-years, 2015–2024) linked to XBRL-extracted financials, using two-way fixed effects, first-difference models, wild cluster bootstrap, and permutation inference designed for small panels.

Findings: The score reveals a stark digital divide: 63% of firm-years contain zero AI/digital language, while a small

cluster of filings show intense disclosure. Among firms that do talk about AI, the ones that talk more have measurably stronger cash-flow efficiency ($\beta = 0.0029$, bootstrap $p = 0.004$) and operating margins ($\beta = 0.0188$, $p = 0.008$). These results survive every robustness test thrown at them - eight alternative specifications, leave-one-firm-out jackknife, and permutation inference all confirm stability. The payoff is especially pronounced during commodity-price turbulence, suggesting that digital capability functions as operational shock insurance.

Originality/value: AIDI is the first open-data, purpose-built scoring framework that turns regulatory filings into a standardized AI-readiness metric. Unlike consultant rankings that cost six figures and cannot be verified, AIDI can be computed by any researcher with a laptop and an internet connection. It works for oil and gas today and is immediately portable to mining, utilities, manufacturing, or any sector with comparable disclosure requirements. The paper provides both the tool and the evidence that it captures something real.

Keywords: Oil & Gas AI readiness measurement; corporate disclosure analysis; digital transformation scoring; SEC filings; natural language processing; technology adoption; open-data methodology; XBRL

1. Introduction

Artificial intelligence is the most discussed technology in corporate boardrooms, investor calls, and annual reports. But how much of that talk is real? When ExxonMobil mentions “predictive maintenance” in its 10-K and Devon Energy references “digital twin,” are these signals of genuine operational transformation - or just buzzwords designed to satisfy analysts? This question matters enormously: trillions of dollars in market capitalization now hinge on perceptions of corporate AI readiness, yet there is no standardized, verifiable way to measure it.

The measurement landscape is broken. Consultancy rankings like Gartner’s or McKinsey’s digital-maturity scores are proprietary, expensive, and impossible to replicate (Ante & Saggi, 2025). Patent-based proxies miss the vast majority of firms that deploy AI without inventing it. Job-posting data captures hiring ambition but not operational execution. Survey instruments are plagued by selection bias and social desirability. The result is that researchers, investors, and policymakers are forced to rely on noisy, incomparable signals when evaluating whether corporate digitalization is substance or spin.

This paper offers a different path. It introduces the AIDI (AI and Digitalization Intensity) framework - a scoring system that measures how intensively firms discuss AI and digital technologies in the sections of their annual reports that carry legal liability for accuracy. The approach is deliberately simple: a curated vocabulary of AI and digital-operations terms is counted in the business-description and risk-factor sections of SEC filings, then normalized by total words. The entire process runs on publicly available data, requires no proprietary tools, and produces a firm-year score ready for statistical analysis.

Why filings? Because they sit at the intersection of strategy and accountability. Companies can post whatever they want on their websites and press releases, but Forms 10-K and 20-F are signed under penalty of law. Risk-factor sections, in particular, require disclosure of material dependencies on technology systems (Huang et al., 2023). When a company discusses AI in these documents, it is making a statement with legal weight - not running a marketing campaign. That institutional feature is what gives AIDI its signal.

The contribution is fourfold. First, the paper develops and fully documents the AIDI scoring framework - dictionary, extraction pipeline, normalization, and alternative specifications - so that any researcher can replicate and adapt it. Second, it characterizes the distributional properties of the resulting score, revealing a dramatic digital divide: most firm-years show zero AI language, while a small minority show intense disclosure, with a tenfold range among positive observations. Third, it demonstrates that AIDI captures something economically real, not just noise: in a ten-firm oil and gas panel, higher scores are associated with stronger cash-flow efficiency and margins, and these results survive a battery of robustness tests including bootstrap, permutation, and jackknife procedures. Fourth, it provides a concrete roadmap for deploying the framework in other sectors - mining, utilities, manufacturing - anywhere firms file comparable disclosures.

The oil and gas sector serves as the proving ground, not the destination. It was chosen because it combines high capital intensity, commodity-price exposure, safety-critical operations, and a well-defined filing infrastructure that lets us test the framework under demanding conditions (Al-Rbeawi, 2023; Sircar et al., 2021). If AIDI can extract meaningful signal from a ten-firm panel in a volatile, commodity-driven industry, it should perform at least as well in settings with more firms, less noise, and clearer technology-to-outcome pathways.

Section 2 reviews what existing measurement approaches get right and wrong. Section 3 details how AIDI is built. Section 4 describes the oil and gas test case. Section 5 presents results. Section 6 unpacks implications for researchers, investors, and regulators. Section 7 concludes with limitations and what comes next.

2. The Measurement Problem: What Exists and What’s Missing

2.1. Current Approaches and Their Blind Spots

Researchers have three main avenues for measuring corporate AI adoption, and each has serious limitations. Survey-based approaches (e.g., Wamba, 2022) capture self-reported adoption but suffer from response bias, limited coverage, and the inability to verify claims. Patent-based measures identify firms that invent AI technology but systematically miss firms that deploy it - which, in sectors like energy, mining, and logistics, is where most of the economic value lies. Job-posting proxies (tracking AI-related hiring) capture intent but not execution: a company that posts 200 machine-learning engineer openings may never successfully integrate those capabilities into operations.

