CFR working paper NO. 26-06

The Green Bond Premium: Evidence from a Multiverse Analysis

T. Bauckloh ● P. Kirsch

centre for financial research
cologne

The green bond premium: evidence from a multiverse analysis

April 2026

Abstract

We study the green bond premium, defined as the yield differential between green and matched conventional bonds in the secondary market. Existing estimates vary widely, raising questions about their robustness. We address this by estimating the premium across more than 500,000 empirical designs spanning common sample and methodological choices. In this multiverse setting, the average premium is -2.59 basis points. It varies systematically with sample composition, with more negative values for municipal bonds, and becomes more negative during periods of heightened climate attention. Finally, we investigate which choices drive variation in premium estimates. We find that it is driven primarily by issuer type and matching choices, while other choices, such as liquidity adjustment, contribute little to overall variation.

JEL Classification Codes: C52, G11, G12, Q54

1 Introduction

The green bond premium, defined as the yield differential between green bonds and matched conventional bonds in the secondary market, has received considerable attention in finance research. Reported estimates range from approximately -60 to +8 basis points (see, e.g., Karpf and Mandel 2018, Nanayakkara and Colombage 2019, Zerbib 2019). These estimates vary with, among other things, sample construction, matching criteria, and liquidity adjustments. As a result, it remains unclear whether the green bond premium reflects a robust pricing effect or primarily differences in empirical design.

In this paper, we address this question by systematically estimating the green bond premium across more than 500,000 distinct empirical designs. These designs span the main sample construction and methodological choices used in the literature and allow us to map the full distribution of plausible green bond premium estimates. This large-scale multiverse design provides a unified and specification-robust assessment of the size of the green bond premium, its cross-sectional drivers, and variation over time. Additionally, it enables us to identify which design choices materially affect its magnitude and which do not.

Our analysis covers green bonds issued by sovereign, municipal, supranational, and corporate issuers across multiple currencies and markets. We vary matching thresholds, matching methodologies, yield definitions, liquidity adjustments, and aggregation procedures. Across the full specification space, the mean green bond premium equals -2.59 basis points. Beyond this average effect, the magnitude of the premium varies systematically across bond characteristics. Municipal and USD-denominated bonds exhibit more negative premia, whereas corporate issuers and EUR-denominated bonds display smaller and more stable premia.

The green bond premium also varies over time. We therefore examine its dynamics in relation to environmental attention and broader market conditions. Theory predicts lower expected returns for sustainable assets either because investors derive non-pecuniary utility from holding them or because green assets hedge climate-related risks (Pástor et al. 2021, Pástor et al. 2022). While both mechanisms can rationalize lower returns in equity markets, the risk-based channel is likely limited in the context of matched green and conventional bonds. Consistent with a demand-based interpretation, we find that the green bond premium co-moves systematically with measures of climate-related attention and financial market uncertainty. Specifically, the premium becomes more negative during periods of heightened climate concern, as measured by the Media Climate Change Concerns index (Ardia et al. 2023) and a Twitter-based climate attention index (Arteaga Garavito et al. 2025). These findings indicate a systematic link between shifts in environmental attention and green bond pricing.

While the average level and time variation of the green bond premium reveal economically meaningful patterns, the large number of plausible empirical design choices raises an additional question: which choices drive the dispersion in estimated premia (which we refer to as "design uncertainty")? We address this question by systematically quantifying how alternative specifications affect the estimated premium. Our results show that a small subset of design choices accounts for most cross-specification dispersion. In particular, issuer type and the choice of matching ratio generate the largest variation in estimates, with mean absolute differences of up to 6.1 basis points, which exceeds the unconditional mean green bond premium of -2.59 basis points in absolute terms. Consistent with this pattern, the matching ratio alone explains 23.0% of the total variation in estimates in a Shapley decomposition, followed by issuance-date thresholds (20.3%) and currency (13.4%). By contrast, green bond certification standards, yield definition, and liquidity adjustments contribute little to overall variation, each accounting for less than 1% of the explained variance.

We contribute to two strands of the literature. First, we contribute to the literature on green bond pricing in the secondary market by providing a comprehensive and robust assessment of both the level and the dynamics of the green bond premium. Prior studies document lower yields for green bonds relative to comparable conventional bonds (e.g., Nanayakkara and Colombage 2019; Zerbib 2019; Partridge and Medda 2020; Immel et al. 2021; Dorfleitner et al. 2022), while others find no statistically significant premium or even higher yields (e.g., Karpf and Mandel 2018; Hachenberg and Schiereck 2018; Bachelet et al. 2019). We reconcile this dispersion by systematically evaluating more than 500,000 empirical designs derived from the existing literature. Across this specification space, the distribution is centered on a small but consistently negative premium, with a mean of -2.59 basis points and a median of -1.66 basis points.

