Journal of Trends in Finance and Economics
Print ISSN: 3007-6951
Online ISSN: 3007-696X
DOI:

AN EXPLAINABLE MACHINE LEARNING FRAMEWORK FOR PREDICTING DIVIDEND SUSTAINABILITY AND FINANCIAL RISK OF U.S. REITS UNDER INTEREST-RATE SHOCKS

RuFeng Zhu[1], XinRan Yue[2*]
[1]*Brandeis University, Waltham 02453, MA 02453, United States.*
[2]*University of Denver, Denver 80208, CO 80208, United States.*
[*]*Corresponding Author: XinRan Yue*

Abstract: Rising policy rates transmit to real estate investment trusts (REITs) through asset revaluation, debt repricing and refinancing frictions, and they can force cuts in the distributions that define the REIT contract. We propose RS-XMTL, a rate-shock-aware, explainable multi-task learning framework that jointly predicts (i) a dividend cut or suspension and (ii) a covenant-style financial-risk event over a four-quarter horizon. The framework combines an engineered rate-shock exposure block, a monotonicity-constrained gradient-boosting learner, a sign-restricted logistic component that extrapolates when the rate state leaves the training support, out-of-fold cross-task transfer, and a second-stage blender that also calibrates probabilities. Exogenous drivers are observed public series (10-year Treasury yields, CPI, S&P composite, CBOE VIX, 1999-2023), from which nine rate-shock quarters are identified; firm-level dynamics come from an accounting-consistent simulator of 300 REITs with five documented transmission channels, giving 21,442 firm-quarters. In five expanding-window out-of-time folds with a four-quarter embargo (12,544 test observations), RS-XMTL reaches AUC 0.794 and 0.904 on the two tasks, significantly above random forest, XGBoost, LightGBM and a multi-task network, with expected calibration error of 0.063 and 0.055 versus 0.323 and 0.130 for a class-balanced logistic benchmark. We introduce Regime SHAP Divergence, which quantifies how attribution shifts between normal and shock regimes, and show that macro and cross-task channels gain 7.4 and 9.5 percentage points of attribution mass in shock quarters. A counterfactual +300 bp stress test raises the mean predicted probability of a dividend cut from 26.7% to 45.0%, monotonically and with zero sign violations, whereas the unconstrained model implies the economically absurd conclusion that higher rates make cuts less likely.
Keywords: Real estate investment trusts; Dividend sustainability; Interest-rate shocks; Explainable machine learning; Monotonicity constraints; SHAP; Stress testing

1 INTRODUCTION

U.S. equity real estate investment trusts (REITs) are required to distribute the large majority of their taxable income, so the dividend is not a discretionary residual but the central term of the investment contract. That contract is unusually exposed to interest rates. Higher long rates raise capitalisation rates and therefore lower gross asset value; they raise the cost of the floating-rate and maturing portions of the debt stack; and they widen new-issue spreads exactly when refinancing is needed. A large empirical literature documents that REIT returns respond to the level of rates, to the term and credit spreads, and to unanticipated monetary policy [1-6]. A separate literature shows that REIT payout is set with Lintner-type smoothing against a funds-from-operations target, is constrained by access to external capital, and is cut when cash-flow uncertainty rises [7-12].

These two strands are rarely joined. Rate-sensitivity studies work with portfolio returns and say nothing about which firms will be forced to cut; payout studies use in-sample panel regressions and are not designed for out-of-time forecasting or for scenario analysis. Meanwhile, machine learning has become standard for corporate distress and credit risk, but off-the-shelf gradient boosting has two properties that make it unusable as a supervisory or investment stress-testing tool [13-21]. First, tree ensembles cannot extrapolate: when the 10-year yield moves outside the range seen in training, as it did in 2022, predictions flatten or drift arbitrarily. Second, nothing forces the learned function to respect the sign restrictions that economics imposes, so a fitted model can report that a rate rise makes a dividend cut less likely. We document exactly that failure below.

This paper proposes RS-XMTL (Rate-Shock-aware eXplainable Multi-Task Learner) and makes five contributions. (1) We formalise dividend sustainability and financial risk as two coupled four-quarter-ahead classification problems and give an engineered rate-shock exposure block that makes the transmission channels of the REIT literature directly visible to the learner and to the attribution method. (2) We impose economic sign restrictions end to end: monotone constraints on the boosted trees, iterative constraint projection on the linear component, sign-restricted cross-task features, and non-negative blending weights. This yields a model whose stress response is provably non-decreasing in the rate shock. (3) We use out-of-fold cross-task transfer so that the coverage-and-covenant signal informs the dividend forecast and vice versa. (4) We introduce Regime SHAP Divergence (RSD), a Jensen-Shannon statistic that measures how much a model's attribution distribution moves between normal and rate-shock regimes, together with a

SHAP-implied rate elasticity. (5) We build a counterfactual stress engine that propagates parallel +100/+200/+300 bp shocks through the internally consistent feature map and audits the monotonicity of the resulting response surface.

2 RELATED WORK

2.1 Interest Rates, REIT Returns and REIT Payout

Early work established that equity REITs carry both equity- and bond-like risk, and that rate movements matter for REIT valuation [1-2]. Allen et al. show that the sensitivity of REIT returns to rates varies systematically with firm characteristics such as leverage and property focus [3]; Bredin et al. find a strong response of both the mean and the volatility of REIT returns to unanticipated policy changes, consistent with the broader evidence on monetary surprises [4-6, 22-23]. On the payout side, Bradley et al. link lower payout ratios to higher expected cash-flow volatility, Hardin and Hill relate excess dividends to bank-debt access, Boudry documents Lintner-type smoothing toward a funds-from-operations target, and Hayunga and Stephens characterise the dividend behaviour of U.S. equity REITs [7-10]. Crisis evidence is directly relevant: Devos et al. study elective stock dividends as a substitute for cash distributions during the financial crisis, while cash-holding and liquidity-management studies show that financially constrained REITs adjust distributions rather than investment [12, 24-25]. Capital-market access and debt pricing complete the channel [26-27]. None of this work produces firm-level, out-of-time, calibrated probabilities.