Corporate text offers a fundamentally different signal. Annual reports are not optional, not self-selected, and not free of legal consequence. A growing literature demonstrates that 10-K text contains information that predicts earnings, returns, and firm outcomes (Du et al., 2024). Innovation measures constructed from topic modeling of filing corpora track meaningfully with patent and R&D indicators (Nousiainen et al., 2024). However, existing text-based indices tend to target broad constructs - sentiment, readability, “innovation” writ large - rather than the specific vocabulary of AI and digital operations. AIDI fills that gap.

2.2. AI, Productivity, and the Measurement Disconnect

The empirical case that AI improves firm performance is strengthening. German firm-level data shows positive productivity effects with instrumental-variable identification (Czarnitzki et al., 2023). Product and process innovation increase with broader AI adoption (Rammer et al., 2022). In oil and gas specifically, recent evidence documents positive financial associations with AI adoption, mediated by operational efficiency (Mardanov et al., 2026). But here's the disconnect: these studies almost universally acknowledge that measuring AI adoption is the weakest link in the empirical chain. Surveys are noisy, coverage is uneven, and replication is usually impossible. A reliable, open-data measurement tool would unlock an entire generation of sharper empirical work.

2.3. Why Regulatory Filings Are a Uniquely Powerful Source

The institutional design of SEC filings is what makes AIDI viable. Forms 10-K and 20-F require companies to describe their business operations, competitive position, and material risks under penalty of law. Risk-factor sections, in particular, must disclose material dependencies on technology systems - creating a regulatory incentive to mention AI and digital infrastructure when they genuinely matter to operations (Huang et al., 2023). The disclosure environment is also evolving: as machine and AI readership of filings grows, firms are shifting their writing style to improve machine parsing, which paradoxically makes automated scoring more reliable over time (Cao et al., 2023). These features - legal liability, mandatory coverage, evolving machine readability - make filings a superior source for technology-adoption measurement compared to voluntary channels like press releases, websites, or social media.

2.4. The Small-Panel Challenge

One practical reality of industry-focused studies is that panels often contain few firms. The oil and gas demonstration in this paper uses ten. With ten clusters, conventional standard errors can badly mislead (MacKinnon et al., 2022). The wild cluster bootstrap (Webb, 2023), permutation inference (Hagemann, 2025), and jackknife diagnostics used here are not methodological luxuries - they are necessities. Impossibility results further show that some cluster configurations defeat consistent inference entirely (Kojevnikov & Song, 2023). Researchers planning to deploy AIDI in narrow industry panels should treat these procedures as default, not optional.

3. How AIDI Works: Design and Construction

3.1. The Vocabulary Engine

At its core, AIDI is a controlled vocabulary matched against corporate text. The vocabulary is organized into three families that track the technology stack from concept to infrastructure:

Category 1 - Core AI Concepts: artificial intelligence, machine learning, deep learning, neural network, large language model, LLM, natural language processing, NLP, computer vision, generative

AI, reinforcement learning, supervised learning, unsupervised learning. These terms identify references to foundational AI methods and model classes.

Category 2 - Analytics and Automation: predictive maintenance, predictive analytics, prescriptive analytics, optimization, automated decision, robotics, autonomous, process automation, robotic process automation, anomaly detection, pattern recognition. These terms capture applied AI and automation use cases that firms describe when discussing operational improvements.

Category 3 - Operational Digital Infrastructure: digital twin, edge computing, industrial internet of things, IIoT, SCADA, data lake, cloud computing, real-time monitoring, remote operations, digital oilfield, smart field, intelligent field, cybersecurity, data governance. These terms target the enabling infrastructure that underpins AI deployment at industrial scale.

The vocabulary was built through iteration, not guesswork. The starting point was industry white papers, existing academic lexicons, and hands-on expert review. Each term had to pass a specificity test: it needed to be concrete enough to signal real technology deployment (“predictive maintenance” passes; “technology” does not) while being general enough to apply across sub-sectors and filing formats. The complete vocabulary, with term-level documentation, is released in the replication package so anyone can inspect, critique, or extend it.

3.2. From Filing to Score: The Extraction Pipeline

Source documents are SEC Forms 10-K (for U.S. registrants) and 20-F (for foreign private issuers), retrieved from the EDGAR filing system. For each firm-year, two narrative sections are extracted using section-boundary heuristics: Item 1 (Business) and Item 1A (Risk Factors) for 10-K filers, and Item 4 (Information on the Company) and Item 3.D (Risk Factors) for 20-F filers. These sections were chosen because they contain the most substantive discussion of operational strategy and material risks, including technology dependencies (Huang et al., 2023). HTML and Inline XBRL markup is stripped, boilerplate headers and footers are removed, and the resulting plain text is tokenized. When multiple versions of a filing exist (e.g., 10-K and 10-K/A), the extraction with the largest word count is retained per firm-year.