Beyond its average level, we document time variation in the green bond premium. Prior work shows that sustainability-related demand varies with climate shocks and attention. For example, Pástor et al. (2022) show that a green equity premium increases when climate concerns intensify, Dragotto et al. (2025) document temporary increases in a green corporate bond premium following major climate policy announcements and climate shocks, and Merli et al. (2025) find that heightened climate attention is associated with larger green premia in European markets. Building on this literature, we demonstrate that the association between climate-related attention and the green bond premium holds across a wide range of samples and empirical designs. By exploiting variation across sample and methodological choices, we show that the documented time-series relationship is not driven by particular samples or methods.

Second, we contribute to the literature on non-standard errors, a term introduced by

Menkveld et al. (2024) to describe variation in empirical findings that arises from researchers' design choices. While prior work documents heterogeneity in green bond premium estimates, little is known about how much of this variation is attributable to sample construction and methodological choices. Adapting the multiverse framework of Walter et al. (2024) to the context of matched bond yield differentials, we systematically vary a large set of empirically justified design choices and characterize the full distribution of green bond premium estimates. Our results quantify which choices materially affect estimated premia and which are largely innocuous. More broadly, our findings relate to the replication and robustness literature in empirical asset pricing, which documents that published results can be sensitive to design choices. In this spirit, we contribute to the growing literature on the so-called replication crisis (e.g., Chen and Zimmermann 2022, Hasler 2023, Hou et al. 2020, and Jensen et al. 2023) by demonstrating that the wide range of green bond premium estimates reported in prior studies can be reproduced within a unified and transparent empirical framework.

The remainder of the paper is organized as follows. Section 2 describes the data, the matched bond-pair framework, and the full set of sample construction and methodological decision forks underlying the multiverse design. Section 3 presents the distribution of green bond premium estimates across specifications and examines both cross-sectional and time-series drivers of the premium. Section 4 analyzes design uncertainty by quantifying the sensitivity of estimates to individual design choices, assessing their relative importance, and comparing distributions under fixed sample and methodological settings. Section 5 concludes.

2 Data and Methodology

2.1 Data

We obtain the universe of green bonds from the LSEG Workspace Advanced Search bond database. For each green bond, we identify all other bonds issued by the same issuer as potential conventional matches. We also retrieve bond-level fundamental data from LSEG Workspace. The dataset includes, for each bond, the issuer ticker and issuer type, the maturity and issuance dates, the issuance amount in both the original currency and USD equivalents, the principal and coupon currency, the coupon type and coupon rate, the bond's collateral and seniority features, its structure, the asset status, and a standardized long-term credit rating.[1] In addition, the dataset contains several bond sustainability characteristics,

---

[1] We use Moody's and Fitch long-term credit ratings and standardize them by mapping all values to the Moody's long-term rating scale. When a Moody's rating is available, we use it directly. When it is missing, we substitute the corresponding Moody's-equivalent value derived from the Fitch rating.

including indicators for alignment with the Climate Bonds Initiative (CBI) and with the ICMA Green Bond Principles (GBP). We supplement the static bond-level fundamentals with time-series data on bid and ask prices and yields obtained from LSEG Workspace. The mid-yield is computed as the simple average of the bid and ask yields. The data are collected from each bond's issuance date until maturity or, at the latest, the end of 2024. We provide definitions for variables in Table A1 in the Appendix.

We restrict the sample to plain-vanilla fixed-coupon bonds and exclude bonds reported to be in default. Following Dick-Nielsen et al. (2012), we remove bonds with an initial time to maturity exceeding 30 years and retain observations only up to one month prior to maturity. We further exclude bonds for which principal and coupon currencies differ, as such mismatches may generate unreliable yield estimates. To ensure internal consistency of quoted prices, we require bid prices to be weakly below ask prices and bid yields to be weakly above ask yields. Observations violating these conditions are removed. Negative bid or ask prices are set to missing. To eliminate remaining outliers, we remove bonds with more than three daily yield observations outside the range of -2% to 40%. Any remaining yields outside this range are set to missing.

After applying these filtering and cleaning steps, we arrive at the final bond sample that underlies all subsequent analyses. The sample consists of all green bonds and the associated universe of potential conventional counterparts. All green bond premium calculations are based on this sample, with bonds and time-series observations selected according to the respective empirical specification. Figure 1 illustrates the evolution in the number of green bonds and issuers in our sample, highlighting a sharp expansion of the market in recent years. The first green bond in our sample was issued on 12 November 2008 (ISIN: XS0398811959), marking the beginning of the period over which green bonds can potentially enter our analysis. In total, the sample comprises 4,889 different green bonds issued by 1,360 distinct issuers. Among these, 3,668 bonds are classified as green according to the ICMA Green Bond Principles, while 3,682 satisfy the CBI alignment criterion. The two classifications overlap substantially but are not identical, reflecting differences in certification criteria and coverage across issuers and time.

Table 1 provides an overview of issuer types and currencies in the green bond sample. The majority of green bonds are issued by corporate entities, followed by municipal and agency issuers, while sovereign and supranational bonds represent the smallest shares of the sample. With respect to currency denomination, U.S. dollars and euros account for the largest fractions of green bond issuance, followed by Chinese yuan, Japanese yen, and South Korean won. Together, these five currencies constitute the dominant currency segments in the sample and reflect broad global coverage across both advanced and emerging green bond

Figure 1: Number of Green Bonds and Issuers Over Time

This figure displays the monthly number of outstanding green bonds and distinct issuers for which time-series data are available. The sample covers the period from the initial issuances in 2008 until end of 2024.