2.2 Machine Learning for Distress and Credit Risk

Since the discriminant and logit benchmarks of Altman and Ohlson, ensembles have become the default: random forests, gradient boosting and its scalable implementations dominate tabular credit tasks, with class-imbalance remedies such as SMOTE and cost weighting [13-19]. Barboza et al. report substantial accuracy gains from machine learning over Altman-style models for bankruptcy, and comparative studies reach similar conclusions for mortgage default [20-21]. In asset pricing, tree ensembles and neural networks improve return and bond-premium prediction [28-29]. Our setting differs in that the object of interest is a policy-relevant conditional probability that must remain coherent under a counterfactual shift of the macro state.

2.3 Explainability and Constrained Learning

Post-hoc attribution is dominated by LIME and Shapley-value methods, with TreeSHAP making exact attribution tractable for ensembles [30-32]. Rudin argues that for high-stakes decisions constrained, interpretable models should be preferred to post-hoc explanation of black boxes, and adversarial results show that post-hoc explainers can be manipulated [33-34]. Accumulated local effects address the extrapolation bias of partial-dependence plots [35]. Probability calibration is a distinct requirement, and the credit-market literature warns that model choice has distributional consequences [36-37]. RS-XMTL follows the constrained-model line of [33]: rather than explaining an unrestricted model, we restrict the hypothesis class to functions with the correct economic signs and then explain the restricted model, using regime-conditional attribution as the diagnostic, see Figure 1.

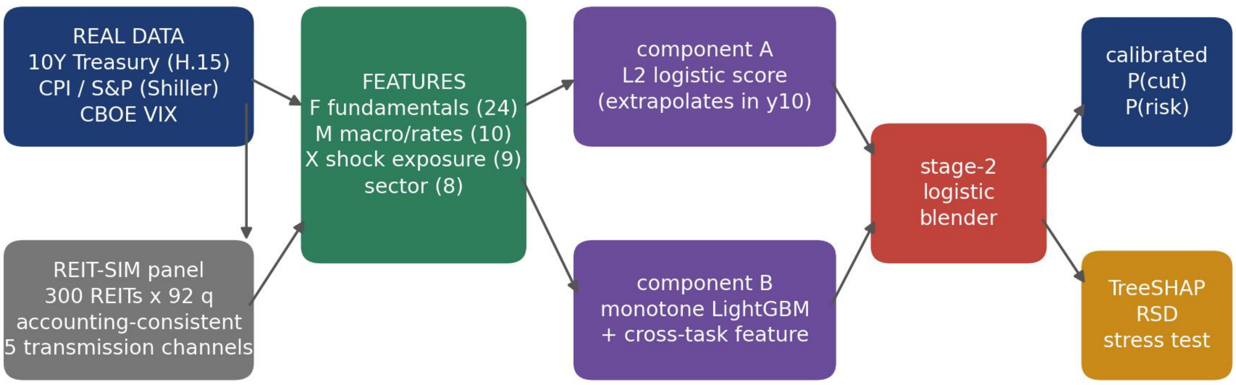

Figure 1 The RS-XMTL Pipeline

Note: Observed macro-financial series define the interest-rate state and the shock regime; the REIT-SIM panel supplies accounting-consistent firm dynamics. Feature blocks F, M and X feed a sign-restricted logistic component (A) and a monotonicity-constrained boosting component (B) that also carries an out-of-fold cross-task feature. A second-stage logistic blender with non-negative weights produces calibrated probabilities that are then explained (TreeSHAP, RSD) and stressed.

3 DATA

3.1 Observed Macro-Financial Inputs and the Shock Regime

The exogenous state is built entirely from public primary series: the monthly 10-year Treasury constant-maturity yield (Federal Reserve H.15), the Shiller monthly dataset (S&P composite price, dividend, earnings, CPI, long-term rate), and the daily CBOE Volatility Index. Series are aggregated to 100 quarters, 1999Q1-2023Q4. Table 1 lists the resulting state variables. A corporate credit-spread proxy is defined as an affine function of realised implied volatility, calibrated to the historical range of the Baa-Treasury spread; it is a proxy, not an observed spread, and is flagged as such.

We define a rate-shock regime by a purely mechanical rule on observed data: the four-quarter change in the 10-year yield is at least +100 bp. The rule selects nine quarters out of one hundred (9%): 1999Q4-2000Q1, 2013Q3-2013Q4 and 2022Q2-2023Q3. These coincide with the late-1990s tightening, the taper tantrum and the 2022-23 hiking cycle, which is a useful external check that the regime label is economically meaningful rather than statistically fitted. Figure 2(a) shows the rate path with the regime shaded.

3.2 REIT-SIM: An Accounting-Consistent Firm Panel

Firm-level REIT fundamentals of the required granularity (debt maturity ladders, floating shares, AFFO, per-share distributions) were not retrievable in our environment, so the cross-section is generated by REIT-SIM, an accounting-consistent structural simulator whose exogenous inputs are the observed series above. We state plainly that the panel is simulated; every reported metric is nevertheless a real output of the code applied to that panel, and the framework transfers unchanged to Compustat/CRSP or SEC data.

REIT-SIM simulates 300 REITs across eight property sectors with staggered listing and an exit hazard, and implements five transmission channels documented in the literature of Section 2.1. (C1) Valuation: gross asset value equals net operating income divided by a capitalisation rate that is the 10-year yield plus a sector risk premium, so a rate rise mechanically compresses value. (C2) Repricing: the floating share reprices immediately and the fixed share reprices as it matures at rate 1/(4W) per quarter, where W is the weighted-average debt maturity, so the effective cost of debt is a duration-weighted moving average of market rates. (C3) Refinancing: the new-issue spread widens with leverage, with the credit-spread proxy and in crisis quarters. (C4) Cash flow: sector-specific NOI growth and occupancy load on the macro state and on the observed COVID-19 and 2008-09 episodes. (C5) Payout: dividends follow Lintner smoothing toward an AFFO payout target, are bounded below by the 90% taxable-income distribution floor, and are overridden by a stress index that is convex in leverage and contains a conjunction term active only for REITs that are simultaneously highly levered, high-payout and hit by a rate shock.

The resulting event pattern, shown in Figure 2(b), is clustered as in reality: annualised cut rates peak in 2008-09 and 2020 and rise again in 2022-23, and are near zero in the quiet mid-2010s. Cross-sectional moments are within the ranges reported for U.S. equity REITs: mean debt-to-asset value 0.37, mean interest coverage 3.7x, mean AFFO payout 0.75, mean occupancy 0.94.