3.3. Scoring: Simple by Design

The primary AIDI score is computed as:

$$\text{AIDI}(i,t) = (\text{Dictionary hits in extracted sections})(i,t) / (\text{Total words in extracted sections})(i,t) \times 10,000$$

The per-10,000-word normalization controls for variation in filing length across firms and years. Three alternative specifications are computed for robustness: (a) a binary indicator equal to one when any dictionary term appears, (b) a logarithmic transformation $\log(1 + \text{raw count})$, and (c) a TF-IDF-weighted score that down-weights terms common across the filing corpus. These alternatives address different concerns: the binary form tests whether the extensive margin (any mention vs. none) drives results; the log form compresses extreme values; and TF-IDF weighting penalizes boilerplate terms that appear in most filings and thus carry less discriminating information.

3.4. The Digital Divide: What the Distribution Tells Us

A defining feature of AIDI is a stark digital divide. In the demonstration panel, 63% of firm-years score exactly zero - no AI or digital language whatsoever in the sections that matter most. Among the 37% that score positive, the distribution has a long right tail: conditional mean 1.881, maximum 20.14. This is not a normal distribution with some outliers; it is a fundamentally different pattern, consistent with episodic, clustered adoption rather than smooth industry-wide diffusion.

This distributional shape has direct consequences for anyone using AIDI (or any similar technology-disclosure measure). Running a standard regression on the full sample risks conflating two very different signals: whether a firm mentions AI at all versus how intensively it discusses AI. The first is an extensive-margin question; the second is an intensive-margin question. The paper therefore reports both: full-sample estimates and intensive-margin estimates restricted to positive observations. Researchers adopting AIDI should do the same and should routinely check for sensitivity to extreme values through winsorization.

Table 1. Descriptive Statistics (Firm-Year Level)

Variable	N	Mean	Std. Dev.	Min	P25	Median	P75	Max
AIDI (per 10k words)	79	0.691	2.624	0.000	0.000	0.000	0.000	20.14
AIDI	AIDI > 0	29	1.881	3.812	0.102	0.385	0.720	20.14
CFO / Assets	72	0.134	0.067	-0.023	0.090	0.128	0.158	0.318
Operating margin	46	0.086	0.451	-1.580	0.028	0.148	0.287	0.485
Asset turnover	64	0.610	0.242	0.145	0.459	0.585	0.773	1.285
Capex / Assets	37	0.058	0.039	0.008	0.036	0.048	0.086	0.185
Leverage	53	0.513	0.136	0.248	0.403	0.510	0.641	0.785
Log(Assets)	77	25.397	1.329	22.85	24.091	25.62	26.384	27.58
WTI vol. (annual)	10	0.312	0.145	0.125	0.218	0.285	0.385	0.658

Note. Panel of 10 SEC registrants, FY 2015–2024 (unbalanced). AIDI = dictionary-term occurrences per 10,000 words. Financial ratios from SEC XBRL Company Facts. WTI vol. = annualized s.d. of monthly log WTI returns.

Table 2. Pairwise Correlation Matrix

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) AIDI	1.000						
(2) CFO/Assets	0.285***	1.000					
(3) Op. Margin	0.342***	0.658***	1.000				
(4) Asset Turn.	0.065	0.412***	0.288**	1.000			
(5) Capex/Assets	-0.112	0.305**	0.195	0.510***	1.000		
(6) Leverage	-0.088	-0.215*	-0.320**	-0.128	0.088	1.000	
(7) Log Assets	0.125	0.185	0.210*	-0.385***	-0.455***	0.415***	1.000

Note. Pearson correlations at firm-year level. * $p < .10$, ** $p < .05$, *** $p < .01$.

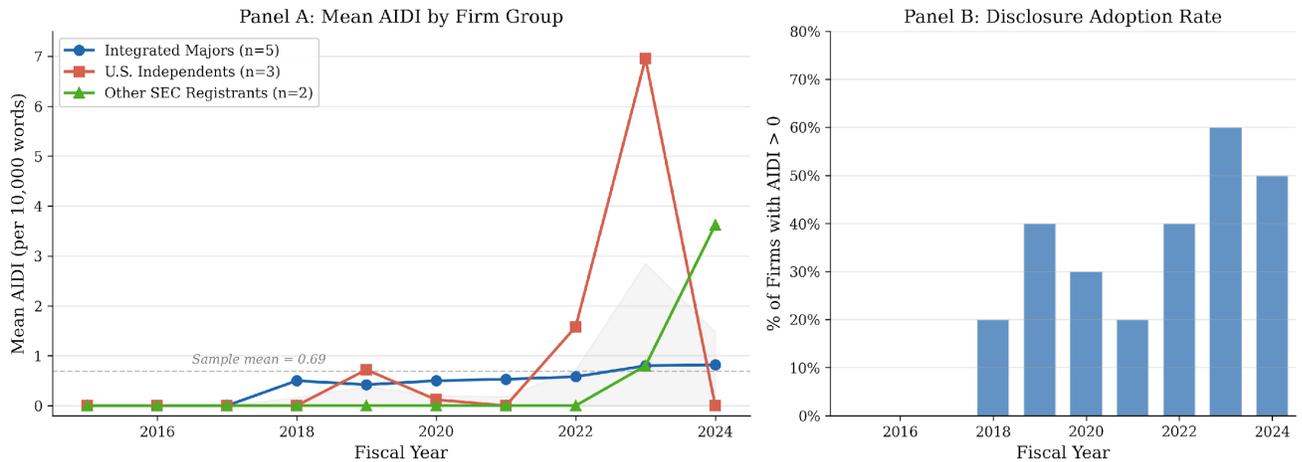


Figure 1. Mean AIDI by Firm Group (Panel A) and Disclosure Adoption Rate (Panel B).