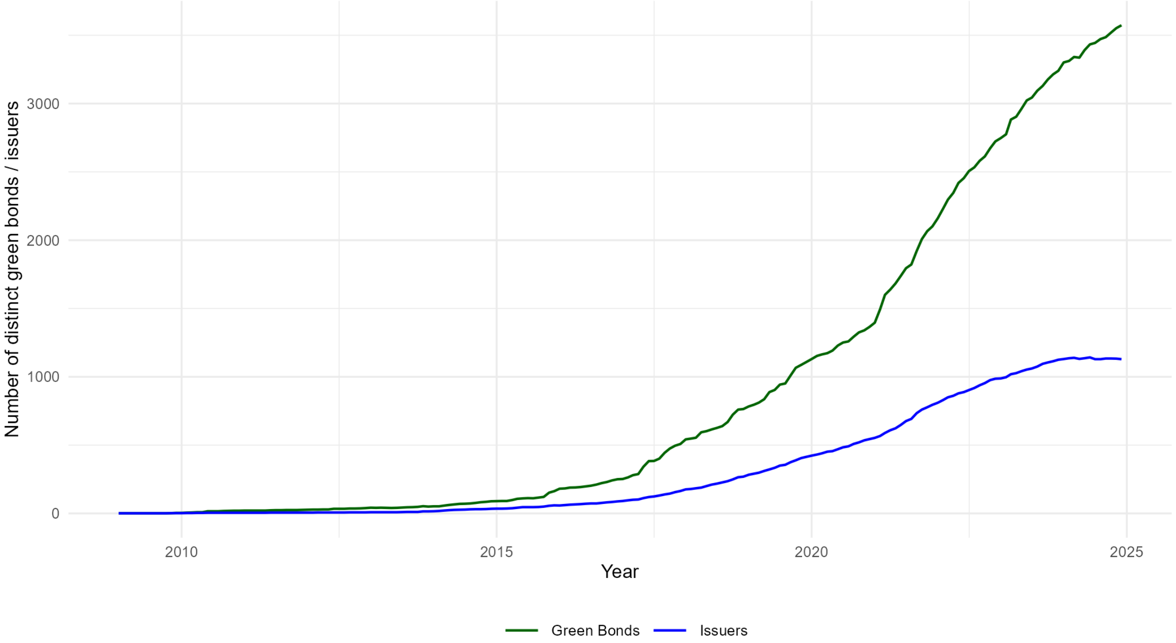

Table 1: Descriptive Statistics - Discrete Variables

This table reports descriptive statistics for discrete variables of the green bond sample. **Panel A** reports statistics by issuer type and **Panel B** by currency. Percentages are calculated relative to the total sample.

	Green Bonds		**Issuers**	
	Number	%	Number	%
Panel A: Issuer Types				
Corporate	3,167	66.0	1,169	84.0
Municipals	700	14.6	97	7.0
Agency	543	11.3	82	5.9
Supranational	330	6.9	16	1.2
Sovereign	61	1.3	27	1.9
Panel B: Top Currencies				
US Dollar	1,119	23.3	283	17.5
Euro	1,061	22.1	300	18.5
Chinese Yuan	791	16.5	363	22.4
Japanese Yen	507	10.6	191	11.8
South Korean Won	341	7.1	93	5.7

markets.

Table 2 reports summary statistics for the continuous variables in the sample, covering both green bond fundamentals and daily secondary-market characteristics. Issuance amounts vary substantially across green bonds, with a mean of USD 116 million and an interquartile range from USD 33 million to USD 500 million. Coupon rates are moderate, with a median of 2.6%. Yields are also concentrated in a moderate range, with a median between 2.8% and 2.9% and a mean of 2.9%. Bid-ask spreads are generally narrow, indicating comparatively tight trading conditions in secondary markets.

Table 2: Descriptive Statistics - Continuous Variables

This table reports descriptive statistics for continuous variables of the green bond sample. **Panel A** shows the distribution of issuance amount (in USD millions) and coupon rate (in %), while **Panel B** shows the distribution of yield variables (in %), prices (in local currency) and the relative bid–ask spread (in %).

	Mean	Min	10%	25%	Median	75%	90%	Max
Panel A: Bond Fundamentals								
Issuance Amount	428.98	0	7.15	32.84	115.68	500	867.60	53,038.34
Coupon Rate	2.66	0	0.34	1.0	2.65	3.90	5.0	11.84
Panel B: Bond Time-Series Variables								
Bid Yield	2.91	-1.97	0.36	1.25	2.86	4.04	5.22	39.91
Ask Yield	2.83	-1.96	0.31	1.20	2.79	3.96	5.12	39.81
Mid Yield	2.87	-0.99	0.34	1.24	2.83	4.00	5.17	39.86
Bid–Ask Spread	0.34	0	0.06	0.13	0.26	0.45	0.62	85.71

2.2 Methodology and Decision Forks

The literature employs several approaches to estimate the green bond premium, including matched bond pairs, regression-based designs, and yield-curve comparisons. Each approach entails a sequence of sample construction and methodological decisions. We focus on the matched bond pair framework because it is arguably the most widely used design and offers a transparent structure for organizing methodological variation.