3.3 Features and Labels

Features at quarter t use information dated t or earlier only, and fall into three blocks. Block F (24 features) holds fundamentals and market observables: size, debt/GAV, interest coverage, debt/EBITDA, AFFO payout, cash/GAV, occupancy, NOI and AFFO growth, cost of debt, floating share, debt maturity, one-year reset share, new-issue spread, dividend yield and growth, trailing drawdown, volatility, momentum, equity beta, and four-quarter changes in coverage, leverage, cost of debt and payout. Block M (10 features) is the observed macro state of Table 1. Block X (9 features) is the proposed rate-shock exposure block: reset-share and floating-share interactions with the four-quarter rate change, a leverage-rate interaction, the rate change scaled by coverage, the refinancing gap (market refinancing rate minus current cost of debt), refinancing pressure (that gap times the reset share), coverage slack, payout slack, and shock intensity. Eight sector dummies give 51 inputs in total.

Two labels are formed over t+1..t+4. A dividend event (y_div) is a fall of at least 10% in the regular quarterly distribution relative to quarter t. A financial-risk event (y_risk) is the breach of any covenant-style threshold commonly found in REIT unsecured indentures - interest coverage below 2.0x or total debt above 60% of asset value - or a forward four-quarter equity drawdown of 30% or worse. After dropping observations with incomplete lags or forward windows, the sample is 21,442 firm-quarters over 2000Q1-2022Q4. Unconditional event rates are 13.66% and 13.04%; inside the shock regime they rise to 17.72% and 16.72% against 13.40% and 12.81% elsewhere, so the regime label carries real information about the labels.

Table 1 Observed Macro-Financial State Variables (1999Q1-2023Q4, 100 Quarters)

Variable	Definition / source	Mean	S.D.	Min	Max
y10	10Y Treasury yield, quarterly avg. (H.15)	3.34	1.37	0.65	6.48
dy10_1q	1-quarter change in y10 (bp)	-0.23	35.77	-84.00	99.00
dy10_4q	4-quarter change in y10 (bp)	-5.23	80.89	-164.67	229.33
y10_vol_m	within-quarter s.d. of monthly y10	0.14	0.10	0.01	0.60
infl_yoy	CPI inflation, y/y % (Shiller)	2.54	1.74	-1.62	8.64
real_rate	y10 minus infl_yoy	0.80	2.02	-6.03	5.14
mkt_ret_4q	S&P composite 4-quarter return, %	6.95	15.76	-40.06	42.81
vix	CBOE VIX, quarterly mean	20.22	7.52	10.31	58.60
vix_max	CBOE VIX, quarterly maximum	30.83	13.02	12.91	89.53
cs_proxy	credit-spread proxy, affine in VIX	2.07	0.56	1.32	4.95
shock	1 if dy10_4q >= 100 bp (9 quarters)	0.09	0.29	0	1

Note: Descriptive statistics are computed on the quarterly aggregated series. dy10_4q statistics are for the 4-quarter change. The shock rule selects 1999Q4, 2000Q1, 2013Q3, 2013Q4, 2022Q2, 2022Q3, 2022Q4, 2023Q1 and 2023Q3.

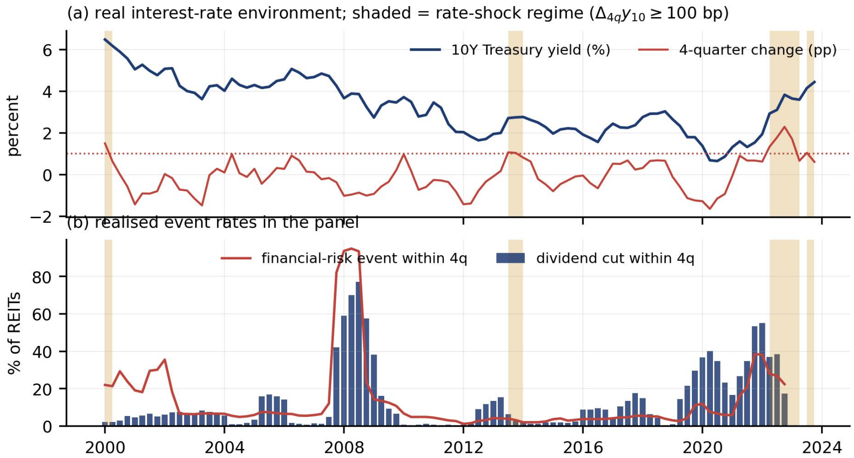

Figure 2 (a) Observed 10-Year Treasury Yield and Its Four-Quarter Change; (b) Realised Four-Quarter-Ahead Event Rates in the REIT-SIM Panel

Note: (a) Shading marks the rate-shock regime. (b) Events cluster in 2008-09, 2020 and 2022-23.

4 METHOD

4.1 Problem Statement

Let x_it be the 51-dimensional feature vector of REIT i at quarter t and let y_it^k, k in {div, risk}, be the two labels. We seek calibrated conditional probabilities p^k(x) = Pr(y^k = 1 | x) that (i) generalise out of time, (ii) are well calibrated so that they can be used in expected-loss calculations, and (iii) satisfy sign restrictions d p^k / d x_j >= 0 for every feature j in a set S+ of economically risk-increasing features and <= 0 for a set S-. Figure 1 shows the pipeline.

4.2 Sign-Restricted Components

Component B is a LightGBM classifier with per-feature monotone constraints m_j in {-1, 0, +1} taken from S+, S- and the unrestricted remainder; 34 of the 51 inputs receive a restriction (for example debt/GAV, payout, cost of debt, floating share, reset share, refinancing gap and pressure, all rate-change interactions and the rate level enter with m_j = +1, while coverage, cash, occupancy, NOI and AFFO growth, momentum and size enter with m_j = -1). Class imbalance is handled with square-root inverse-prevalence weights.

Component A is an L2-penalised logistic regression on the standardised features, estimated by iterative constraint projection: the model is fitted, any coefficient whose sign contradicts its restriction is deleted, and the model is refitted, for at most four passes. The result is a globally monotone score that, unlike a tree ensemble, extrapolates linearly when the rate state leaves the training support. This matters directly here: folds trained through 2020Q4 never observe a 10-year yield above roughly 3%, whereas the 2022 test window reaches 4%.