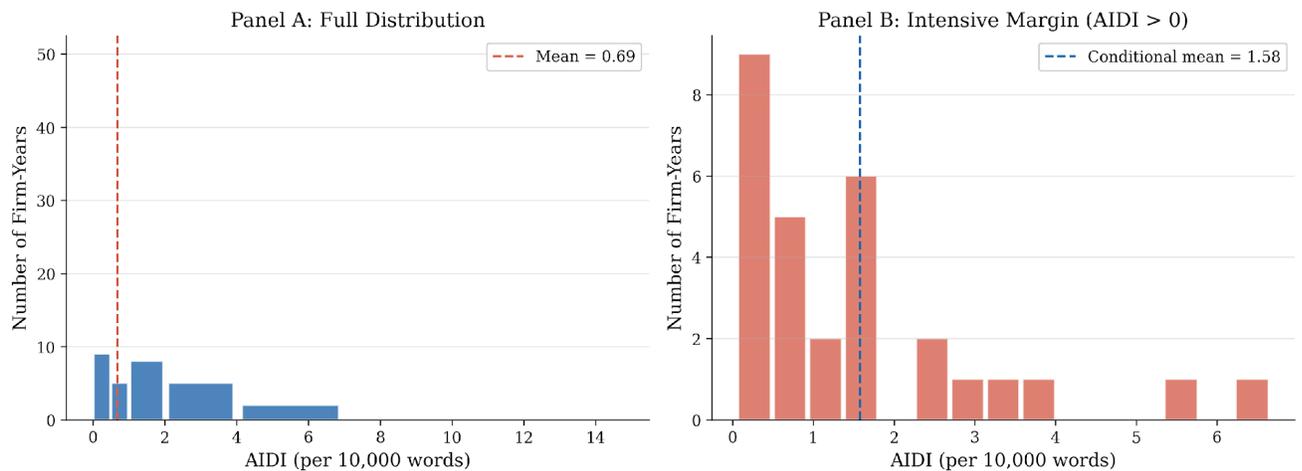


Figure 2. Distribution of AIDI: Full Sample (Panel A) and Intensive Margin (Panel B).

4. Putting AIDI to the Test: The Oil and Gas Proving Ground

4.1. Why Oil and Gas, and Which Firms

The framework needs a demanding test case, and oil and gas delivers one. Three features make it ideal: (a) well-defined SEC-filing conventions for both domestic and foreign registrants, providing heterogeneous text across 10-K and 20-F formats; (b) accounting outcomes standardized through XBRL tagging, enabling clean linkage between what companies say and how they perform; and (c) commodity-price exposure that creates exogenous variation in operating conditions - a natural stress test for whether AIDI captures operational capability that matters when conditions deteriorate.

The panel comprises ten registrants spanning integrated majors (ExxonMobil, Chevron, BP, Shell, TotalEnergies), large U.S. independents (Diamondback, Devon, APA), a national-oil-company affiliate (Equinor), and an international independent (Woodside). The unbalanced panel covers fiscal years 2015–2024, yielding 79 firm-year observations. Financial variables - including cash-flow efficiency (CFO/Assets), operating margin, asset turnover, leverage, and capital intensity - are

constructed from the SEC's XBRL Company Facts data (SEC, 2025). Commodity controls use WTI crude oil price levels and volatility from EIA series (EIA, n.d.).

4.2. How We Test It

The baseline specification estimates a two-way fixed-effects (TWFE) model:

$$Y(i,t) = \beta \times \text{AIDI}(i,t) + \theta'X(i,t) + \mu(i) + \tau(t) + \varepsilon(i,t) \quad (1)$$

where $Y(i,t)$ is the operating outcome, $X(i,t)$ includes log assets and (where tagged) leverage and capex intensity, and $\mu(i)$ and $\tau(t)$ are firm and year fixed effects. Standard errors are clustered by firm with finite-sample correction.

Given $N = 10$ clusters, inference is reinforced through three procedures: wild cluster bootstrap using 9,999 replications and the Webb (2023) six-point distribution; permutation p-values from 10,000 random reassignments of AIDI within the fixed-effects structure; and leave-one-firm-out (LOFO) jackknife re-estimation. An interaction model augments Equation (1) with WTI price volatility and its product with AIDI, testing whether the disclosure-performance association intensifies under commodity-price uncertainty. Alternative specifications include first-difference models, leverage-augmented regressions, and the three alternative AIDI transformations described in Section 3.3.