At each decision fork, researchers choose among alternative choices. Each combination of choices defines a distinct green bond premium estimation path, and the set of all feasible combinations spans the universe of green bond premium estimates. In the following, we describe each fork and its associated choices in detail. We distinguish between two classes of forks. Sample forks determine the composition of the underlying green bond universe. Methodological forks arise from the implementation of the matched-pair estimation procedure. The classification of forks and their respective choices is derived from the leading bond green bond premium literature. Table 3 summarizes all forks and their associated choices.

We consider four sample decision forks. First, green bonds can be identified in different ways: the database label (e.g., Nanayakkara and Colombage 2019, Dorfleitner et al. 2022), compliance with the ICMA green bond principles (e.g., Zerbib 2019), or alignment with the standards of the CBI (e.g., Bachelet et al. 2019). Second, we vary currency restrictions by allowing either the full multi-currency universe (e.g., Nanayakkara and Colombage 2019, Zerbib 2019) or subsamples limited to EUR- or USD-denominated issuances (e.g., Karpf and Mandel 2018, Partridge and Medda 2020). Third, we partition the issuer universe and estimate the green bond premium separately for corporate (e.g., Nanayakkara and Colombage 2019, Dragotto et al. 2025), municipal (e.g., Karpf and Mandel 2018, Partridge and Medda 2020), and supranational, sovereign, and agency (SSA) issuers (e.g., Pástor et al. 2022), or pool all issuer types (e.g., Bachelet et al. 2019, Zerbib 2019). Finally, we vary the time horizon by either employing the full sample period (e.g., Dorfleitner et al. 2022, Dragotto et al. 2025) or focusing on bonds issued until the end of 2017 (e.g., Hachenberg and Schiereck 2018, Zerbib 2019) versus after 2017 (e.g., Merli et al. 2025).

Methodological decision forks govern how the green bond premium is computed once the underlying sample has been defined. We require all potential conventional matches to share the same issuer, currency, bond structure, coupon type, and seniority or collateral status as the corresponding green bond (e.g., Zerbib 2019, Dorfleitner et al. 2022). In addition to these fixed criteria, we treat credit rating as the first methodological decision fork. If exact matching is selected (e.g., Zerbib 2019, Dorfleitner et al. 2022), the green bond and its conventional counterparts must have identical alphanumeric ratings (e.g., AAA matched only with AAA). If not selected (e.g., Hachenberg and Schiereck 2018), credit rating is excluded from the matching conditions.

Moreover, we introduce threshold forks for bond characteristics for which exact matching is not feasible. Thresholds define the maximum allowed deviation between a green bond and a potential conventional match. The first threshold concerns the issuance amount. A conventional bond is eligible only if its issuance amount is at least one-half (or one-quarter) and at most twice (or four times) the issuance amount of the green bond (e.g., Zerbib 2019, Kapraun et al. 2021, Dorfleitner et al. 2022). The second threshold applies to the maturity date. The absolute difference in maturity dates must not exceed one year (e.g., Zerbib 2019, Merli et al. 2025), two years (e.g., Bachelet et al. 2019, Zerbib 2019), or may remain unrestricted (e.g., Hachenberg and Schiereck 2018, Dragotto et al. 2025). The third threshold restricts issuance dates. The issuance dates of the green bond and its conventional counterpart must lie within either two (e.g., Zerbib 2019, Kapraun et al. 2021)

Table 3: Possible Forks and Choices in the Green Bond Premium Calculation

This table provides an overview of the forks and choices in the green bond premium calculations. N gives the number of possible choices per fork, and the possible fork-paths result from the combination of all possible choices. **Panel A** reports all sample forks, i.e., forks that change the underlying bond sample. **Panel B** reports all methodology forks, i.e., forks that affect the matching procedure and the calculation of the green bond premium.

Fork	**Classification**	**Possible Choices**	**N**
		Panel A: Sample	
Green Bond definition	-	Green Bond Principles (ICMA) CBI aligned By database	3
Currency	-	Only EUR Only USD All currencies	3
Issuer Types	-	Corporates Municipals Sovereign, supranational, and agency (SSA) All issuers	4
Time Horizon	-	before 2018 after 2017 All years	3
		Panel B: Methodology	
Credit Rating	Exact Matching	Yes No	2
Issuance Amount	Threshold	Log(2) Log(4)	2
Maturity Date	Threshold	1 year 2 years No threshold	3
Issuance Date	Threshold	2 years 6 years No threshold	3
Coupon Rate	Threshold	0.25 percentage points No threshold	2
Method	Matching	Propensity Score Matching (PSM) Closest maturity	2
Ratio	Matching	1:1 matching 1:2 with forced interpolation 1:2 allowing for extrapolation	3
Yield	Estimation	Ask yield Bid yield Mid yield	3
Liquidity Adjustment	Estimation	Yes No	2
Aggregation	Estimation	Time-specific Bond-specific	2
Possible fork-paths			**559,872**