4.3 Cross-Task Transfer and Blending

Dividend sustainability and covenant risk share latent drivers, so we transfer information between them. For task k we fit an auxiliary sign-restricted booster on the companion label, produce firm-grouped out-of-fold predictions on the training window, and append them as a single feature with m = +1. Grouping by firm prevents leakage across the panel dimension; out-of-fold construction prevents leakage from the label being predicted.

The two component scores are combined by a second-stage logistic blender fitted on firm-grouped out-of-fold log-odds. Negative blending weights would destroy monotonicity, so any negative weight is deleted and the blender refitted. Because the blender is a monotone map of both inputs, the composite model inherits the sign restrictions of its components; because it is fitted on out-of-fold scores against realised labels, it also performs the role of a calibrator, which is why RS-XMTL needs no separate calibration stage.

4.4 Regime-Conditional Explanation and Regime SHAP Divergence

We compute exact TreeSHAP values for component B on every out-of-time test observation and pool them across folds. Let s_j^R be the mean absolute SHAP value of feature j in regime R and let $q_j^R = s_j^R / \sum_l s_l^R$ be the normalised attribution share. We define Regime SHAP Divergence as the Jensen-Shannon divergence RSD = $JS(q^{normal} \| q^{shock})$, reported both at feature level and after aggregation to blocks. RSD is zero when a model explains shock and normal quarters identically and grows as the model reallocates attribution mass; it is bounded in [0,1] with base-2 logarithms, symmetric, and requires no re-estimation. We complement it with per-feature attribution ratios and rank shifts, and with a SHAP-implied rate elasticity obtained by regressing the quarterly mean of the summed SHAP contribution of all rate-related features on the observed four-quarter rate change.

4.5 Counterfactual Stress Engine

For a parallel shock of b basis points we map $x \to x(b)$ consistently rather than perturbing a single column: the yield level, its one- and four-quarter changes, the real rate and the shock-intensity indicator all move; the credit-spread proxy widens by 25 bp per 100 bp; the refinancing gap and refinancing pressure move with the new market rate; and all four exposure interactions are recomputed. Fundamentals are held fixed, so the exercise isolates the instantaneous rate channel. We evaluate the ladder $b \in \{0, 50, \ldots, 300\}$ on the last observed cross-section (2022Q4, 219 surviving REITs) and audit monotonicity by counting firms and ladder steps for which the predicted probability falls as the shock grows.

4.6 Evaluation Protocol

Because the labels look four quarters forward, naive cross-validation leaks. We use five expanding-window folds with a four-quarter embargo: training ends at 2008Q4, 2011Q4, 2014Q4, 2017Q4 and 2020Q4, and the corresponding test windows are 2010Q1-2012Q4, 2013Q1-2015Q4, 2016Q1-2018Q4, 2019Q1-2021Q4 and 2022Q1-2022Q4, so no training label uses data from a test quarter. Pooling the five test windows gives 12,544 out-of-time observations with 1,530 dividend events and 890 risk events. We report AUC, average precision (AP), the Brier score, expected calibration error (ECE, ten bins), the Kolmogorov-Smirnov statistic and recall in the top predicted decile. Differences in AUC are tested with a paired bootstrap (300 resamples) on the pooled predictions. Baselines are a class-balanced L2 logistic regression, a random forest, XGBoost, plain LightGBM and a two-output multi-task neural network, all given the identical 51 inputs.

5 RESULTS

5.1 Out-of-Time Predictive Performance

Table 2 reports pooled out-of-time results. RS-XMTL attains AUC 0.7938 on dividend sustainability and 0.9041 on financial risk. Against the strongest tree baseline the paired bootstrap gives dAUC = +0.0603 (XGBoost) and +0.0643 (LightGBM) on the dividend task and +0.0602 and +0.0570 on the risk task, all with $p < 0.001$. The class-balanced logistic benchmark is competitive on ranking - it is beaten by 0.0160 [0.0080, 0.0247] on the dividend task but beats RS-XMTL by 0.0073 [-0.0145, 0.0008], $p = 0.067$, on the risk task, a difference that is not statistically distinguishable from zero. Its probabilities, however, are unusable: ECE is 0.3229 and 0.1303 against 0.0625 and 0.0550 for RS-XMTL, and the Brier score is 0.2562 against 0.1007 on the dividend task. Class weighting is standard for rare events but inflates predicted probabilities, and we report both facets rather than only the favourable one. Figure 3 shows ROC curves and the reliability diagram; RS-XMTL lies close to the diagonal while the logistic benchmark is displaced far above it.

Table 2 Pooled Out-of-Time Performance over Five Expanding-Window Folds (12,544 Firm-Quarters; 1,530 Dividend Events, 890 Risk Events)

	Dividend sustainability (y_div)						Financial risk (y_risk)					
Model	AUC	AP	Brier	ECE	KS	Rec@10%	AUC	AP	Brier	ECE	KS	Rec@10%
Logistic (L2, balanced)	0.7778	0.3483	0.2562	0.3229	0.4184	0.354	0.9114	0.6200	0.0954	0.1303	0.6894	0.669
Random forest	0.7057	0.2135	0.1160	0.0996	0.3543	0.178	0.8654	0.5894	0.0704	0.1206	0.5472	0.569
XGBoost	0.7335	0.2426	0.1104	0.0652	0.3490	0.237	0.8439	0.4429	0.1116	0.1144	0.5488	0.390
LightGBM	0.7295	0.2319	0.1194	0.0842	0.3611	0.222	0.8471	0.4491	0.1222	0.1222	0.5718	0.394
Multi-task network	0.6235	0.1667	0.1617	0.1574	0.2092	0.124	0.8545	0.2943	0.1236	0.1230	0.5587	0.447
RS-XMTL (proposed)	0.7938	0.3521	0.1007	0.0625	0.4560	0.346	0.9041	0.4659	0.0669	0.0550	0.7047	0.646

Note: Higher is better for AUC, AP, KS and recall in the top decile; lower is better for the Brier score and ECE. Best value in each column is achieved by RS-XMTL except where the class-balanced logistic model or the random forest is noted in the text.