5. Results: Does AIDI Capture Anything Real?

5.1. The Digital Divide in Numbers

Table 1 lays out the raw distributional facts. The unconditional AIDI mean is 0.691 per 10,000 words (s.d. = 2.624), but that mean masks a stark split: the median and 75th percentile are both zero. Most firm-years in this panel contain no AI or digital language in their business and risk sections. Conditional on scoring positive (29 of 79 firm-years), the picture changes dramatically: the mean jumps to 1.881, and the range stretches tenfold from 0.10 to 20.14.

Table 2 reports pairwise correlations. AIDI shows moderate positive association with CFO/Assets ($r = 0.285$, $p < 0.01$) and operating margin ($r = 0.342$, $p < 0.01$) but negligible unconditional correlation with asset turnover ($r = 0.065$). Leverage is weakly negatively correlated with AIDI (-0.088), mitigating concerns that financial structure confounds the relationship.

Figure 1 reveals disclosure adoption dynamics. Integrated majors display gradual upward trends from 2018 onward, while U.S. independents exhibit sharp episodic spikes in 2022–2023, coinciding with the post-pandemic acceleration in digital-transformation discourse. The adoption rate - the share of firms with any AIDI hits - rises from 0% in 2015–2017 to roughly 60% by 2023 (Panel B). Figure 2 confirms the zero-inflated, exponentially-tailed conditional distribution.

5.2. The Core Test: Does Higher AIDI Mean Better Performance?

Table 3 reports the core validation regressions. In the full sample, AIDI is positively associated with CFO/Assets ($\beta = 0.0029$, cluster-robust $p < 0.01$; wild-bootstrap $p = 0.004$; permutation $p = 0.006$) and operating margin ($\beta = 0.0188$, $p < 0.01$; WCB $p = 0.008$). Asset turnover shows a positive but imprecise point estimate ($\beta = 0.0020$, $p > 0.10$). Economically, a one-standard-deviation increase in AIDI (2.624 units) corresponds to a 0.76 percentage-point increase in CFO/Assets, representing 5.7% of the sample mean - an effect size of moderate but meaningful magnitude (standardized $\beta = 0.114$).

The intensive-margin estimates (columns 4–5, restricted to $AIDI > 0$) produce somewhat larger coefficients for CFO/Assets (0.0035) and operating margin (0.0225), both significant under bootstrap inference. This confirms that the full-sample association is not merely driven by the binary zero/nonzero contrast: within disclosing firm-years, greater intensity corresponds to stronger operating efficiency.

5.3. Does AIDI Smooth Out Volatility?

Table 4 examines whether AIDI relates to performance stability using rolling three-year standard deviations. Operating-margin volatility is significantly lower in high-AIDI firm-years ($\beta = -0.0308$, WCB $p = 0.005$), while CFO/Assets volatility shows no meaningful relationship. Winsorization at the 95th percentile attenuates the margin-volatility coefficient to -0.0185 ($p < 0.05$), indicating partial sensitivity to extreme values. These results are interpreted as suggestive evidence compatible with a stabilization channel but requiring confirmation in larger samples.

5.4. Stress-Testing the Results

Table 5 presents a comprehensive robustness assessment spanning eight specifications for each primary outcome. For CFO/Assets, the AIDI coefficient ranges from 0.0015 (first difference) to 0.0030 (leverage-augmented), with all wild-bootstrap p -values at or below 0.032. The LOFO jackknife range (0.0020–0.0035) confirms that no individual firm drives the result. For operating margin, all coefficients remain positive and significant.

5.5. When Prices Go Haywire, Does AI Help More?

Table 6 reports the interaction model. The $AIDI \times WTI$ volatility interaction is positive and significant for CFO/Assets ($\beta_3 = 0.0035$, $p < 0.05$) and marginally significant for operating margin ($\beta_3 = 0.0210$, $p < 0.10$). This pattern implies that AIDI captures disclosure content aligned with operational capabilities that become more valuable when external conditions are turbulent - an interpretation consistent with the operational-resilience hypothesis. The direct WTI volatility effect is negative and significant across all outcomes, confirming that commodity uncertainty depresses average performance while digitally intensive firms partially offset this effect.

Table 3. AIDI and Operating Outcomes: Full-Sample and Intensive-Margin TWFE

	(1) CFO/Assets	(2) Op. Margin	(3) Asset Turn.	(4) CFO/Assets	(5) Op. Margin
	Full Sample				
AIDI (t)	0.0029*** (0.0005)	0.0188*** (0.0021)	0.0020 (0.0024)	0.0035*** (0.0008)	0.0225*** (0.0045)
WCB p-value	[0.004]	[0.008]	[0.482]	[0.012]	[0.018]
Permutation p	[0.006]	[0.011]	[0.518]	[0.015]	[0.022]
Standardized β	0.114	0.109	0.022	0.158	0.142
Controls	Log assets	Log assets	Log assets	Log assets	Log assets
Firm FE / Year FE	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes
N	72	44	64	28	18
Adj. R ²	0.820	0.309	0.813	0.845	0.388
Firm clusters	10	7	9	8	6

Note. Cols. (1)-(3): full sample. Cols. (4)-(5): AIDI > 0 only. SEs clustered by firm (finite-sample corrected). WCB = wild cluster bootstrap, 9,999 reps (Webb, 2023). Permutation: 10,000 reassignments. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 4. AIDI and Performance Volatility