or six years (e.g, Zerbib 2019, Dorfleitner et al. 2022) of each other, or may be unrestricted (e.g., Hachenberg and Schiereck 2018, Bachelet et al. 2019). The fourth threshold pertains to the coupon rate. A conventional bond qualifies as a match only if the absolute coupon difference does not exceed 25 basis points (e.g., Bachelet et al. 2019), unless no coupon restriction is imposed (e.g., Nanayakkara and Colombage 2019, Zerbib 2019). All threshold conditions apply jointly: a conventional bond is considered a valid match only if it satisfies every active threshold in addition to the always-imposed exact-match requirements. If no conventional bond meets all criteria, the respective green bond is excluded under the given fork-path.

After identifying the universe of eligible conventional bonds that satisfies the exact-match and threshold criteria, we consider two alternative matching procedures to determine the final matched counterparts. The first procedure is nearest-neighbour matching based on the threshold characteristics issuance amount, maturity date, and issuance date (e.g., Bachelet et al. 2019, Dragotto et al. 2025). Specifically, for each green bond, we select the conventional bond with the closest propensity score among all eligible candidates. The second procedure is closest-maturity matching (e.g., Hachenberg and Schiereck 2018, Zerbib 2019). Eligible conventional bonds are ranked by the absolute difference between their maturity date and that of the green bond. In the event of ties, we first select the bond with the smaller issuance amount deviation and, if necessary, the one with the smaller issuance date deviation. We further vary the matching ratio. Under 1:1 matching, each green bond is paired with a single conventional bond (e.g., Bachelet et al. 2019, Dragotto et al. 2025). Under 1:2 matching with forced interpolation, each green bond is matched to one shorter- and one longer-maturity conventional bond, imposing a strict interpolation structure around the green bond's maturity (e.g., Hachenberg and Schiereck 2018, Dorfleitner et al. 2022). Under 1:2 matching with extrapolation, each green bond is matched to its two closest conventional bonds, which may both lie on the same side of the green bond in terms of maturity (e.g., Zerbib 2019). This specification permits yield extrapolation beyond the green bond's maturity.

Whenever we employ 1:2 matching, we construct a synthetic conventional yield at the exact maturity of the green bond as done by Zerbib (2019). Specifically, we fit a linear function through the yields and maturities of the two matched conventional bonds, obtaining slope $a_{i,t}$ and intercept $b_{i,t}$. The synthetic conventional yield evaluated at the green bond's maturity is then given by

$$\text{Yield}_{i,t}^{\text{syn}} = a_{i,t} \cdot \text{Maturity}_i + b_{i,t} \tag{1}$$

where $\text{Yield}_{i,t}^{\text{syn}}$ denotes the synthetic conventional yield for green bond i on day t, and Maturity_i refers to the maturity date of the respective green bond.

We now define the yield differential $\Delta\text{Yield}_{i,t}$ as the green bond's yield minus the yield of its matched conventional counterpart, i.e.,

$$\Delta\text{Yield}_{i,t} = \text{Yield}^{\text{green}}_{i,t} - \text{Yield}^{\text{conv}}_{i,t} \tag{2}$$

where $\text{Yield}^{\text{conv}}_{i,t}$ denotes the yield of the matched conventional bond in the 1:1 case or the synthetic conventional yield in the 1:2 case.

Lastly, we introduce three decision forks that govern the estimation of the green bond premium itself. The first fork concerns the yield input. We consider three alternatives: ask yield (e.g., Bachelet et al. 2019, Zerbib 2019), bid yield (e.g., Dorfleitner et al. 2022), and mid yield (e.g., Kapraun et al. 2021, Pástor et al. 2022). The second fork determines whether the green bond premium is adjusted for liquidity (e.g., Zerbib 2019, Merli et al. 2025). We measure liquidity using the relative bid-ask spread, defined as the difference between ask and bid prices scaled by the average of the bid and ask price as in Zerbib (2019). Under 1:1 matching, the liquidity differential is simply the difference between the bid-ask spread of the green bond and that of its matched conventional counterpart. Under 1:2 matching, however, a synthetic bid-ask spread must first be constructed for the conventional side to ensure consistency with the synthetic yield used in the yield differential. Specifically, we compute the synthetic bid-ask spread $\text{BA}^{\text{syn}}_{i,t}$ as a maturity-distance-weighted average of the bid-ask spreads of the two matched conventional bonds (Zerbib 2019):

$$\text{BA}^{\text{syn}}_{i,t} = \frac{d_{2,i}}{d_{1,i} + d_{2,i}} \text{BA}_{\text{CB1},i,t} + \frac{d_{1,i}}{d_{1,i} + d_{2,i}} \text{BA}_{\text{CB2},i,t} \tag{3}$$

where $\text{BA}_{\text{CB1},i,t}$ and $\text{BA}_{\text{CB2},i,t}$ denote the bid-ask spreads of the two matched conventional bonds. $d_{1,i}$ is the absolute difference between the maturity date of the green bond and that of the first conventional match; $d_{2,i}$ is defined analogously for the second match. The liquidity difference is then defined as the difference between the green bond's bid-ask spread and the synthetic bid-ask spread.

