Heterogeneous Market Timing: Evidence from the Catastrophe Bond Market

Despoina Makariou[*], Pauline Barrieu[†], Yining Chen[‡]

February 10, 2025

Abstract

This paper investigates the heterogeneous outcomes of market timing decisions using a causal machine learning approach, offering more flexibility than traditional methods. Leveraging the diverse transaction profiles of the catastrophe bond market, we show that issuance timing materially affects primary market spreads, with earlier-issued catastrophe bonds priced more favorably. Capital availability in the reinsurance market and issuance size drive this variability, among other factors. Our findings highlight the importance of accounting for heterogeneity in market timing effects, providing new insights into how strategic timing decisions influence pricing dynamics, with broad implications for asset pricing, investment strategies, and risk management.

Keywords: market timing, security issuance, causal forest, catastrophe bonds

JEL Classification: G11, G12, G14, C50

[*]University of St Gallen, Email: despoina.makariou@unisg.ch
[†]London School of Economics and Political Science, Email: p.m.barrieu@lse.ac.uk
[‡]London School of Economics and Political Science, Email: y.chen101@lse.ac.uk

Introduction

Market timing effects can vary, raising the question of whether traditional econometric approaches adequately capture this variability. There are several[1] angles from which market timing can be interpreted, as evidenced by the extensive body of literature exploring its role across various financial contexts—from specific investment vehicles like mutual funds, hedge funds, venture capital and private equity to broader corporate finance decisions and market dynamics. Some examples of research in this direction, include but they are not limited to Evans et al. (2024), Busse et al. (2023), Glasserman et al. (2023), Cao et al. (2022), Dixon and Kelley (2022), Lewis and Tan (2016), Bolton et al. (2013), Ball et al. (2011), Chen and Liang (2007), Jiang et al. (2007), Bollen and Busse (2001), Fung and Hsieh (2001), Goetzmann et al. (2000), Ferson and Schadt (1996), Fung and Hsieh (1997).

However, a common assumption in much of the literature employing traditional econometric approaches, such as linear regression, is that the effects of market timing are homogeneous[2], un-

[1]For instance, in corporate finance, it often refers to the practice of issuing securities, such as equity or debt, based on expectations of market trends to achieve favourable financial outcomes (Gao and Lou, 2013).illustrating that In contrast, in investment contexts, market timing typically pertains to the ability to predict market movements and adjust trading strategies accordingly, which applies only to secondary market activities and is unrelated to security issuance. This research adopts the corporate finance perspective of market timing, although the core idea could be extended to explore other stances of market timing, such as those in investment strategy.

[2]To clarify, this notion of homogeneity aligns with the concept of causal effects, or treatment effects, which refer to the impact of a binary variable (the treatment) on an outcome variable of interest (the control variable). The treatment effect is estimated to determine how the treatment influences the outcome. When the magnitude and direction of the treatment effect are consistent across all units, regardless of other characteristics or subgroup differences, this is referred to as a homogeneous treatment effect model. On the contrary, a heterogeneous treatment effect model allows for varying effects across units or subgroups in a population and may provide a more accurate understanding of the impact, see Holland (1986).

less the sample groups or observational units for which these effects are examined are explicitly divided into distinct categories in advance to empirically capture some of the heterogeneity in the market timing impact. While such an approach can provide valuable insights into variations in effects, there are potential limitations. For instance, observed differences may be influenced by factors beyond the group characteristics, such as related variables. Additionally, this approach depends on the subjective choice of defining groups or interactions, which may limit the discovery of unexpected forms of variability. To address these challenges, semi-parametric and non-parametric methods, such as nearest-neighbour matching (Crump et al., 2008), kernel methods (Lee, 2009), Bayesian regression trees (Hill, 2011), and the causal forest (Athey et al., 2019), can be used to provide greater flexibility in estimating heterogeneous treatment effects. Causal forest, in particular, has demonstrated potential in finance and economics, with Gulen et al. (2024) employing the method to improve inferences on debt covenant defaults through heterogeneous estimates and Davis and Heller (2017) utilising it to identify subgroup differences in responses to employment policy interventions.

Given the advantages of estimating heterogeneous market timing effects using flexible methods like causal forest, markets with diverse transaction profiles provide an ideal setting for such analysis. The insurance-linked securities (ILS) market, which is tied to insured loss events, attracts large institutional investors such as pension funds, sovereign wealth funds, multi-asset investment firms, endowments, and some family offices. Within this market, catastrophe bonds— a prominent type of ILS— allow investors to indirectly participate in providing insurance coverage for catastrophic

events like natural disasters. Catastrophe bonds, issued by insurers, reinsurers, or governments, transfer a portion of the risk from catastrophic events to investors in exchange for periodic coupon payments, with the risk of losing some or all of the principal if predefined triggering events, such as hurricanes or earthquakes, occur. Though still relatively small compared to global bond and equity markets, the catastrophe bond market has seen notable growth (Swiss Re, 2024) reflecting its increasing role in risk diversification and its alignment with the rising demand for environmentally and socially responsible investments. However, the market remains opaque, with many transactions occurring over-the-counter. The complexity introduced by customised deal terms, varying risk profiles, and diverse pricing structures (Barrieu et al., 2025) creates a unique opportunity to explore heterogeneous market timing effects.

This study contributes to the broader market timing literature, where traditional models often assume homogeneous effects, limiting their ability to account for variability in market timing decisions. By leveraging the flexibility of the causal forest method, we identify these heterogeneous effects more effectively than standard econometric approaches. Additionally, our application—analysing market timing effects in the catastrophe bond market—is novel, as this area remains relatively sparse in both general finance market timing research and catastrophe bond-specific literature[3]. In this context, studying issuance timing effects on primary catastrophe bond spreads is crucial, as it directly impacts issuers' risk transfer costs and investors' return potential and risk assessment.

[3]For a broader discussion on market timing, see the first paragraph of this section. A more detailed review of catastrophe bond-related literature can be found in Barrieu et al. (2025).

In this research, we examine how catastrophe bond issuance timing affects primary market spreads and what drives heterogeneity in these effects. Using the causal forest method, we first establish that traditional models, which assume uniform effects, are inadequate in this setting. We then identify and rank the key factors shaping heterogeneity in issuance timing effects, assessing how they are influenced by capital availability in the reinsurance and ILS markets, as well as broader market conditions. To support our analysis, we construct a comprehensive dataset, integrating catastrophe bond pricing, structural characteristics, and market sentiment indicators with financial variables from the reinsurance, ILS, and broader capital markets.

Our findings highlight the importance of capturing heterogeneity in market timing effects, revealing how strategic issuance timing influences catastrophe bond pricing dynamics. We show that bonds issued earlier in the year tend to have lower spreads, with this effect varying based on capital availability in the reinsurance and ILS markets and issuance size. Traditional econometric models assume uniform market timing effects, limiting their ability to detect these variations. By applying a flexible, data-driven approach, we uncover nuanced issuance timing patterns that standard methods overlook, demonstrating that issuance decisions are shaped by complex, market-dependent factors rather than following a fixed effect. These insights are critical for both issuers and investors. Issuers can time issuance strategically to lower risk transfer costs, while investors gain a deeper understanding of pricing dynamics and market conditions, improving risk assessment and capital allocation. By demonstrating the heterogeneous nature of market timing effects, our study advances the understanding of market timing in asset pricing and underscores the neces-

sity of flexible estimation techniques. This framework extends beyond catastrophe bonds, offering insights into other financial markets where issuance timing plays a crucial role in price formation.

The remainder of this paper is organised as follows. Section I. outlines the methodology. Section II. presents the empirical analysis and results, covering an overview of the data, testing the homogeneity assumption in market timing, and applying the causal forest method to examine heterogeneous market timing in the context of catastrophe bond issuance, followed by a discussion of the findings. Finally, Section III. concludes the paper.

I. Methodology

In this section, we outline the causal forest methodology, detailing its core principles, definition, and an associated measure used to identify variables driving heterogeneity in treatment effects (Section I. A.). We also describe the causal forest training[4] process (Section I. B.) and explain how a trained causal forest is used to estimate heterogeneous treatment effects (Section I. C.) from an applied perspective. The mathematical aspects of causal forest can be found in Athey et al. (2019).

[4]Training a machine learning model refers to the process of using a dataset, known as the training dataset, to adjust the model's parameters or structure to improve predictions. The model learns patterns from the data, enabling accurate predictions on new, unseen data.

A. Causal Forest

Causal forest builds upon the random forest algorithm (Breiman, 2001), which is an ensemble of decision trees[5] that are randomised to improve prediction accuracy and generalisation for regression and classification tasks[6]. Causal forest adapts this approach by integrating methods from causal inference to estimate heterogeneous treatment effects, particularly the impact of a treatment or intervention on an outcome variable. Unlike random forest, which focus on minimising prediction error, causal forest aims to maximise the heterogeneity of treatment effects across different subgroups in a population. This difference affects how tree splits are made: in random forest, splits are chosen based on their ability to reduce prediction error, while in causal forest, splits are selected to enhance heterogeneity in treatment effects.