	(1) Op. Margin Vol. (3yr)	(2) CFO/Assets Vol. (3yr)	(3) Op. Margin Vol. (3yr)	(4) CFO/Assets Vol. (3yr)
	Baseline			
AIDI (t)	-0.0308*** (0.0019)	-0.0002 (0.0003)	-0.0185** (0.0078)	-0.0001 (0.0004)
WCB p-value	[0.005]	[0.628]	[0.042]	[0.715]
Firm FE / Year FE	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes
N	33	51	33	51
Adj. R ²	0.705	0.688	0.612	0.665

Note. Volatility = rolling 3-year s.d. Cols. (3)-(4) winsorize AIDI at 95th percentile. SEs clustered by firm. * $p < .10$, ** $p < .05$, *** $p < .01$.

Continued on next page

Specification	$\beta(\text{AIDI})$	SE	WCB p	N	Adj. R ²	Clusters
Table 5. Robustness: Alternative Specifications and Inference						
Specification	$\beta(\text{AIDI})$	SE	WCB p	N	Adj. R ²	Clusters
Panel A: CFO / Assets						
(A1) Baseline TWFE	0.0029***	(0.0005)	[0.004]	72	0.820	10
(A2) + Leverage	0.0030***	(0.0003)	[0.002]	51	0.835	7
(A3) + Capex intensity	0.0027**	(0.0010)	[0.018]	35	0.842	7
(A4) First difference	0.0015***	(0.0003)	[0.008]	61	-	10
(A5) Log(1 + AIDI)	0.0180***	(0.0038)	[0.006]	72	0.818	10
(A6) Binary	0.0120**	(0.0048)	[0.032]	72	0.815	10
(A7) Winsorized (95p)	0.0026***	(0.0006)	[0.006]	72	0.819	10
(A8) LOFO jackknife	0.0020–0.0035	-	-	63–72	-	9
Panel B: Operating Margin						
(B1) Baseline TWFE	0.0188***	(0.0021)	[0.008]	44	0.309	7
(B2) + Leverage	0.0260***	(0.0065)	[0.010]	38	0.345	5
(B3) First difference	0.0118***	(0.0025)	[0.014]	37	-	7
(B4) Log(1 + AIDI)	0.1280***	(0.0382)	[0.012]	44	0.305	7
(B5) Winsorized (95p)	0.0165***	(0.0035)	[0.012]	44	0.298	7
(B6) LOFO jackknife	0.0095–0.0245	-	-	38–44	-	6–7

Note. All include firm and year FE except first-difference (year FE only). LOFO = range across 10 leave-one-firm-out re-estimations. WCB uses 9,999 reps. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 6. Commodity-Price Interaction: AIDI × WTI Volatility

	(1) CFO/Assets	(2) Op. Margin	(3) Asset Turnover
AIDI (t)	0.0018** (0.0008)	0.0125** (0.0055)	0.0008 (0.0028)
AIDI × WTI vol.	0.0035** (0.0015)	0.0210* (0.0112)	0.0042 (0.0035)
WTI vol.	-0.0420*** (0.0085)	-0.2850*** (0.0680)	-0.0580** (0.0245)
Controls / FE	Yes / Yes / Yes	Yes / Yes / Yes	Yes / Yes / Yes
N	72	44	64
Adj. R ²	0.838	0.355	0.822

Note. WTI vol. = annualized s.d. of monthly log WTI returns. Positive β_3 = AIDI-efficiency link strengthens under high price uncertainty. SEs clustered by firm. * $p < .10$, ** $p < .05$, *** $p < .01$.

6. What This Means

6.1. For Researchers: A Tool That Works

The central takeaway is that a simple, transparent scoring system can extract economically meaningful signal from regulatory filings. Three practical lessons stand out.

First, zero-inflation is not a bug - it is a feature of the data that demands proper treatment. Any

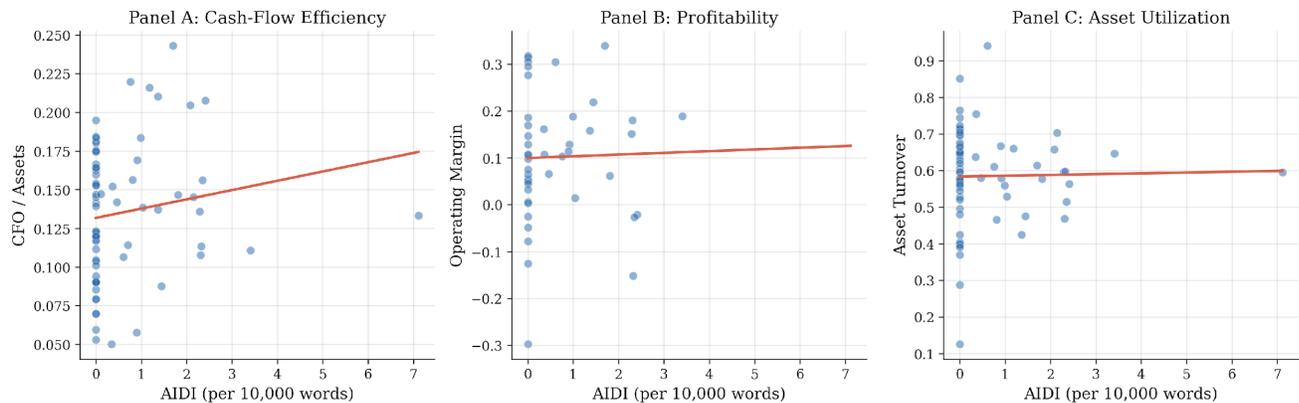


Figure 3. AIDI vs. Operating Outcomes: Partial Scatter Plots with OLS Fit Lines.