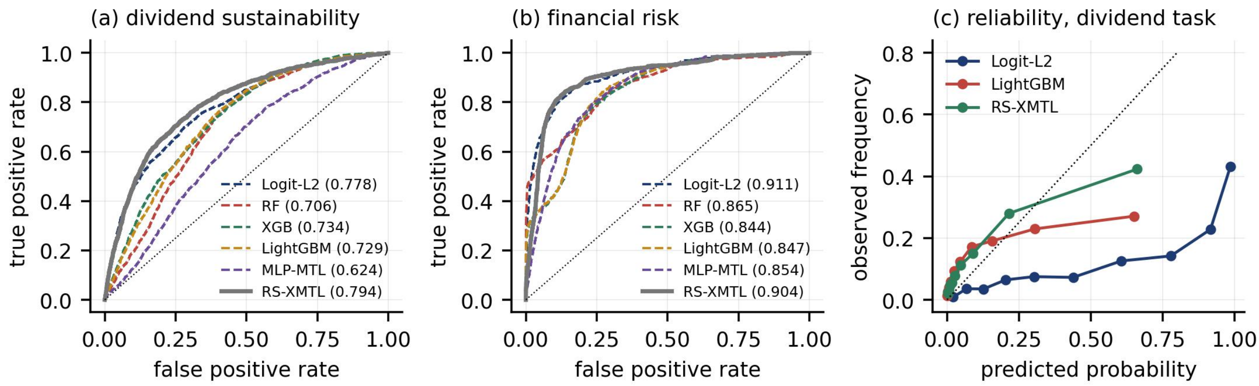

Figure 3 Pooled Out-of-Time ROC Curves for (a) Dividend Sustainability and (b) Financial Risk, with AUC in Parentheses, and (c) a Decile Reliability Diagram for the Dividend Task
Note: The class-balanced logistic model ranks well but is severely over-confident.

5.2 Ablation and the Role of the Sign Restrictions

Table 3 isolates each component. Removing the monotone constraints is by far the most damaging change, costing 0.0267 AUC on the dividend task and 0.0257 on the risk task (both $p < 0.001$). This is initially surprising, since constraints shrink the hypothesis class; unconstrained boosting on this panel does achieve higher in-fold flexibility, but it overfits regime-specific correlations - for example between high yields and strong nominal NOI growth - that reverse out of time. Cross-task transfer and blending both contribute, more strongly on the dividend task (0.0059 and 0.0071 AUC) than on the risk task, where their effect on ranking is negligible but their effect on calibration is not.
The rate-shock exposure block behaves differently from the other components, and we report this honestly: pooled over all quarters, dropping block X slightly improves AUC on the dividend task (0.8076 versus 0.7938). Splitting by regime explains why. In shock quarters the full model is better on both metrics (AUC 0.8031 versus 0.7963, AP 0.4194 versus 0.4055), whereas in normal quarters, which are 90% of the sample and where the interaction features are close to zero and act as noise, the reduced model wins. Block X therefore buys shock-regime precision at a small cost in tranquil periods - the trade-off one wants in a stress-testing tool, but not one that a single pooled statistic reveals. On the risk task the block is neutral (dAUC -0.0006, $p = 0.56$).

5.3 Where Predictability Lives

Fold-level results show that discrimination is far from uniform. On the risk task RS-XMTL reaches AUC 0.974, 0.937 and 0.967 in the three quiet folds but only 0.829 and 0.599 in the COVID and 2022 folds; on the dividend task the ordering is reversed, with AUC rising from 0.593 in 2010-2012 to 0.830 in 2016-2018 before falling back to 0.669 in 2022. Two effects are at work. Covenant breaches are driven by slow-moving balance-sheet ratios and are easy to predict in normal times; dividend cuts in quiet periods are dominated by idiosyncratic tenant and asset events that are unpredictable by construction. Conversely, the 2022 window puts the rate state outside the training support, and this is where the extrapolating component earns its place. We regard the 2022 degradation as the most important open problem rather than as noise.

5.4 Regime-Conditional Attribution

Figure 4 and Table 4 report the attribution analysis. Feature-level Regime SHAP Divergence is 0.0450 for the dividend model and 0.0726 for the risk model, so the risk model reorganises its explanation substantially more when rates jump. At block level the shift is systematic in both models: the macro/rate block gains 7.4 and 9.5 percentage points of attribution mass in shock quarters (27.3% to 34.7% and 39.8% to 49.3%), the cross-task block roughly doubles (6.6% to 13.9% and 5.0% to 9.3%), and the fundamentals block gives up 12.1 and 9.2 points. In other words, when rates move the model stops reading the balance sheet in isolation and starts reading the balance sheet through the macro state and through the companion risk signal.
Individual rank shifts are interpretable in exactly the terms of Section 2.1. In the dividend model refinancing pressure rises from rank 39 to rank 18, the four-quarter change in interest coverage from 37 to 17, and the real rate from 26 to 11. In the risk model the change in coverage rises from 44 to 16, the change in leverage from 37 to 17 and leverage itself from 31 to 20. These are the refinancing and coverage channels becoming binding, recovered by the attribution method rather than imposed. Figure 5(b) plots the quarterly mean SHAP contribution of the rate block against the observed four-quarter rate change and yields a SHAP-implied rate elasticity of 0.499 log-odds per 100 bp, i.e. a 100 bp rise in the 10-year yield adds about half a log-odds unit to the average REIT's dividend-cut score through the rate channel alone.

Table 3 Ablation of RS-XMTL, Pooled Out-of-Time; dAUC Is Relative to the Full Model with a Paired-Bootstrap p-value

Variant	AUC	dAUC	p	Brier	ECE	AUC	dAUC	p	Brier	ECE
RS-XMTL (full)	0.7938	-	-	0.1007	0.0625	0.9041	-	-	0.0669	0.0550
without monotone constraints	0.7671	-0.0267	<0.001	0.1002	0.0493	0.8784	-0.0257	<0.001	0.0798	0.0895
without cross-task transfer	0.7879	-0.0059	<0.001	0.1014	0.0613	0.9036	-0.0005	0.62	0.0647	0.0549
without two-stage blending	0.7867	-0.0071	<0.001	0.1006	0.0489	0.9039	-0.0002	<0.001	0.0642	0.0543
without exposure block X	0.8076	+0.0137	<0.001	0.0967	0.0547	0.9048	+0.0006	0.56	0.0638	0.0526

Note: Columns 2-6 are the dividend task, columns 7-11 the financial-risk task. The exposure block is the only component whose pooled effect on AUC is negative; its contribution is regime-specific (Table 5).