In what follows, we provide a more formal description of the causal forest method. We work with training data[7] which consist of tuples $Z_i = (X_i, Y_i, W_i)$ for $i = 1, ..., n$ with n being the number of observations. Here, $X_i = \{X_{i1}, X_{i2}, ..., X_{iP}\}$ is a feature vector containing P features and it is an element of $\mathbb{R}^P$, $Y_i \in \mathbb{R}$ is the response and $W_i \in \{0, 1\}$ is the treatment assignment

[5] A decision tree is a type of predictive model that organises data by making a series of decisions. It begins with a "root node," representing the starting point containing the entire dataset. At each step, the data is split into smaller groups, referred to as "child nodes," based on specific criteria related to input features. This splitting process continues until the data reaches "leaf nodes", where no further divisions occur, and the final prediction or decision is assigned. The root node sets the direction, while the branches and nodes represent decision points along the way. A widely used method for constructing decision trees is CART (Classification and Regression Trees), introduced by Breiman et al. (1986).

[6] Regression tasks involve predicting a continuous outcome variable, while classification tasks involve predicting a discrete category or class. In both cases, decision trees in random forests are used to partition data and make predictions based on observed patterns in the features.

[7] Training data refers to the dataset used during the training process to help the model learn patterns and relationships for making accurate predictions.

which takes value 1 if treatment is assigned and 0 otherwise. We aim to estimate the heterogeneous conditional average treatment effect $\tau(x) = \mathbb{E}[Y_i^{(1)} - Y_i^{(0)} | X_i = x]$ at a pre-specified test point x. The potential outcomes $Y^{(1)}$ and $Y^{(0)}$ represent the two cases with or without receiving the treatment. Since only one outcome can be observed under each scenario of receiving the treatment or not, we cannot estimate $\tau(x)$ directly from the observed data Z_i. We build a causal forest CF, as shown in the equation below, by averaging estimates τ obtained by training honest causal trees Γ over random sub-samples $\mathcal{I} \in \{i = 1, ..., n\}$ of size $|\mathcal{I}| = s$ whilst $\xi \sim \Xi$ is a source of auxiliary randomness.

$$CF(x; Z_1, ..., Z_n) = \binom{n}{s}^{-1} \sum_{1 \leq i_1, i_2, ..., i_s \leq n} \mathbb{E}_{\xi \sim \Xi}[\Gamma(x; Z_1, ..., Z_n)]$$

The term honesty here means that a causal tree is not permitted to look at the responses Y_i when deciding where to make a split, however it can look at the treatment assignments W_i in order to reduce bias in tree predictions (Athey et al., 2019).

A common way to quantify the impact of each feature X_i on the heterogeneity of treatment effects is to compute a weighted average of how many times a given predictor was split on at each depth in the causal forest, see Athey et al. (2019), where a higher weight indicates higher importance. This is the approach we adopt here, although alternative methods exist, as discussed in Hines et al. (2022), Boileau et al. (2023), and Bénard and Josse (2023) for instance.

B. Training

The causal forest is an ensemble of decision trees, so when describing the training process we primarily focus on the development of the individual trees. In the beginning, a random sub-sample is drawn from the full training dataset through sampling without replacement[8]. The first node of the tree, called the root node, encompasses this sub-sample. This root node is recursively split into child nodes, and the process continues until further splits are no longer possible, resulting in leaf nodes. The construction of child nodes depends on how splits are determined. At each level, in the causal tree construction process, a random subset of covariates is selected, and for each covariate, all possible values are considered as potential split points. The best split is chosen based on its ability to increase the heterogeneity of treatment effects within the child nodes—the split that maximizes the difference in treatment effects. Instances with predictor values greater than or equal to the split value are assigned to the right child node, while the rest go to the left. If a split does not improve the model fit, the node becomes a leaf node.

The causal forest algorithm requires the selection of several hyperparameters[9] similar to those in the random forest algorithm (Breiman, 2001). These include the number of trees (typically 2,000), the number of variables considered at each split, and the minimum leaf node size, which determines when the tree-building process terminates. In the causal forest, this minimum size ensures that each node has sufficient treated and control samples, enabling balanced treatment effect

[8]Sampling without replacement refers to the process where each data point is selected only once. This ensures each tree is trained on a unique subset of the data, which reduces overfitting and promotes diversity among the trees.

[9]Hyperparameters are key settings that guide the learning process and are tuned to optimize performance.

estimation (Athey and Wager, 2019). Causal forests also introduce additional hyperparameters, not found in the original random forest, that influence the model's honesty[10] (Athey et al., 2019). One such parameter, α, controls the maximum allowed imbalance for splits, penalizing splits that are too imbalanced to avoid bias. The specific hyperparameters used in this study will be discussed in Section II. C..

C. Estimation

The causal forest provides individualised treatment effect estimates, revealing how treatment impacts vary across observations and allowing researchers to identify key drivers of variability in treatment outcomes. To estimate treatment effects at specified test points[11], the causal forest follows the following process: each test point is passed through a causal tree to identify its corresponding leaf. The treatment effect is estimated by taking into account the treatment status of all observations in that particular leaf. This effect is aggregated across all trees to produce a final estimate for each test point. Here, we use out-of-bag[12] (OOB) data as test examples to compute estimates for each causal tree.

[10] Honesty refers to the use of different data for splitting nodes and populating the leaf nodes, reducing bias in the treatment effect estimates.

[11] A test point refers to an observation not used during the training process of a machine learning algorithm and is typically used for validation or performance evaluation.

[12] Out-of-bag (OOB) examples are data points excluded from the random sub-sample used to train a specific tree within a forest-based algorithm. These examples serve as a validation dataset for evaluating model performance during training; see Breiman (2001).

II. Empirical Analysis and Results

The empirical analysis and results are outlined in Section II. In particular, Section II. A. introduces the catastrophe bond dataset, accompanied by summary statistics. In Section II. B., we assess the suitability of a traditional econometric method to address the market timing problem in the context of catastrophe bond issuance. Section II. C. outlines the development of the causal forest model using the catastrophe bond data, while Section II. D. presents the results of the estimation in terms of average and heterogeneous treatment effects. Finally, Section II. E. delves into an analysis of the factors driving variation in market timing effects.

A. Data

This study uses a dataset compiled from both private and public sources, with the aim of representing, to the best of available knowledge, all nonlife catastrophe bonds issued in the primary market between December 2009 and May 2018 totaling 934 transactions. The dataset draws on earlier work by Makariou et al. (2021) but expanded and refined to overall include, data on investors' returns, securitized risk loss potential, risk transfer design features, credit ratings, risk modelling firms involved, issuance timing, sponsor's reputation related details, risk modelling company details, capital availability in the ILS and reinsurance markets, and returns in the BB corporate bond market. The exact variables used in our study along with their definitions and role in our study are described in Table 1. Summary statistics for all variables are presented in Table 2, Table 3, and

Figure 1. Some clarifications on the choice of certain variables follow.

The issuance timing in the catastrophe bond market–the treatment effect variable–is represented by whether a catastrophe bond was issued in the first or second half of the year. This approach was chosen for several reasons. Firstly, splitting the year in half provides a simple and interpretable framework for comparing outcomes across a well-defined temporal division. Secondly, this division allows us to account for seasonal factors that vary between the two periods, such as the major reinsurance contract renewal season in the first half of the year and the hurricane season in the second half. Finally, balancing the two classes of the treatment variable is important for robust estimation of treatment effects. In our dataset, 58% of catastrophe bonds were issued in the first half of the year, while 42% were issued in the second half, ensuring a reasonable balance. Details on some of the variables used in our study follow.

The catastrophe bond spread is the response variable given its industry importance for pricing purposes. It reflects the compensation that investors could earn on the top of the risk free rate which in a catastrophe bond context reflects the return on low-risk collateral, typically invested in safe assets like Treasury bills. The spread above this rate compensates investors for the risk of the bond, such as the likelihood of a natural disaster triggering a payout. The spread at issuance is here calculated, by subtracting the money market rate component from the coupon (or implied coupon in the case of zero coupon bonds).

Incorporating whether a catastrophe bond issuance covers a peak region peril combination, such as US wind related exposures, is important because it indicates the level of risk associated

with the catastrophe bond for the investors. Peak perils are typically linked to more understood perils whilst non peak perils to more niche and less understood exposures. There is no standard definition but the industry tend to think as peak perils being US wind and earthquake, Japanese wind and earthquake and Northern European wind, see Lloyd's (2024). The inclusion of peak perils is critical in defining the risk-return profile of the catastrophe bond, helping investors assess whether the potential returns justify the increased risk of loss.

Moreover, sponsor reputation can be considered a key market sentiment factor. A well-known and reputable sponsor (e.g., Swiss Re) may attract more investors, driving higher demand and resulting in lower spreads. This reputation serves as a market perception factor, influencing investor confidence and their willingness to engage in the catastrophe bond issuance, therefore we include it in our analysis.

TABLE 1. Overview of variables used in the analysis.