When liquidity adjustment is applied, we estimate the following regression and save the estimated parameters:

$$\Delta\text{Yield}_{i,t} = \beta\, \Delta\text{Liquidity}_{i,t} + \gamma_i + \varepsilon_{i,t} \tag{4}$$

where $\Delta\text{Yield}_{i,t}$ is the yield difference from Equation (2), $\Delta\text{Liquidity}_{i,t}$ is the liquidity difference of the matched pair, γ_i denotes bond fixed effects, and $\varepsilon_{i,t}$ is the error term.

Finally, aggregation determines whether the green bond premium is computed at the bond level (cross-sectional measure) or at the time level (time-series measure). A bond-specific green bond premium summarizes the yield difference for each green bond across time, while a time-specific green bond premium summarizes the cross-sectional yield difference across all green bonds observed on a given day. Table 4 summarizes this final step in the calculation of the green bond premium.

Table 4: Calculation of the Green Bond Premium

This table summarizes how the green bond premium is computed depending on (i) whether yields are adjusted for liquidity and (ii) whether aggregation is performed at the bond level (averaging per bond over time) or at the day level (averaging across bonds observed on a given day). i denotes a green bond and t denotes the day of observation. Without liquidity adjustment, the yield difference $\Delta\text{Yield}_{i,t}$ from Equation (2) is used. With liquidity adjustment, Equation (4) is estimated to obtain bond fixed effects γ_i and residuals $\varepsilon_{i,t}$.

	Liquidity Adjustment	
Aggregation Level	No	Yes
Bond-level: green bond premium$_i$	$\frac{1}{T_i}\sum_{t=1}^{T_i}\Delta\text{Yield}_{i,t}$	γ_i
Day-level: green bond premium$_t$	$\frac{1}{N_t}\sum_{i=1}^{N_t}\Delta\text{Yield}_{i,t}$	$\frac{1}{N_t}\sum_{i=1}^{N_t}(\gamma_i+\varepsilon_{i,t})$

Without liquidity adjustment, the green bond premium corresponds directly to the yield difference $\Delta\text{Yield}_{i,t}$ (e.g., Hachenberg and Schiereck 2018, Partridge and Medda 2020). For a bond-specific green bond premium, we average $\Delta\text{Yield}_{i,t}$ over all days t for each green bond i. For a time-specific green bond premium, we average $\Delta\text{Yield}_{i,t}$ across all green bonds observed on day t.

With liquidity adjustment, Equation (4) decomposes the yield difference into a bond fixed effect γ_i and a residual component $\varepsilon_{i,t}$. We use these components to construct liquidity-adjusted green bond premiums.

For a bond-specific green bond premium, the fixed effect γ_i serves as the liquidity-adjusted green bond premium for bond i (e.g., Zerbib 2019). For a time-specific green bond premium (e.g., Dorfleitner et al. 2022), the sum of the fixed effect γ_i and the residuals $\varepsilon_{i,t}$ serves as the time-varying liquidity-adjusted green bond premium for bond i on day t. We then average this measure across all green bonds observed on day t to compute a time-specific liquidity-adjusted green bond premium.

Taken together, these forks generate 559,872 distinct fork-paths for estimating the green bond premium. 41,610 (7.43%) of the fork-paths are infeasible. Infeasibility arises primarily in the pre-2018 subsample due to the small number of outstanding green bonds, which limits matching under restrictive specifications. Further infeasibility results from requirements for exact credit-rating matches, as rating information is frequently missing. For the time-specific green bond premium, feasible fork-paths contain on average 1,646.28 daily observations, based on an average of 100.77 green bonds contributing each day. For the bond-specific green bond premium, feasible fork-paths include on average 260.42 matched bond-pairs, with each pair providing an average of 665.02 daily observations. For each successfully estimated fork-path, we store a single summary statistic: the average green bond premium implied by that specification. When the green bond premium is modeled as time-varying (bond-day level), this summary measure corresponds to the time-series mean of the daily green bond premium estimates. When the green bond premium is specified as bond-level and time-invariant, the summary measure corresponds to the cross-sectional mean across bonds. In addition, for each fork-path we record the associated t-statistic and the Wilcoxon signed-rank statistic, both testing the null hypothesis that the mean green bond premium implied by the respective specification equals zero.

3 Results

3.1 The Multiverse Green Bond Premium

Table 5 reports summary statistics for the 518,262 green bond premium estimates, while Figure 2 visualizes their distribution together with the distribution of the corresponding t-statistics. The mean green bond premium across specification paths equals -2.59 bps, with a median of -1.66 bps. The 25th and 75th percentiles are -4.27 bps and 0.12 bps, respectively, implying an interquartile range of 4.39 bps. Notably, the interquartile range exceeds the absolute value of the mean premium, indicating that plausible specification choices can materially alter the estimated magnitude and even the sign of the premium. Based on a t-test of the null hypothesis that a green bond premium equals zero, 61.24% of all paths yield a statistically significant green bond premium estimate at the 5% level, of which 78.87% are negative. The Wilcoxon signed-rank test produces a similar pattern: 73.21% of green bond premium estimates are significant at the 5% level, and 79.12% of those are negative.