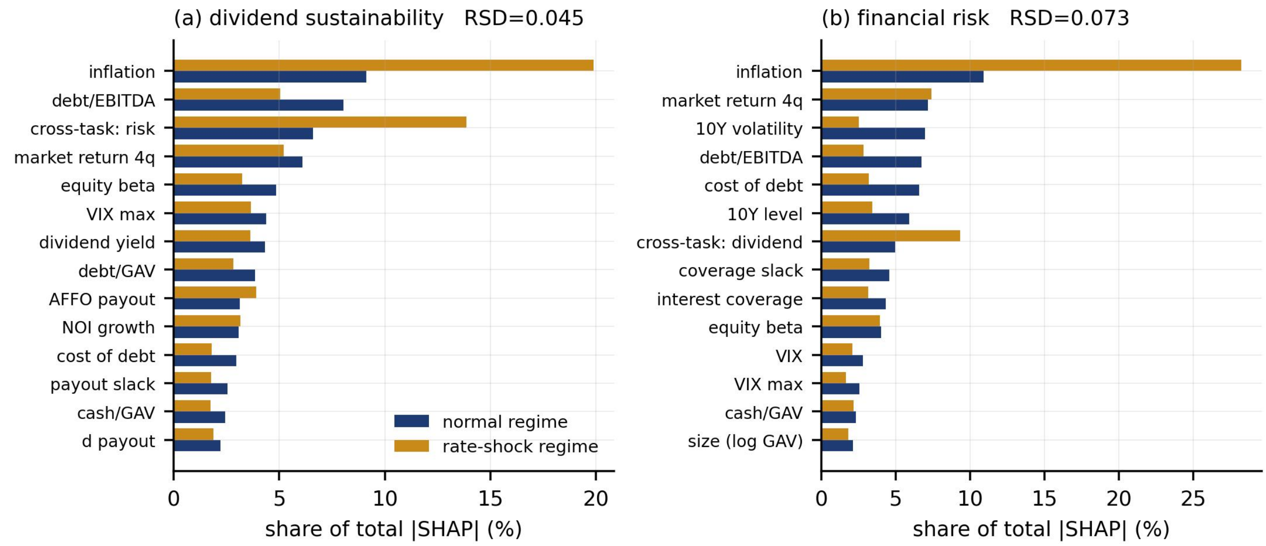

Figure 4 Share of Total Absolute TreeSHAP Attribution for the Fourteen Most Important Features, Split by Regime, with Feature-Level Regime SHAP Divergence Reported in the Panel Titles

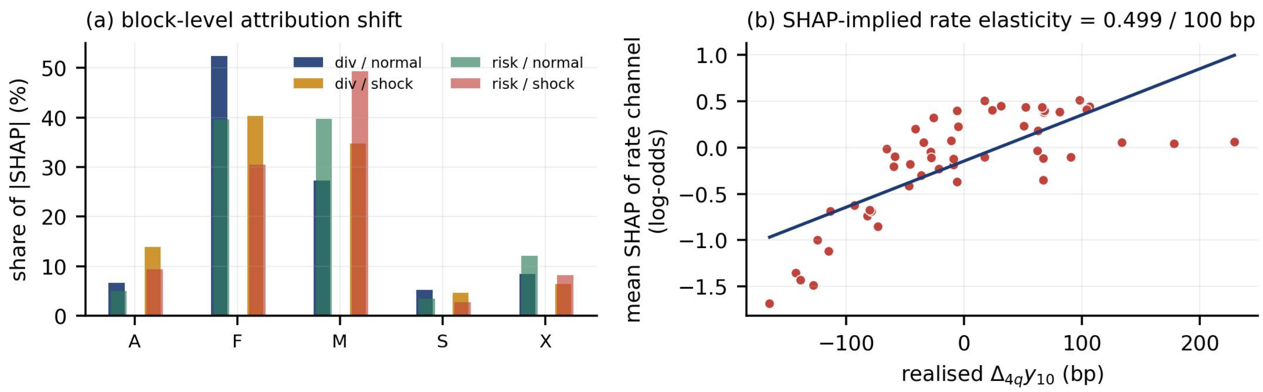

Figure 5 (a) Block-Level Attribution Shares by Regime and Task; (b) Quarterly Mean SHAP Contribution of the Rate Block against the Observed Four-Quarter Change in the 10-Year Yield

Note: (b) The fitted slope defines the SHAP-implied rate elasticity.

Table 4 Block-Level Attribution Shares by Regime and Regime SHAP Divergence

Block	normal	shock	shift	normal	shock	shift
F fundamentals	0.524	0.404	-0.121	0.397	0.305	-0.092
M macro / rates	0.273	0.347	+0.074	0.398	0.493	+0.095
X shock exposure	0.084	0.064	-0.020	0.121	0.082	-0.040
A cross-task	0.066	0.139	+0.073	0.050	0.093	+0.044
S sector	0.053	0.046	-0.006	0.035	0.027	-0.008
RSD (feature level)	0.0450			0.0726		
RSD (block level)	0.0195			0.0161		

Note: Columns 2-4 are the dividend model, columns 5-7 the financial-risk model. Shares are computed on 1,160 shock and 11,384 normal out-of-time observations.

5.5 Counterfactual Stress Test and Coherence Audit

Table 5 and Figure 6 report the stress test on the 2022Q4 cross-section. A parallel +300 bp shock raises the mean predicted probability of a dividend cut over the following year from 26.7% to 45.0%, and the share of REITs above a 25% probability threshold from 43.4% to 78.1%; the probability of a covenant-style risk event rises from 24.7% to 46.9%. The response is strongly ordered by rate exposure: the top exposure quartile, defined by the sum of the floating share and the one-year reset share, gains 24.0 percentage points of cut probability against 15.5 points for the bottom quartile. Sector patterns are economically sensible without having been imposed on the model: at +200 bp the mean cut probability is highest for hotel (62.2%) and office (49.2%) REITs and lowest for residential (33.5%).