Variable	Definition	Type	Role
catbond_spread	Represents the amount of interest earned above the risk-free rate reflecting the return on low risk collateral such as US Treasury Bills.	Continuous	Response
early_year_issue	Indicates whether a catastrophe bond was issued during the first six months of the year.	Categorical	Treatment
AP	(**A**ttachment **P**robability). Reflects the likelihood of losses exceeding the attachment point. For parametric triggers, AP represents the probability of measured parameters surpassing the trigger point.	Continuous	Covariate
BB_spread	U.S. High Yield BB Option-Adjusted Spread, calculated as the difference between a BB-rated yield index and the Treasury spot curve.	Continuous	Covariate
coverage	Indicates if protection applies to a series of loss events or a single event.	Categorical	Covariate
EL	(**E**xpected **L**oss). The expected annual loss for a given layer, divided by the layer size.	Continuous	Covariate
rating	A dummy variable that identifies the credit rating of the bond or if it is unrated.	Categorical	Covariate
loc_peril	Represents a location-peril combination.	Categorical	Covariate
RoL	(**R**ate **o**n **L**ine). Quarterly values from the Lane Financial LLC Synthetic Rate on Line Index, reflecting rate levels in the ILS and reinsurance markets.	Continuous	Covariate
size	The nominal amount of the catastrophe bond.	Continuous	Covariate
term	The duration from issuance to maturity, in years.	Continuous	Covariate
trigger	The mechanism used to activate loss payments.	Categorical	Covariate
vendor	The catastrophe risk modelling software provider.	Categorical	Covariate
peak_peril	Indicates whether the issuance covers an industry standard or a more exotic region-peril combination (Lloyd's, 2024).	Categorical	Covariate
Swiss_Re	Indicates whether the catastrophe bond was issued by Swiss Re, a leading sponsor in many catastrophe bond issuances. It can serve as a proxy for the sponsor's reputation (Braun, 2016).	Categorical	Covariate

FIGURE 1. Histograms of continuous variables (defined in Table 1), showing the proportion of total observations (y-axis) as a percentage across ranges of values (x-axis) for each continuous variable in our study.

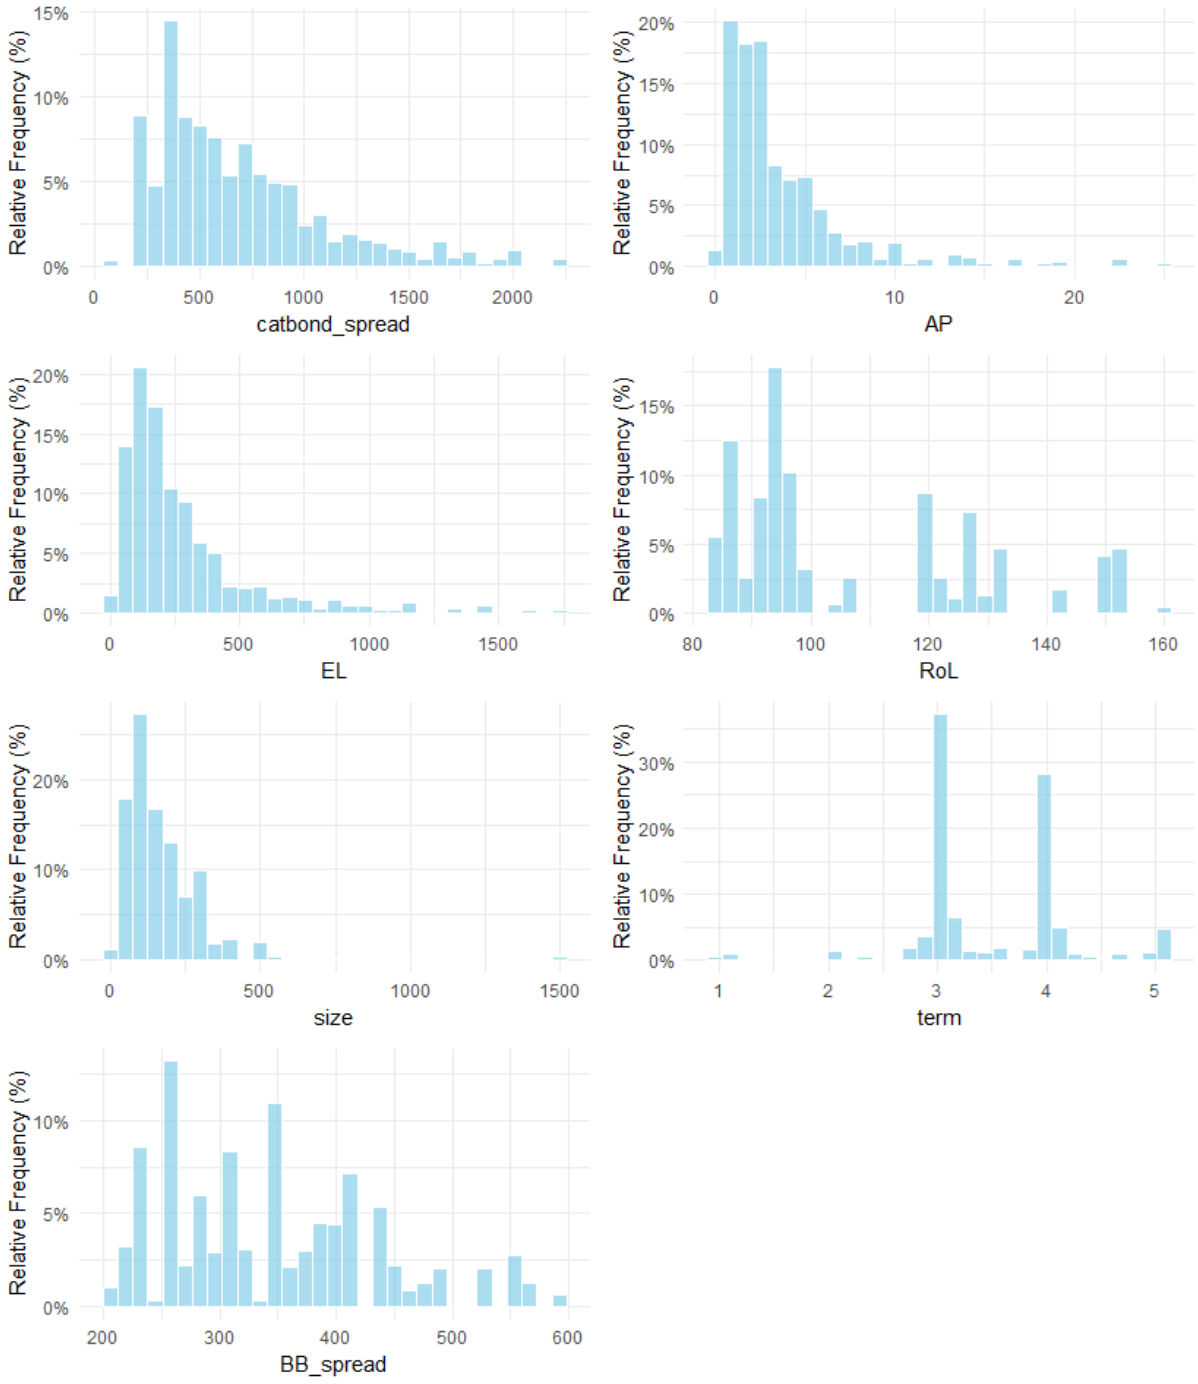


TABLE 2. Summary statistics for continuous variables (defined in Table 1), with units provided for each variable in parentheses. Basis points (bps) represent 1/100th of a percentage point, or 0.01%, commonly used to measure changes in financial metrics.

Variable (Unit)	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
catbond_spread (bps)	50.00	375.00	590.00	687.70	871.50	2200.00
EL (bps)	1.00	111.00	188.50	274.60	333.80	1735.00
AP (%)	0.02	1.36	2.51	3.72	4.68	25.04
RoL (bps)	83.57	91.73	96.06	106.97	124.57	159.73
BB_spread (bps)	212	261	341	346	414	598
size (USD million)	3.00	75.00	130.00	164.70	200.00	1500.00
term (years)	1.00	3.02	3.18	3.49	4.02	5.12

TABLE 3. Summary statistics for categorical variables (as defined in Table 1), presenting the number of observations and corresponding percentages for each level. Explanations on categorical variable levels not presented or not implied from the definitions given in Table 1 are provided below. Coverage types mainly include per occurrence (catastrophe bond payout is triggered by a single large loss) and aggregate (triggered by multiple events). With respect to loc_peril, abbreviations are explained in brackets as follows: APAC (Asia Pacific), Typh (Typhoon), NA (North America), SA (South America), Quake (Earthquake), Multi_Peri (various perils in listed regions). Credit ratings are based on S&P Global; others are converted to S&P Global range unless not rated (nr). With regards to trigger types: Indemnity (catastrophe bond payout is triggered based on actual loss level), Parametric (triggered by predefined event parameters such as wind speed or earthquake magnitude), Industry Loss Index (trigger based on total industry losses), Parametric Index (trigger which combines event parameters with an impact index), Model (trigger which is based on modelled losses), and Multiple (a combination of different trigger types). Finally, vendor levels include the following catastrophe risk modelling firms: RMS (Moody's RMS), AIR (AIR Worldwide), EQECAT, AON (Aon plc), or "pp" for unknown vendors due to private placement.