Table 5: Summary Statistics of Green Bond Premium Estimates

This table reports the distribution of the green bond premium estimates across all feasible fork-paths in %. IQR represents the interquartile range.

Mean	Min	1%	25%	Median	75%	99%	Max	IQR
-0.0259	-15.3823	-0.2867	-0.0427	-0.0166	0.0012	0.2109	2.7028	0.0439

Figure 2: Distribution of the Green Bond Premium and Corresponding t-Statistics

This figure shows the distribution of the green bond premium across all feasible fork-paths and the corresponding distribution of the associated t-statistics.

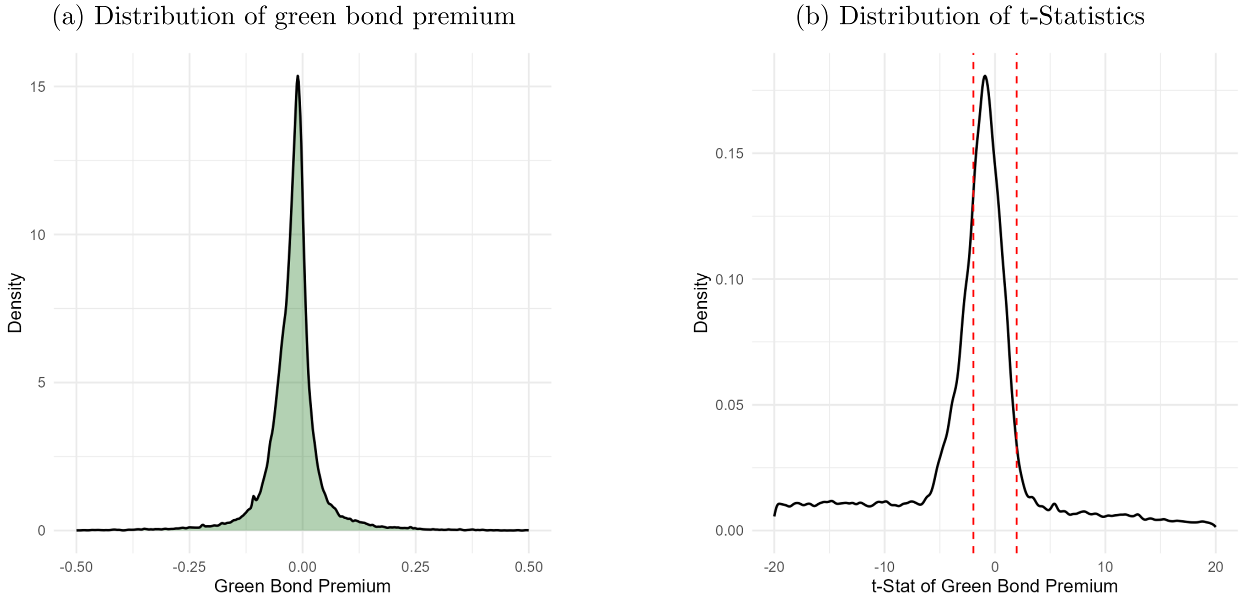
(a) Distribution of green bond premium

(b) Distribution of t-Statistics

3.2 Cross-sectional Drivers

Having documented the overall distribution of the green bond premium, we next examine its magnitude and sign in the cross-section. For each cross-sectional sample decision fork, we fix one choice and visualize the resulting green bond premium distribution across all remaining specification combinations.

Table 6 reports summary statistics conditional on cross-sectional sample forks. The choice of green bond definition, whether based on the database classification, the ICMA GBP, or the CBI standard, has only a limited effect on the mean green bond premium. In contrast, currency exerts a more pronounced influence. Restricting the sample to EUR-denominated bonds yields a less negative mean green bond premium (-1.64 bps) than USD-denominated bonds (-3.61 bps). This pattern also holds for the median premium, although the difference is somewhat smaller, because USD-denominated green bonds exhibit more extreme premia than their EUR-denominated counterparts. Across issuer types, municipal bonds generate the most negative mean green bond premium (-3.78 bps), whereas corporate and SSA bonds exhibit more moderate values.

We supplement Table 6 by providing a regression-based assessment of how cross-sectional sample forks influence the magnitude and sign of green bond premium estimates. We estimate the following regression model:

Table 6: Distribution of the Green Bond Premium Estimates - Cross Section

This table reports the distribution of the green bond premium estimates (in %) when fixing a specific cross-sectional sample fork. The remaining forks each vary in their choices as defined in Table 3. IQR represents the interquartile range. **Panel A** reports results by green bond definition, **Panel B** by currency, and **Panel C** by issuer type.