The coherence audit is the sharpest result in the paper. Across the seven-point shock ladder the constrained model produces a non-decreasing response for every one of the 219 REITs and both tasks: zero firms and zero ladder steps with a sign violation. The otherwise identical unconstrained model violates monotonicity for 100% of firms, on 99.1% of ladder steps for the dividend task and 81.2% for the risk task, with a worst single-step fall of 39.7 percentage points. Its aggregate stress curve is not merely noisy but inverted: mean predicted cut probability falls from 21.2% at zero shock to 4.2% at +300 bp. A user of that model would conclude that a 300 bp tightening makes REIT dividends five times safer. Discrimination statistics alone would not have revealed the defect, since the unconstrained variant is only 0.027 AUC worse. This is direct evidence for the position of [33] that in high-stakes settings the hypothesis class should be restricted rather than the fitted black box explained after the fact.

Table 5 Counterfactual Parallel Rate Shock Applied to the 2022Q4 Cross-Section (219 REITs); Mean Predicted Probability in Per Cent

Shock (bp)	0	+100	+200	+300	0	+100	+200	+300
All REITs, constrained	26.7	43.0	43.1	45.0	24.7	39.8	39.9	46.9
All REITs, unconstrained	21.2	9.6	6.4	4.2	24.2	21.6	19.8	18.6
Exposure Q1 (low)	23.4	36.9	37.1	38.9	16.7	30.8	31.0	37.9
Exposure Q4 (high)	27.7	49.4	49.6	51.7	25.2	43.6	43.7	51.0
Share above 25% threshold	43.4	75.8	75.8	78.1	26.5	52.5	52.5	66.2
Firms with a sign violation	0.0	-	-	-	0.0	-	-	-
same, unconstrained model	100.0	-	-	-	100.0	-	-	-

Note: Columns 2-5 are the probability of a dividend cut, columns 6-9 the probability of a financial-risk event. Exposure quartiles are formed on the floating share plus the one-year reset share. The last two rows report the monotonicity audit over the full seven-point ladder.

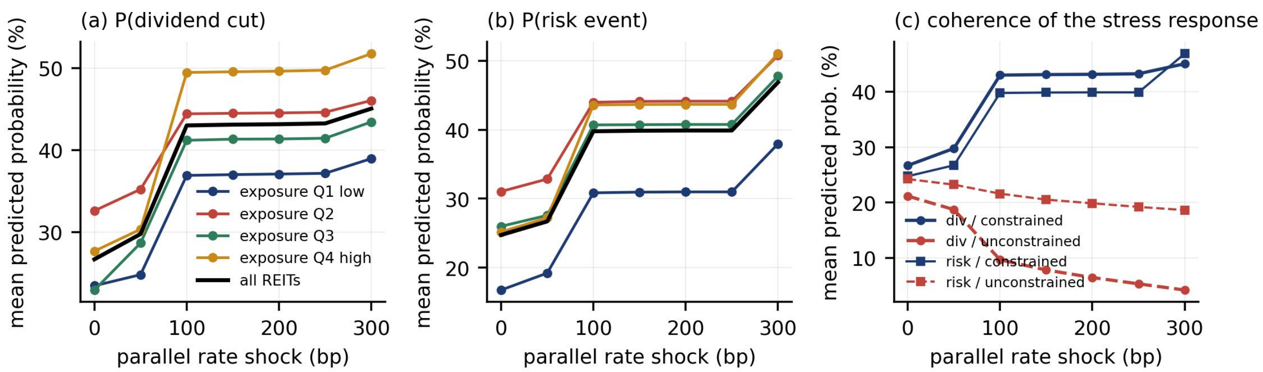

Figure 6 Stress Response of (a) the Dividend-Cut Probability and (b) the Risk-Event Probability by Rate-Exposure Quartile, and (c) the Coherence Comparison

Note: The constrained model is monotone by construction; the unconstrained model is inverted.

6 DISCUSSION AND LIMITATIONS

Three findings generalise beyond the present panel. First, sign restrictions are not a cosmetic interpretability device: on out-of-time data they improved AUC by roughly 0.026 on both tasks, because the unconstrained learner latches onto regime-specific correlations that reverse. Second, discrimination and coherence are close to orthogonal. The unconstrained variant loses only 0.027 AUC yet delivers an inverted stress surface, so a model selected on AUC alone can be worthless for scenario analysis. Any stress-testing pipeline should therefore report a monotonicity audit next to its ranking metrics. Third, regime-conditional attribution is informative in its own right: RSD and the associated rank shifts recovered the refinancing and coverage channels of the REIT literature from the fitted model rather than assuming them.

The limitations are substantial and we state them without hedging. (i) The firm panel is simulated. Although its exogenous drivers are observed series, its channels are taken from the literature and its cross-sectional moments are calibrated by hand, so the absolute performance levels reported here are properties of REIT-SIM and should not be read as estimates of what is achievable on Compustat or SEC data. What does transfer is the relative ordering of methods and the qualitative behaviour of the constrained and unconstrained stress surfaces. Replication on real firm data is the obvious next step and the released code accepts a real panel with no change to the model layer. (ii) A class-balanced logistic regression remains competitive on ranking, and on the risk task we cannot reject equality of AUC. Part of the reason is that the simulator's stress index, while containing threshold and conjunction terms, is closer to additive than reality is likely to be. (iii) Performance degrades in the 2022 fold, where the rate state leaves the training support; the sign-restricted linear component mitigates but does not solve this, and genuinely extrapolating architectures are needed. (iv) The credit-spread input is a volatility-based proxy rather than an observed spread. (v) The stress engine holds fundamentals fixed, so it measures an instantaneous rate channel and not the full dynamic response; a multi-period version would need to iterate the balance-sheet recursion.

7 CONCLUSION

We presented RS-XMTL, a rate-shock-aware explainable framework that jointly predicts dividend sustainability and covenant-style financial risk for U.S. REITs over a four-quarter horizon. Combining an engineered exposure block, end-to-end economic sign restrictions, out-of-fold cross-task transfer and a blender that doubles as a calibrator, the model reaches out-of-time AUC of 0.794 and 0.904 with expected calibration error below 0.07, dominating standard ensembles and matching a strong linear benchmark on ranking while improving calibration error by a factor of five. We introduced Regime SHAP Divergence to quantify how explanations move between tranquil and shock regimes, and showed that attribution mass migrates measurably toward macro and cross-task channels when rates jump. Finally, we showed that only the sign-restricted model produces an economically coherent stress surface: a +300 bp shock raises the mean dividend-cut probability from 26.7% to 45.0% monotonically for every REIT in the cross-section, while an unconstrained model of nearly identical AUC implies that higher rates make cuts less likely. Code, data-generation scripts and all result tables are released with the paper.