Categorical Variable	Level	No. of Observations	Percentage (%)
early_year_issue	Yes	542	58.03
	No	392	41.97
coverage	Aggregate	303	32.40
	Occurrence	627	67.10
	Both	4	0.50
loc_peril	APAC_Quake	51	5.46
	APAC_Typh	22	2.36
	Europe_APAC_Multi_Peril	2	0.21
	Europe_Quake	12	1.28
	Europe_Wind	54	5.78
	NA_APAC_Multi_Peril	26	2.78
	NA_Europe_APAC_Multi_Peril	36	3.85
	NA_Europe_Multi_Peril	39	4.18
	NA_Multi_Peril	425	45.50
	NA_Quake	80	8.57
	NA_Wind	184	19.70
	SA_Quake	3	0.32
rating	B	141	15.09
	BB	286	30.62
	BBB	4	0.43
	CCC	4	0.43
	nr (Not Rated)	499	53.43
trigger	Indemnity	511	54.70
	Parametric	29	3.10
	Industry Loss Index	325	34.80
	Parametric Index	23	2.50
	Model	22	2.40
	Multiple	24	2.60
vendor	AIR	741	79.30
	AON	4	0.40
	EQECAT	42	4.50
	RMS	141	15.10
	pp	6	0.60
peak_peril	Yes	905	96.89
	No	29	3.11
Swiss_Re	Yes	306	33.00
	No	628	67.00

B. On Modelling Market Timing using Linear Regression

Linear regression is a widely used method for estimating treatment effects, assuming homogeneity, meaning the effect is consistent across all observations in a dataset. While this may be suitable in some contexts, in others, such as catastrophe bond issuance timing, interactions with multiple bond attributes can be complex, making linear regression too restrictive. Without interaction terms, the model assumes a uniform effect, and even with interactions, capturing all potential relationships can be challenging. Given the potential for nonlinear and high-dimensional interactions, more flexible approaches, such as causal forests, may better suited to account for heterogeneity in issuance timing effects.

For illustration purposes, we build several linear models to evaluate whwther the homogeneity assumption holds for our analysis. We begin by estimating the effect of issuing a catastrophe bond in the first half of the year versus the second using a linear regression model, as shown in Table 4. Interaction terms are then introduced to examine whether issuance timing effects vary based on key catastrophe bond characteristics—such as expected loss, trigger type, vendor, and location-peril. To statistically assess these interactions, we perform an analysis of variance (ANOVA), which provides an empirical test for heterogeneity in the issuance timing effect. The ANOVA results in Table 5 show significant interaction terms, indicating that issuance timing effects vary across catastrophe bonds. This challenges the homogeneity assumption in linear regression and highlights the need for more flexible methods like causal forests to capture this heterogeneity more

effectively.

TABLE 4. In-sample fit of the linear regression model with multiple predictors. The coefficient of the treatment variable early_year_issue represents the average treatment effect for the full sample in this linear modelling framework.

	Estimate	**Std. Error**
Swiss_Re	6.649	(13.822)
RoL	8.093***	(0.388)
BB_spread	46.570***	(8.477)
term	−29.397***	(8.682)
size	0.000***	(0.000)
triggerindemnity	−23.661	(14.611)
triggermodel	−86.014**	(39.013)
triggermultiple	−37.079	(42.859)
triggerpar_index	−77.544*	(42.107)
triggerparametric	−217.160***	(36.694)
coverageboth	43.407	(82.489)
coverageoccurrence	−44.091***	(13.505)
vendorAON	150.164*	(88.099)
vendorEQECAT	−1.192	(34.027)
vendorPP	11.905	(71.535)
vendorRMS	45.034**	(18.671)
AP	−12.364**	(5.913)
EL	1.226***	(0.085)
loc_perilAPAC_Typh	−69.353	(43.242)
loc_perilEurope_APAC_Multi_Peril	−33.422	(126.002)
loc_perilEurope_Quake	−31.963	(75.991)
loc_perilEurope_Wind	−190.532***	(41.902)
loc_perilNA_APAC_Multi_Peril	52.829	(47.312)
loc_perilNA_Europe_APAC_Multi_Peril	95.890**	(40.964)
loc_perilNA_Europe_Multi_Peril	115.517***	(40.080)
loc_perilNA_Multi_Peril	148.571***	(27.868)
loc_perilNA_Quake	−47.749	(34.291)
loc_perilNA_Wind	89.258***	(30.176)
loc_perilSA_Quake	61.851	(118.470)
rate_1_excl_minplusBB	−134.694***	(18.441)
rate_1_excl_minplusBBB	−341.847***	(83.656)
rate_1_excl_minplusCCC	−9.984	(94.364)
rate_1_excl_minplusnr	−12.422	(18.643)
early_year_issue	−75.720***	(13.731)
peak_peril	−51.453	(54.719)
Constant	−439.530***	(93.603)
Observations	934	
R^2	0.850	
Adjusted R^2	0.844	
Residual Std. Error	161.592 (df = 898)	
F Statistic	145.506*** (df = 35; 898)	
Notes:	***Significant at the 1% level.	
	**Significant at the 5% level.	
	*Significant at the 10% level.	

TABLE 5. ANOVA results for testing the homogeneity of treatment effects. The first model is the baseline regression including all the predictors from Table 1. The remaining models include additionally an interaction term between issuance timing and another variable listed below.

Linear Model	Res.Df	RSS	Df	Sum of Sq	F	Pr(>F)
baseline	898	23,448,442				
EL	897	23,148,893	1	302,612	11.607	< 0.001 ***
trigger	893	22,846,433	5	602,009	4.706	< 0.001 ***
vendor	896	23,291,432	2	157,010	3.020	0.038 *
loc_peril	889	22,907,251	9	541,191	2.334	0.013 *

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1.

C. Causal Forest Development

The causal forest was developed based on the catastrophe bond dataset presented in Section II A.. Following the notation outlined in Section I A., we have $n = 934$ training examples and $P = 13$ predictors. Our treatment effect variable, W_i, is the binary indicator which takes the value 1 if a catastrophe bond is issued in the first half of the year, and 0 otherwise. The response variable, Y_i, represents the catastrophe bond spread at issuance, while the explanatory variables in $X_i = \{X_{i1}, X_{i2}, ..., X_{iP}\}$ are listed in Table 1.

With respect to the choice of causal forest hyperparameters, we primarily use the default values suggested by Athey et al. (2019), which have been empirically shown to perform well for our dataset too. These include a causal forest size of 2000 trees, an honesty parameter of 0.5, a maximum causal tree imbalance of 0.05, and a minimum node size of 5. The hyperparameter controlling the number of predictors to be considered at each split in the causal tree development process has been optimised using a cross-validation-based tuning process.

D. Estimation of Market Timing Effects using the Causal Forest

We present the estimated effects of catastrophe bond issuance timing on primary market spreads, derived using the causal forest method. Our analysis begins with an evaluation of the average[13] effects of issuance timing on spreads followed by comparison of these results to those obtained from the linear regression benchmark (see Table 4). We then investigate the heterogeneity of issuance timing effects by examining how they vary across different subgroups or characteristics within the catastrophe bond population. Notably, the causal forest method enables the identification of these subgroups without requiring prior knowledge.

D.1 Average Treatment Effects

Initially, we estimate the average treatment effect (ATE)[14] to assess the overall impact of early issuance timing on catastrophe bond spreads across the entire population of catastrophe bonds in the primary market. Next, we compute the average treatment effect on the treated sample (ATT)[15], which represents the treatment effect averaged only across catastrophe bonds issued in the first half of the year. ATT evaluates the direct impact of early catastrophe bond issuance on primary market spreads since it focuses on the subset of catastrophe bonds that were actually issued early

[13]The estimation of average treatment effects is based on a variant of augmented inverse propensity weighting described in Athey et al. (2019).

[14]In a randomised control trial setting, ATE is known as the overall effectiveness of the intervention.

[15]In a randomised control trial setting, ATT measures the effectiveness of the intervention for the part of the population who receive it.

in the year. Finally, we compute the average treatment effect on the control sample (ATC)[16], which estimates the average treatment effect for catastrophe bonds issued in the second half of the year, assuming they had been issued in the first half instead. Essentially, ATC shows the potential impact of early issuance timing on late-issued catastrophe bond spreads, answering the counterfactual question of what if these catastrophe bonds had been issued earlier.

The average issuance timing effect estimates obtained from the causal forest via ATE, ATT, and ATC are presented in Table 6. All measures are negative, indicating that early issuance timing generally reduces catastrophe bond spreads at issuance. The ATE estimate of -88.02 basis points reflects the average reduction in spreads across the entire market, including both treated and control groups, with a standard error of 8.247 suggesting reasonable precision. The ATT estimate of -92.607 is slightly more negative, highlighting a greater reduction in spreads for catastrophe bonds already issued early in the year (treated group) compared to the overall market. Meanwhile, the ATC estimate of -104.028 is the largest in absolute terms, suggesting that catastrophe bonds typically issued later in the year could have seen even greater reductions in spreads if issued earlier.

The differences between the ATC, ATT, and ATE estimates (ATC > ATT > ATE) signal heterogeneity in the treatment effect of issuance timing on catastrophe bond spreads. This indicates that the effect is not uniform across the population but varies by subgroup of the catastrophe bond population, even at this broad level (treated versus control). The smaller ATT estimate suggests that early-issued catastrophe bonds may have already incorporated the effects of early-year is-

[16]In a randomised control trial setting, ATC measures the average effect of a treatment on the control group—the units or individuals that did not receive the treatment.

suance timing, limiting further spread reductions. The larger ATC estimate implies that late-issued catastrophe bonds might have benefited more from earlier issuance, potentially due to variations in market dynamics or underlying factors affecting these bonds during different issuance periods.