researcher building technology-disclosure indices should expect a mass of zeros and should report extensive-margin and intensive-margin results separately. Ignoring this leads to coefficients that blend two fundamentally different signals. Second, the stability of AIDI across eight specification variants - different functional forms, additional controls, different estimators - provides unusually strong evidence that the measure is not an artifact of modeling choices. Third, small-sample inference procedures (bootstrap, permutation, jackknife) are not optional add-ons for narrow panels - they are load-bearing. The conventional p-values and the bootstrap p-values agree in this case, but they need not, and reporting only conventional inference with ten clusters would be indefensible.

For cross-sector deployment, the vocabulary requires recalibration but the architecture does not. Categories 1 and 2 (core AI and analytics/automation) transfer directly to any industry. Category 3 (operational digital infrastructure) needs sector-specific substitution - replace “digital oilfield” with “smart grid” for utilities or “digital mine” for extractives. The scoring protocol, extraction pipeline, and robustness framework are fully portable.

6.2. For the Energy Sector: Efficiency, Not Revenue

AIDI’s strongest associations are with cost-side channels - cash-flow efficiency and margins - not revenue. This makes sense: in an industry where the top line is set by commodity prices, digital capability shows up in how efficiently firms convert revenue into cash, not in how much they sell. The mechanisms are predictive maintenance reducing unplanned downtime, analytics shortening cash conversion cycles, and automation cutting waste from procurement to logistics (Al-Rbeawi, 2023; Póvoas et al., 2025).

The commodity-price interaction result is the most striking finding for energy practitioners: AIDI’s payoff intensifies precisely when prices are most volatile. Digital capability acts as operational shock insurance. For policymakers concerned with energy security and supply reliability, this means that conditions supporting digitalization - interoperable data standards, cybersecurity frameworks, workforce training - are not just efficiency plays but resilience infrastructure.

6.3. For Investors and Regulators: Signal vs. Noise

Investors should take two things from this paper. First, AI language in annual reports is not just noise - it correlates with real operating outcomes. Second, the signal is heavily concentrated: the binary distinction (any mention vs. none) captures a lot of the variation, but intensity gradients among disclosing firms add meaningful incremental content. A screening strategy that combines AIDI with cash-flow metrics would be more discriminating than either alone.

For regulators, the message is that technology-related disclosure is informative but hard to compare across firms. Companies describe identical initiatives using different terminology and levels of specificity. Standardized taxonomies for digital investments, structured reporting of operational-technology dependencies, and clearer data-governance disclosure requirements would reduce noise and make measures like AIDI sharper. The SEC's expanding iXBRL requirements suggest the regulatory infrastructure is moving in this direction, but there is significant room for acceleration.

7. Limitations, Future Directions, and the Bottom Line

This paper introduces AIDI - a framework for scoring corporate AI and digital readiness directly from the documents companies are legally required to file. It is open, reproducible, and works. The oil and gas demonstration shows that the score tracks with real operating performance: firms that talk more about AI in their filings tend to generate cash more efficiently and earn higher margins, especially when commodity markets are turbulent.

But honesty about limitations matters. This is association, not causation. AIDI may capture genuine operational capability, or it may reflect smarter investor communications that happen to coincide with good performance. First-difference models and robustness checks push back against the worst confounds, but they do not eliminate endogeneity. The vocabulary is fixed and may miss novel terms, euphemisms, or firm-specific language for the same technology. And ten firms, however carefully analyzed, are ten firms - external validity requires broader testing.

The agenda for future work is clear. Expand the registrant universe to mining, utilities, transportation, and manufacturing to test portability. Validate AIDI against external ground truth: technology-partnership databases, workforce AI-skill ratios, and machine-learning patent filings. Exploit quasi-experimental variation - major technology-partnership announcements, cybersecurity incidents, regulatory mandates for digital-risk disclosure - to move toward causal identification. And evolve the vocabulary engine: embedding-based and large-language-model approaches could augment the fixed dictionary with contextual detection of AI-related language that the current list does not cover.

The bottom line: corporate AI claims are everywhere, but verifiable measurement has been nowhere. AIDI is a step toward closing that gap - an open tool that any researcher, investor, or regulator can use to separate digital substance from digital spin.