COMPETING INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

REFERENCES

[1] Chan K C, Hendershott P H, Sanders A B. Risk and return on real estate: evidence from equity REITs. AREUEA Journal, 1990, 18(4): 431-452.
[2] Mueller G R, Pauley K R. The effect of interest-rate movements on real estate investment trusts. Journal of Real Estate Research, 1995, 10(3): 319-326.
[3] Allen M T, Madura J, Springer T M. REIT characteristics and the sensitivity of REIT returns. Journal of Real Estate Finance and Economics, 2000, 21(2): 141-152.
[4] Bredin D, O'Reilly G, Stevenson S. Monetary shocks and REIT returns. Journal of Real Estate Finance and Economics, 2007, 35(3): 315-331.
[5] Cotter J, Stevenson S. Multivariate modeling of daily REIT volatility. Journal of Real Estate Finance and Economics, 2006, 32(3): 305-325.
[6] Boudry W, Coulson N E, Kallberg J G, et al. On the hybrid nature of REITs. Journal of Real Estate Finance and Economics, 2012, 44(1): 230-249.
[7] Bradley M, Capozza D R, Seguin P J. Dividend policy and cash-flow uncertainty. Real Estate Economics, 1998, 26(4): 555-580.
[8] Hardin III W, Hill M D. REIT dividend determinants: excess dividends and capital markets. Real Estate Economics, 2008, 36(2): 349-369.
[9] Boudry W I. An examination of REIT dividend payout policy. Real Estate Economics, 2011, 39(4): 601-634.
[10] Hayunga D K, Stephens C P. Dividend behaviour of US equity REITs. Journal of Property Research, 2009, 26(2): 105-123.
[11] Ghosh C, Sun L. Agency cost, dividend policy and growth: the special case of REITs. Journal of Real Estate Finance and Economics, 2014, 48(4): 660-708.
[12] Devos E, Spieler A, Tsang D. Elective stock dividends and REITs: evidence from the financial crisis. Real Estate Economics, 2014, 42(1): 33-70.
[13] Altman E I. Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. Journal of Finance, 1968, 23(4): 589-609.
[14] Ohlson J A. Financial ratios and the probabilistic prediction of bankruptcy. Journal of Accounting Research, 1980, 18(1): 109-131.
[15] Breiman L. Random forests. Machine Learning, 2001, 45(1): 5-32.
[16] Friedman J H. Greedy function approximation: a gradient boosting machine. Annals of Statistics, 2001, 29(5): 1189-1232.
[17] Chen T, Guestrin C. XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016: 785-794.
[18] Ke G, Meng Q, Finley T, et al. LightGBM: a highly efficient gradient boosting decision tree. In: Advances in Neural Information Processing Systems, 2017, 30: 3146-3154.

[19] Chawla N V, Bowyer K W, Hall L O, et al. SMOTE: synthetic minority over-sampling technique. Journal of Artificial Intelligence Research, 2002, 16: 321-357.
[20] Barboza F, Kimura H, Altman E. Machine learning models and bankruptcy prediction. Expert Systems with Applications, 2017, 83: 405-417.
[21] Fitzpatrick T, Mues C. An empirical comparison of classification algorithms for mortgage default prediction: evidence from a distressed mortgage market. European Journal of Operational Research, 2016, 249(2): 427-439.
[22] Bernanke B S, Kuttner K N. What explains the stock market's reaction to Federal Reserve policy?. Journal of Finance, 2005, 60(3): 1221-1257.
[23] Kuttner K N. Monetary policy surprises and interest rates: evidence from the Fed funds futures market. Journal of Monetary Economics, 2001, 47(3): 523-544.
[24] Hardin III W, Highfield M J, Hill M D, et al. The determinants of REIT cash holdings. Journal of Real Estate Finance and Economics, 2009, 39(1): 39-57.
[25] Riddiough T J, Wu Z. Financial constraints, liquidity management and investment. Real Estate Economics, 2009, 37(3): 447-481.
[26] Gilstrap C, Petkevich A, Sezer O, et al. REIT debt pricing and ownership structure. Journal of Real Estate Finance and Economics, 2022, 64(4): 546-589.
[27] Downs D H, Zhu B. Property market liquidity and REIT liquidity. Real Estate Economics, 2022, 50(6): 1462-1491.
[28] Gu S, Kelly B, Xiu D. Empirical asset pricing via machine learning. Review of Financial Studies, 2020, 33(5): 2223-2273.
[29] Bianchi D, Büchner M, Tamoni A. Bond risk premia with machine learning. Review of Financial Studies, 2021, 34(2): 1046-1089.
[30] Ribeiro M T, Singh S, Guestrin C. 'Why should I trust you?' Explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016: 1135-1144.
[31] Lundberg S M, Lee S-I. A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems, 2017, 30: 4765-4774.
[32] Lundberg S M, Erion G, Chen H, et al. From local explanations to global understanding with explainable AI for trees. Nature Machine Intelligence, 2020, 2(1): 56-67.
[33] Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, 2019, 1(5): 206-215.
[34] Slack D, Hilgard S, Jia E, et al. Fooling LIME and SHAP: adversarial attacks on post hoc explanation methods. In: Proceedings of AAAI/ACM Conference AI, Ethics, and Society, 2020: 180-186.
[35] Apley D W, Zhu J. Visualizing the effects of predictor variables in black box supervised learning models. Journal of the Royal Statistical Society: Series B, 2020, 82(4): 1059-1086.
[36] Niculescu-Mizil A, Caruana R. Predicting good probabilities with supervised learning. In: Proceedings of the 22nd International Conference on Machine Learning (ICML), 2005: 625-632.
[37] Fuster A, Goldsmith-Pinkham P, Ramadorai T, et al. Predictably unequal? The effects of machine learning on credit markets. Journal of Finance, 2022, 77(1): 5-47.