TABLE 6. Average treatment effects derived from the causal forest for the impact of catastrophe bond issuance timing on their spreads. ATE (Average Treatment Effect) represents the average effect of early-year catastrophe bond issuance on spreads across the entire catastrophe bond population, including both treated (early-year issued catastrophe bonds) and control (late-year issued catastrophe bonds) units. It quantifies the expected change in spreads, measured in basis points, due to issuance timing across the entire market. ATT (Average Treatment Effect on the Treated) focuses on the treated group, measuring the effect of early issuance timing specifically on catastrophe bonds that were issued early in the year. It quantifies the direct impact of early issuance on this subset of catastrophe bonds. ATC (Average Treatment Effect on the Control) estimates the potential effect of early issuance timing on catastrophe bonds that were issued later in the year, offering insights into how spreads might have been impacted if these late-year catastrophe bonds had been issued earlier. Standard errors (std. err) are provided for each estimate.

Causal Inference Metric	Causal Forest Estimate	std.err
ATE	-88.017	8.247
ATT	-92.607	23.002
ATC	-104.028	13.089

For comparison purposes, we look into the linear regression benchmark with multiple predictors presented in Table 4. In this case, the ATE is estimated by including the treatment indicator variable and the relevant covariates as independent variables in the regression model and interpreting the coefficient of the treatment variable as the average treatment effect on the outcome variable, controlling for the other predictors and covariates included in the model. However, since linear regression does not allow direct estimation of subgroup-specific effects like ATT or ATC without additional stratification or interaction terms, any comparison between average treatment effects derived from the causal forest and linear regression alike is limited to ATE.

The ATE estimates from the causal forest and linear regression models align in the direction

of the issuance timing effect on catastrophe bond spreads but differ in magnitude: -88.017 for the causal forest versus -75.720 for linear regression. This, roughly, 12 basis point difference may stem from heterogeneity in the data or the causal forest's ability to capture non-linear patterns and complex interactions that linear regression overlooks. Although seemingly modest, a 12-basis-point difference can have implications in the catastrophe bond market. For issuers, a reduction in spreads directly translates into lower cost of capital, which can result in savings when issuing bonds with large principal amounts, often in the hundreds of millions of dollars. These savings could influence issuers' decisions about timing future issuances to optimise funding costs. For investors, even small variations in spreads can alter the relative appeal of catastrophe bonds in comparison to other fixed-income instruments, potentially affecting portfolio allocation strategies and risk-return profiles. These insights illustrate the practical value of the causal forest capacity to uncover differences in how issuance timing impacts catastrophe bond pricing at issuance, which might not be evident from simpler models like linear regression.

D.2 Heterogeneous Treatment Effects

The aggregate treatment effect measures—ATE, ATT, and ATC—provide valuable insights into the overall impact of issuance timing on catastrophe bond spreads and points to studying heterogeneity further. However, these metrics do not directly capture the heterogeneity of the treatment effect across different subgroups or characteristics within the catastrophe bond population. To better understand and identify such heterogeneity, we turn to the conditional treatment effects derived

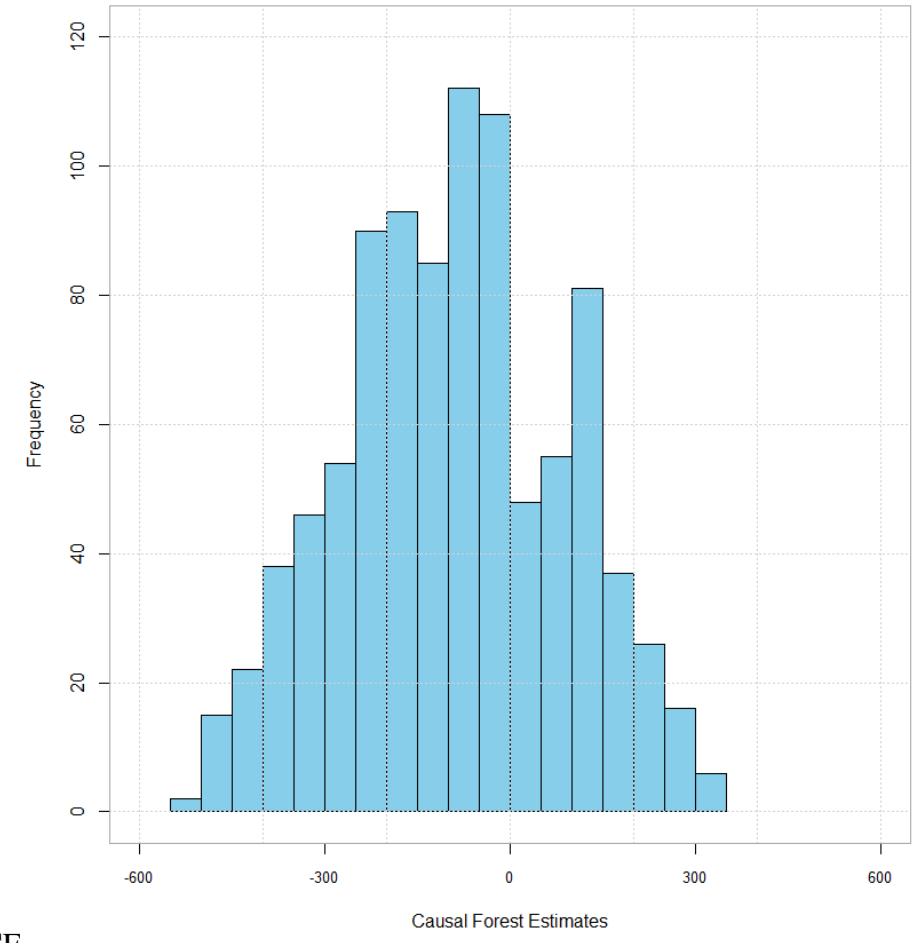
from the trained causal forest.

Unlike traditional linear regression, which requires the specification of interaction terms and can become computationally intractable as the number of variables and interactions increases, the causal forest method allows for the efficient estimation of treatment effect heterogeneity across various subgroups without explicitly modelling every interaction. For instance, with 14 independent variables (including 13 covariates and the treatment variable), considering second-degree interactions would result in 105 terms, and higher-order interactions would significantly increase this number, making the model computationally burdensome. The causal forest, by contrast, automatically accounts for complex interactions and non-linear relationships, providing a more scalable and effective way to capture treatment effect heterogeneity.

A visual representation of the varying impacts of issuance timing on catastrophe bond spreads,as derived from the causal forest, is provided in Figure 2. The histogram displays the distribution of spread outcomes across the catastrophe bond population after treatment-issuance in the first half of the year. The height of each bar reflects the frequency of catastrophe bonds experiencing specific spread outcomes. Although the overall trend suggests that early issuance is generally associated with lower spreads, the distribution reveals considerable variability in the estimated treatment effects. Notably, some subgroups of catastrophe bonds exhibit a positive treatment effect, indicating that for certain bonds, issuing earlier in the year results in higher spreads, or a "spread markup". That said, the median estimated conditional average treatment effect (CATE) is -85.098, whilst the minimum estimated CATE is -523.677 and the maximum one is 326.983. This illustrates the

complexity of how issuance timing impacts catastrophe bond pricing and emphasises the value of using the causal forest to reveal this heterogeneity.

FIGURE 2. Histogram of causal forest estimates of heterogeneous treatment effects for the impact of catastrophe bond issuance timing on spreads. Frequency distribution of estimated spread differences for early-year issuances versus the counterfactual.