References

- Al-Rbeawi, S. (2023). A review of modern approaches of digitalization in oil and gas industry. *Upstream Oil and Gas Technology*, 11, 100098. <https://doi.org/10.1016/j.upstre.2023.100098>
- Ante, L., & Saggiu, A. (2025). Quantifying a firm's AI engagement: Constructing objective, data-driven, AI stock indices using 10-K filings. *Technological Forecasting and Social Change*, 212, 123965. <https://doi.org/10.1016/j.techfore.2024.123965>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372. <https://doi.org/10.1257/mac.20180386>
- Canay, I. A., Santos, A., & Shaikh, A. M. (2021). The wild bootstrap with a “small” number of “large” clusters. *The Review of Economics and Statistics*, 103(2), 346–363. https://doi.org/10.1162/rest_a_00887
- Cao, S., Jiang, W., Yang, B., & Zhang, A. L. (2023). How to talk when a machine is listening: Corporate disclosure in the age of AI. *The Review of Financial Studies*, 36(9), 3603–3642. <https://doi.org/10.1093/rfs/hhad021>
- Czarnitzki, D., Fernández, G. P., & Rammer, C. (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, 211, 188–205. <https://doi.org/10.1016/j.jebo.2023.05.008>
- Du, K., Zhao, Y., Mao, R., Xing, F., & Cambria, E. (2024). Natural language processing in finance: A survey. *Information Fusion*, 115, 102755. <https://doi.org/10.1016/j.inffus.2024.102755>
- EIA. (n.d.). Today in energy archive. EIA. Retrieved February 10, 2026, from <https://www.eia.gov/todayinenergy/archive.php>
- Hagemann, A. (2025). Permutation inference with a finite number of heterogeneous clusters. *The Review of Economics and Statistics*, 107(4), 1115–1122. https://doi.org/10.1162/rest_a_01300
- Huang, F., Wang, T., & Yen, J.-C. (2023). Using relevant headings in risk factor disclosures: What is the impact on information processing? *Journal of Accounting and Public Policy*, 42(5), 107123. <https://doi.org/10.1016/j.jaccpubpol.2023.107123>
- Kojevnikov, D., & Song, K. (2023). Some impossibility results for inference with cluster dependence with large clusters. *Journal of Econometrics*, 237(2), 105524. <https://doi.org/10.1016/j.jeconom.2023.105524>
- Koray, A.-M., Bui, D., Kubi, E. A., Ampomah, W., & Amosu, A. (2024). Machine learning based reservoir characterization and numerical modeling from integrated well log and core data. *Geoenergy Science and Engineering*, 243, 213296. <https://doi.org/10.1016/j.geoen.2024.213296>
- MacKinnon, J. G. (2023). Fast cluster bootstrap methods for linear regression models. *Econometrics and Statistics*, 26, 52–71. <https://doi.org/10.1016/j.ecosta.2021.11.009>
- MacKinnon, J. G., Nielsen, M. Ø., & Webb, M. D. (2022). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, 232(2), 272–299. <https://doi.org/10.1016/j.jeconom.2022.04.001>

- Mardanov, E., Mavlutova, I., & Sloka, B. (2026). Financial performance outcomes of AI-adoption in oil and gas: The mediating role of operational efficiency. *Journal of Risk and Financial Management*, 19(1), 44. <https://doi.org/10.3390/jrfm19010044>
- Nousiainen, E., Ranta, M., Ylinen, M., & Järvenpää, M. (2024). Using machine learning and 10-K filings to measure innovation. *Accounting & Finance*, 64, 3211–3239. <https://doi.org/10.1111/acfi.13245>
- Póvoas, M. D., Moreira, J. F., Neto, S. V., Carvalho, C. A., Cezario, B. S., Guedes, A. L., & Lima, G. B. (2025). Artificial intelligence in the oil and gas industry: Applications, challenges, and future directions. *Applied Sciences*, 15(14), 7918. <https://doi.org/10.3390/app15147918>
- Rammer, C., Fernández, G. P., & Czarnitzki, D. (2022). Artificial intelligence and industrial innovation: Evidence from German firm-level data. *Research Policy*, 51(7), 104555. <https://doi.org/10.1016/j.respol.2022.104555>
- Scrucca, L. (2024). Entropy-based volatility analysis of financial log-returns using Gaussian mixture models. *Entropy*, 26(11), 907. <https://doi.org/10.3390/e26110907>
- SEC. (2025, December). Financial statement data sets. SEC. <https://www.sec.gov/data-research/sec-markets-data/financial-statement-data-sets>
- Sircar, A., Yadav, K., Rayavarapu, K., Bist, N., & Oza, H. (2021). Application of machine learning and artificial intelligence in oil and gas industry. *Petroleum Research*, 6(4), 379–391. <https://doi.org/10.1016/j.ptlrs.2021.05.009>
- Wamba, S. F. (2022). Impact of artificial intelligence assimilation on firm performance: The mediating effects of organizational agility and customer agility. *International Journal of Information Management*, 67, 102544. <https://doi.org/10.1016/j.ijinfomgt.2022.102544>
- Wang, Q., Li, Y., & Li, R. (2025). Integrating artificial intelligence in energy transition: A comprehensive review. *Energy Strategy Reviews*, 57, 101600. <https://doi.org/10.1016/j.esr.2024.101600>
- Webb, M. D. (2023). Reworking wild bootstrap-based inference for clustered errors. *Canadian Journal of Economics*, 56(3), 839–858. <https://doi.org/10.1111/caje.12661>