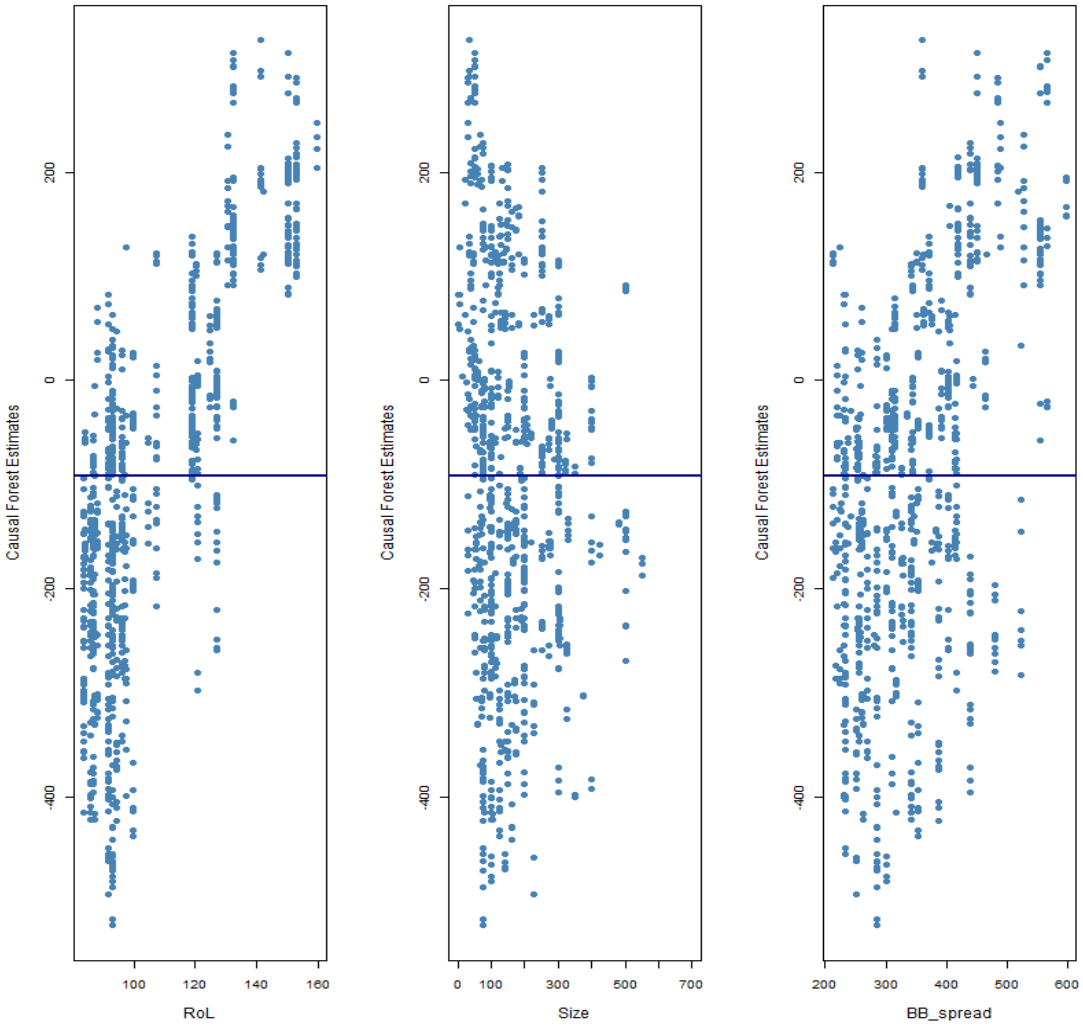


CF

Further understanding of this variability is essential for tailoring the timing of catastrophe bond issuances. Following the example of Willke et al. (2012), we examine the impact of heterogeneous treatment effects on "averages", as illustrated in Figure 3. Specifically, we plot the estimated

treatment effects on the y-axis against the values of some selected predictors on the x-axis for simplicity. The blue line represents the mean CATE as estimated by the causal forest. If treatment effects were homogeneous, they would be randomly distributed and independent of the catastrophe bond characteristic of interest. In such a scenario, the conditional estimated treatment effect would be more or less constant across all catastrophe bonds and independent of the values of other covariates. However, in Figure 3, we observe that the variation in treatment effects is not random and appears to be influenced by the distribution of the continuous covariates. As a result, a decision to issue a catastrophe bond in the first half of the year based on the average effect estimated for the entire sample may sometimes underestimate or overestimate the average effect for a specific subsample.

FIGURE 3. Exploring heterogeneity in causal forest treatment effect estimates for some covariates. For expository purposes, this figure illustrates the variabilty in estimated conditional treatment effects (CATE) for different continuous covariates, here RoL, size, and BB_spread, with the blue line representing the mean CATE, as derived from causal forest. The observed heterogeneity highlights that relying solely on the average treatment effect for the entire sample may result in inaccurate estimations for specific catastrophe bond subgroups, either underestimating or overestimating the issuance timing impact depending on the covariate values. A more in-depth analysis of the factors driving heterogeneity in catastrophe bond issuance timing effects is provided in Section III E..



E. Drivers of Variation in Market Timing Effects

We proceed by delving deeper into the specific input variables that contribute to the variations in the effects of issuance timing on catastrophe bond spreads identified by the causal forest.

E.1 Variable Importance Ranking

As mentioned in Section I. A., a potential way to calculate the impact of every covariate on the heterogeneity of treatment effects is to compute a weighted average of how many times a given predictor was split on at each depth in the causal forest, where a higher weight indicates higher importance. By using this approach, in Table 7, we show the most important variables, as derived by the causal forest, in determining the heterogeneity of issuance timing effects on catastrophe bond spreads.

At a high level, we observe that different predictors have non-equal contribution in explaining the differences in spreads at issuance between catastrophe bonds issued in the first half of the year and those issued in the second half. This finding further supports the notion that issuance timing impacts spreads heterogeneously. Were the effects homogeneous, all the variables would have exhibited comparable levels of importance in the ranking.

TABLE 7. Variable importance measures derived from causal forest. A weighted average of how many times a covariate was split on at each depth in the causal forest. The average importance per class is considered for categorical variables.

Variable	Importance
RoL	0.223
size	0.189
AP	0.157
EL	0.124
term	0.106
BB_spread	0.054
loc_peril	0.070
rating	0.049
trigger	0.008
coverage	0.008
Swiss_Re	0.006
vendor	0.003
Total	1.000

By delving into the variable importance rankings (Table 7), it is evident that market conditions in the ILS and reinsurance markets is the most important driver of heterogeneity in issuance timing effects on spreads. The state of the ILS and reinsurance markets is influenced by a complex interplay of evolving factors that affect market liquidity throughout the year, making it difficult to easily explain their impact on driving heterogeneity in the effect of issuance timing. Nevertheless, some timing-related factors, such as renewal seasons in the (re)insurance industry and reactions to post-catastrophe events, can provide a framework for interpreting these heterogenous effects. In particular, reinsurance renewal seasons, which occur at the beginning of the year (January 1st) and mid-year (June 1st or July 1st), influence the availability of capital for the following months, potentially affecting catastrophe bond spreads. Moreover, early in the year, fresh data from the previous year and the end of the US hurricane season can shape risk perceptions, influencing spreads

on early issuances. The second most important factor in driving heterogeneity in issuance timing effects is the catastrophe bond issuance size—the amount of coverage offered by the catastrophe bond. Issuing large bonds early in the year may be cheaper due to fresh investor capital, increased liquidity, and lower immediate risk perceptions as the main natural disaster seasons (such as US hurricanes) lie later in the year.

In addition to these top heterogeneity contributors, other factors like attachment probability and expected loss, which combined can reflect a catastrophe bond risk profile, also drive spread heterogeneity with respect to issuance timing. As the year progresses, the increased likelihood of catastrophic events during main disaster seasons widens spreads for later issuances. Other factors, such as alternative investment options (for example, BB-rated corporate bonds) and the sponsor's reputation, play a smaller role in influencing the heterogeneous effects of issuance timing.

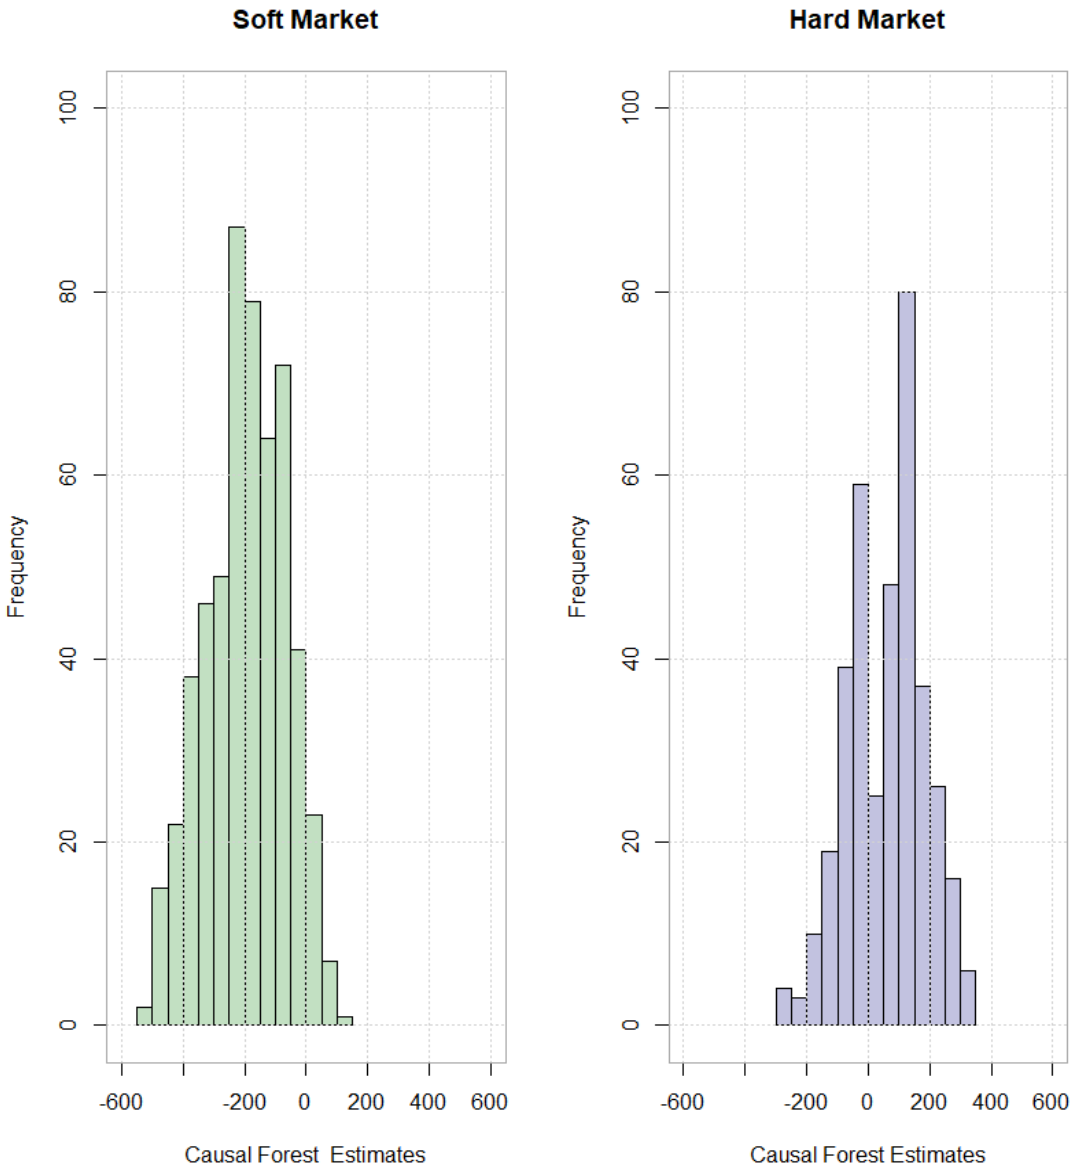
E.2 Investigating the Importance of Market Conditions

In the following section, we delve deeper into how market conditions influence the heterogeneous effect of issuance timing on catastrophe bond spreads at issuance. As shown in Table 7, the reinsurance and ILS market conditions are the primary drivers of heterogeneity in our analysis, represented by the Return on Line index (defined as RoL in Table 1).

According to industry standards, RoL values of 100 or below indicate a *soft* market environment, characterised by ample capital, flat or declining rates, increased competition, and more relaxed underwriting criteria. In contrast, RoL values above 100 signify a *hard* market, where the

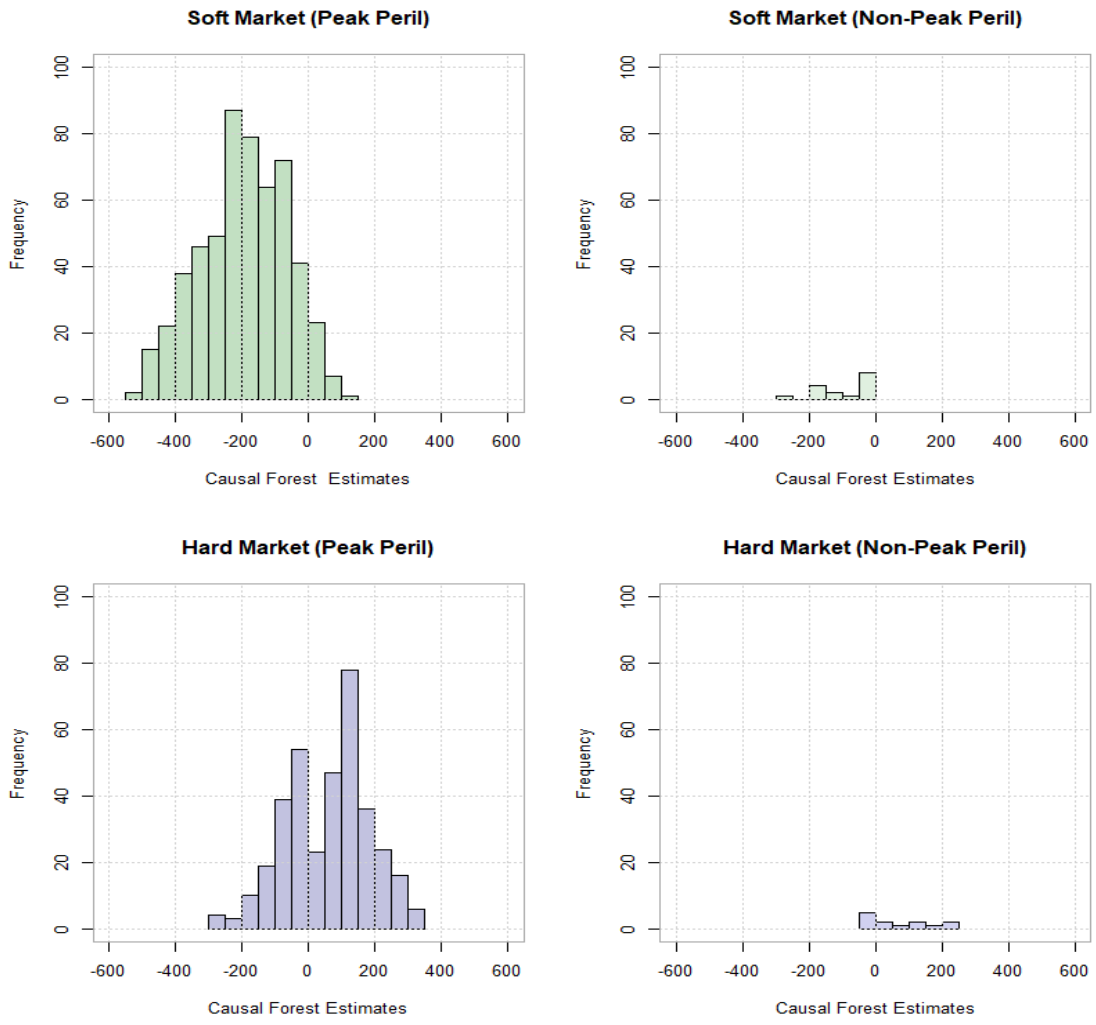
opposite conditions prevail. In our dataset, approximately 40% of catastrophe bonds were issued in hard markets, while 60% were issued in soft markets. Using this classification, we extract the causal forest's conditional treatment effect estimates, filter them by market condition (soft or hard), and create histograms to show the distribution of differences in catastrophe bond spreads based on issuance timing (early-year vs. late-year issuance), as presented in Figure 4. The findings suggest that the negative impact of early-year issuance on catastrophe bond spreads is more pronounced under soft market conditions than in hard market conditions. Additionally, in hard market periods, the relationship between issuance timing and spreads appears to be more variable. This implies that in soft markets, the timing of issuance plays a clearer role in determining spreads, whereas in hard markets when market participants face uncertainty, other factors may influence spreads more strongly.

FIGURE 4. Histogram of causal forest estimates of heterogeneous treatment effects for the impact of catastrophe bond issuance timing on spreads. Frequency distribution of estimated spread differences for early-year issuances versus the counterfactual, under soft (RoL $\leq$ 100) and hard (RoL > 100) market conditions.



In Figure 5, we examine how the timing of issuance (early or late in the year) affects catastrophe bond spreads under soft and hard market conditions for peak and non-peak perils. Due to the limited number of non-peak peril observations (Table 3), the findings should be interpreted as indicative trends. Investigating the interaction between peril type and market conditions helps reveal how both the nature of the risk and prevailing market dynamics influence the issuance timing effect on catastrophe bond spreads.

FIGURE 5. Histogram of causal forest estimates of heterogeneous treatment effects for the impact of catastrophe bond issuance timing on spreads. Frequency distribution of estimated spread differences for early-year issuances versus the counterfactual, under soft (RoL $\leq$ 100) and hard (RoL > 100) market conditions for peak and non-peak peril coverage.



In soft markets, the negative effect of issuance timing on spreads prevails, regardless of peril type. However, in hard markets, the effect of early issuance on spreads is more variable. For peak perils, early issuance can lead to either higher or lower spreads, but the latter is less frequent. For non-peak perils, early issuance consistently results in higher spreads in hard markets. This pattern is likely driven by capital scarcity and heightened risk perceptions in hard markets, where non-peak perils are seen as more uncertain and riskier. Investors, perhaps being cautious early in the year, may demand higher spreads for these risks, reflecting the lack of fully assessed risk data. Finally, it is noteworthy that although the peak peril covariate was not found to drive heterogeneity of issuance timing effects at all on its own (Table 7), its interaction with market conditions shows the power of the causal forest methodology in capturing complex relationships between factors influencing the impact of issuance timing on catastrophe bond spreads at issuance.

F. Interpretation

Our research findings reveal that issuance timing has a heterogeneous effect on the primary catastrophe bond pricing. We show that relying on traditional econometric methods would have not fully account for directly exploring such heterogeneity and we employ a causal machine learning method namely causal forest to capture these variations instead. Our results indicate that catastrophe bonds issued earlier in the year tend to have lower spreads, with an even stronger effect if late-year issuances had occurred earlier. The strength of this timing effect, however, is influenced by several factors such as capital availability in the reinsurance and ILS markets and

the issuance size. In particular, larger issuances, which attract institutional investors and enhance market liquidity, often reduce perceived risk and result in lower spreads. Similarly, periods of abundant capital availability amplify the timing related spread discount. Some of the additional factors driving the issuance timing heterogeneity is the probability of investors incurring initial loss on their investment and the expected loss of the underlying securitised disaster.

The insights derived from the causal forest approach can inform issuers and investors alike. For issuers, understanding how market liquidity and capital availability influence the issuance timing effects can lead to better decisions regarding when to enter the market. By aligning issuance timing with periods of abundant capital and stronger investor demand, issuers can take advantage of the pricing discounts that early-year issuances often offer. This could help reduce the cost of transferring catastrophe risk, an increasingly vital consideration in an era of escalating natural disasters driven by climate change.

For investors, understanding the heterogeneous effects of issuance timing enhances their ability to assess risk and pricing more effectively. Recognising that factors such as market liquidity, capital availability, and issuance size influence pricing dynamics allows investors to align their strategies with their specific objectives. For risk-averse investors, this knowledge can help identify opportunities, such as early-year issuances with typically lower spreads, while those seeking higher returns may focus on issuances offering greater yield potential. This improved understanding supports more targeted portfolio management and risk mitigation strategies, ultimately fostering a more efficient allocation of capital in the catastrophe bond market.

G. Broader implications

Building on these findings, we discuss the broader implications of failing to detect for heterogeneity in financial decision-making. Traditional econometric approaches often rely on uniformity assumptions or broad subgroup analyses, which can obscure critical variations in how market participants respond to changing conditions. In contrast, causal forests provide a data-driven, non-parametric approach that allows for the detection of non-linear interactions and previously unobserved heterogeneity.

The failure to account for heterogeneity in market timing can have important consequences for financial decision-making. In corporate finance, overlooking heterogeneous timing effects can lead to suboptimal issuance strategies, mispricing of risks, and inefficient capital allocation. Similar challenges exist in monetary policy and asset pricing, where uniform-effect assumptions may fail to capture market segmentation and investor behavior shifts. By identifying context-specific issuance timing effects, flexible machine learning approaches, such as causal forests, allow for more precise risk assessments and better-informed decision-making across financial markets. These insights highlight the need for financial researchers and policymakers to consider flexible estimation techniques when evaluating financial market timing strategies.

III. Conclusion

In this paper, we argue that market timing effects may be heterogeneous, emphasising the advantages of flexible methods for estimating these varying effects over traditional approaches that assume uniformity. To demonstrate this, we apply the causal forest—a causal machine learning method designed for estimating heterogeneous treatment effects—and use the diverse catastrophe bond market as a case study in market timing. Our analysis examines how issuance timing heterogenously affects primary catastrophe bond spreads, identifies the key factors influencing these variations, and assesses the role of overall market conditions in shaping the extent of these differences.

Our analysis reveals that the impact of issuance timing on primary catastrophe bond spreads is far from uniform. Early-year issuances are generally associated with lower spreads at issuance, but the extent of this effect varies greatly depending on market conditions in the reinsurance and ILS markets and specific characteristics of the catastrophe bonds such as issuance size. We find that, in soft market conditions, early issuances tend to benefit from a more pronounced spread discount, particularly for catastrophe bonds linked to peak peril-region combinations. However, this timing advantage diminishes in hard market conditions, where the opposite effect can sometimes occur, especially for niche perils.

While our study focuses on the catastrophe bond market, its implications extend to broader financial markets where timing strategies exhibit variable impacts. By employing the causal forest

approach, we demonstrate the importance of accounting for heterogeneity in timing effects—an insight that can benefit corporate bond issuances, structured finance, and alternative asset pricing models. Future research could build on this work by exploring additional dimensions of timing strategies, including the role of behavioural finance in issuance decisions or the impact of climate risk perceptions on catastrophe bond pricing.

Our results provide valuable insights for both issuers and investors, emphasising the importance of timing strategies that align with market liquidity conditions, investor sentiment, and capital flows. As financial markets become increasingly complex, machine learning-based approaches offer a promising alternative to traditional econometric methods for understanding dynamic, heterogeneous financial behaviours.

References

Athey, S., Tibshirani, J., and Wager, S. (2019). Generalized random forests. *The Annals of Statistics*, 47(2):1148 – 1178.

Athey, S. and Wager, S. (2019). Estimating treatment effects with causal forests: An application. *Observational Studies*, 5(2):37–51.

Ball, E., Chiu, H. H., and Smith, R. (2011). Can VCs time the market? An analysis of exit choice for venture-backed firms. *The Review of Financial Studies*, 24(9):3105–3138.

Barrieu, P., Braun, A., and Makariou, D. (2025). *Catastrophe Bonds*, pages 169–195. Springer Nature Switzerland, Cham.

Bénard, C. and Josse, J. (2023). Variable importance for causal forests: breaking down the heterogeneity of treatment effects. Available at: `https://arxiv.org/abs/2308.03369`. Accessed: 2024-11-10.

Boileau, P., Qi, N. T., van der Laan, M. J., Dudoit, S., and Leng, N. (2023). A flexible approach for predictive biomarker discovery. *Biostatistics*, 24(4):1085–1105.

Bollen, N. P. and Busse, J. A. (2001). On the timing ability of mutual fund managers. *The Journal of Finance*, 56(3):1075–1094.

Bolton, P., Chen, H., and Wang, N. (2013). Market timing, investment, and risk management. *Journal of Financial Economics*, 109(1):40–62.

Braun, A. (2016). Pricing in the primary market for cat bonds: new empirical evidence. *Journal of Risk and Insurance*, 83(4):811–847.

Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.

Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1986). *Classification and Regression Trees*. Wadsworth Brooks/Cole, Monterey, CA.

Busse, J. A., Ding, J., Jiang, L., and Tang, Y. (2023). Artificial market timing in mutual funds. *Journal of Financial and Quantitative Analysis*, 58(8):3450–3481.

Cao, S. S., Wang, Z., and Yeung, P. E. (2022). Skin in the game: Operating growth, firm performance, and future stock returns. *Journal of Financial and Quantitative Analysis*, 57(7):2559–2590.

Chen, Y. and Liang, B. (2007). Do market timing hedge funds time the market? *Journal of Financial and Quantitative Analysis*, 42(4):827–856.

Crump, R. K., Hotz, V. J., Imbens, G. W., and Mitnik, O. A. (2008). Nonparametric tests for treatment effect heterogeneity. *The Review of Economics and Statistics*, 90(3):389–405.

Davis, J. M. and Heller, S. B. (2017). Using causal forests to predict treatment heterogeneity: An application to summer jobs. *American Economic Review*, 107(5):546–550.

Dixon, P. N. and Kelley, E. K. (2022). Business cycle variation in short selling strategies: Picking during expansions and timing during recessions. *Journal of Financial and Quantitative Analysis*, 57(8):3018–3047.

Evans, K. P., Leung, W. S., Li, J., and Mazouz, K. (2024). ETF ownership and seasoned equity offerings. *Journal of Financial and Quantitative Analysis*, 59(4):1821–1848.

Ferson, W. E. and Schadt, R. W. (1996). Measuring fund strategy and performance in changing economic conditions. *The Journal of Finance*, 51(2):425–461.

Fung, W. and Hsieh, D. A. (1997). Empirical characteristics of dynamic trading strategies: The case of hedge funds. *The Review of Financial Studies*, 10(2):275–302.

Fung, W. and Hsieh, D. A. (2001). The risk in hedge fund strategies: Theory and evidence from trend followers. *The Review of Financial Studies*, 14(2):313–341.

Gao, P. and Lou, D. (2013). Cross-market timing in security issuance. AFA 2012 Chicago Meetings Paper. Available at SSRN: `https://ssrn.com/abstract=1787187`. Accessed: 2025-01-15.

Glasserman, P., Li, F., and Mamaysky, H. (2023). Time variation in the news–returns relationship. *Journal of Financial and Quantitative Analysis*, page 1–37.

Goetzmann, W. N., Ingersoll, J., and Ivković, Z. (2000). Monthly measurement of daily timers. *Journal of Financial and Quantitative Analysis*, 35(3):257–290.

Gulen, H., Jens, C., and Page, B. (2024). Balancing external vs. internal validity: An application of causal forest in finance. Available at SSRN: https://ssrn.com/abstract=3583685. Accessed: 2025-02-10.

Hill, J. L. (2011). Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240.

Hines, O., Diaz-Ordaz, K., and Vansteelandt, S. (2022). Variable importance measures for heterogeneous causal effects. Available at: https://arxiv.org/abs/2204.06030. Accesed: 2025-01-03.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945–960.

Jiang, G. J., Yao, T., and Yu, T. (2007). Do mutual funds time the market? Evidence from portfolio holdings. *Journal of Financial Economics*, 86(3):724–758.

Lee, M.-j. (2009). Non-parametric tests for distributional treatment effect for randomly censored responses. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 71(1):243–264.

Lewis, C. M. and Tan, Y. (2016). Debt-equity choices, R&D investment and market timing. *Journal of Financial Economics*, 119(3):599–610.

Lloyd's (2024). LCM Forecast Returns 2024: Guidance and Instructions. Available at:

https://assets.lloyds.com/media/a589a726\protect\penalty\z@-6776-4367\protect\penalty\z@-bc3adb9a6eb9/2.LCM-Forecast-Returns-2024.pdf. Accessed: 2025-02-06.

Makariou, D., Barrieu, P., and Chen, Y. (2021). A random forest based approach for predicting spreads in the primary catastrophe bond market. *Insurance: Mathematics and Economics*, 101:140–162.

Swiss Re (2024). Insurance-linked securities market insights. Available at: https://www.swissre.com/dam/jcr:bb189e59-a15f-49df-a250-07b2c6b2d9bd/2024-02-sr-ILS-market-insights-feb-2024.pdf. Accessed: 2025-02-06.

Willke, R. J., Zheng, Z., Subedi, P., Althin, R., and Mullins, C. D. (2012). From concepts, theory, and evidence of heterogeneity of treatment effects to methodological approaches: A primer. *BMC Medical Research Methodology*, 12:1–12.