	Mean	Min	1%	25%	Median	75%	99%	Max	IQR
Panel A: Green Bond Definition									
CBI	-0.0283	-15.3823	-0.3070	-0.0435	-0.0178	-0.0007	0.2301	1.7163	0.0428
ICMA	-0.0234	-3.1809	-0.2654	-0.0413	-0.0153	0.0025	0.1759	1.7163	0.0438
By database	-0.0259	-15.0113	-0.2932	-0.0433	-0.0167	0.0018	0.2264	2.7028	0.0451
Panel B: Currency									
US Dollar	-0.0361	-15.3823	-0.3448	-0.0522	-0.0247	0.0010	0.2361	2.7028	0.0532
Euro	-0.0164	-2.5309	-0.2546	-0.0280	-0.0122	0.0016	0.2107	1.7163	0.0297
All currencies	-0.0240	-9.4688	-0.2621	-0.0434	-0.0161	0.0009	0.1732	1.2186	0.0443
Panel C: Issuer Type									
SSA	-0.0164	-1.0855	-0.2335	-0.0301	-0.0120	0.0012	0.1538	2.7028	0.0313
Municipals	-0.0378	-1.9085	-0.2829	-0.0644	-0.0369	-0.0115	0.2251	0.7341	0.0529
Corporate	-0.0247	-15.3823	-0.6033	-0.0250	-0.0075	0.0121	0.2545	1.7163	0.0371
All issuers	-0.0260	-5.8475	-0.2638	-0.0446	-0.0197	-0.0024	0.1584	2.5707	0.0422

$$\text{green bond premium}_p = \alpha + \sum_{k=1}^{K} \beta_k \, \mathbf{1}\{\text{ForkChoice}_{k,p}\} + \varepsilon_p. \tag{5}$$

Equation (5) relates the mean green bond premium implied by path p to indicator variables capturing the cross-sectional sample fork choices that define that path. The term $\mathbf{1}\{\text{ForkChoice}_{k,p}\}$ is an indicator variable that equals one if specification path p selects choice k and zero otherwise. The index $k = 1, \ldots, K$ runs over all cross-sectional sample fork choices. The coefficients β_k therefore measure how selecting a given choice shifts the implied green bond premium relative to the omitted reference category. The error term ε_p captures residual variation across specification paths.

In addition to the baseline specification in Equation (5), we estimate one alternative model. We include fixed effects for all remaining fork-paths, isolating variation driven exclusively by cross-sectional sample forks. Together, these specifications allow us to assess the relative importance of cross-sectional sample fork choices in shaping green bond premium estimates.

Table 7 quantifies how individual cross-sectional sample choices shape green bond premium estimates across specification paths. The regression results confirm the patterns documented in Table 6. Currency and issuer type emerge as the most influential dimensions. Relative to the unrestricted benchmarks (i.e., all currencies and all issuers), restricting the sample to USD-denominated bonds lowers the green bond premium by 1.28 bps, while municipal bonds are associated with a 1.23 bps more negative estimate. By comparison, the choice of green bond taxonomy has economically small effects.

Table 7: Cross-Sectional Drivers

This table reports regressions of green bond premium estimates on the cross-sectional fork variables (green bond definition, currency, and issuer type), treated as categorical factors. The most inclusive category of each factor serves as the reference group. For each feasible fork-path, the green bond premium corresponds to the average of the green bond premium estimates across that path. Model (1) reports heteroskedasticity-robust standard errors (in parentheses), Model (2) includes fixed effects at the remaining fork-path level with clustered standard errors (in parentheses) at the same level. Statistical significance is indicated as *p<.10, **p<.05, and ***p<.01.

	Green Bond Premium	
	(1)	(2)
Intercept	-0.0241***	
	(0.0005)	
Green Bond: ICMA	0.0025***	0.0023***
	(0.0005)	(0.0007)
Green Bond: CBI aligned	-0.0024***	-0.0024***
	(0.0007)	(0.0001)
Currency: only Euro	0.0062***	0.0038***
	(0.0004)	(0.0007)
Currency: only USD	-0.0119***	-0.0128***
	(0.0006)	(0.0007)
Issuer Type: Corporates	0.0016*	0.00009
	(0.0009)	(0.0013)
Issuer Type: Municipals	-0.0103***	-0.0123***
	(0.0004)	(0.0007)
Issuer Type: SSA	0.0093***	0.0084***
	(0.0003)	(0.0005)
Fork-Path FE	No	Yes
Observations	518,262	518,262
R^2	0.00383	0.16905

3.3 Time-series Drivers

Having examined how cross-sectional specification choices shape the distribution of green bond premium estimates, we now turn to its temporal dynamics. Figure 3 plots the mean and median daily green bond premium over the full sample period. The early part of the sample exhibits substantial time-series variation. Over time, dispersion declines and both the mean and median converge toward values fluctuating moderately around zero, consistent with a more stable green bond premium in later years. This pattern is partly attributable to the expanding bond universe, which increases cross-sectional breadth, improves statistical precision, and reduces estimation noise. Nevertheless, some variation persists in more recent periods. Such residual volatility may reflect shifts in investor attention to climate-related risks, which can temporarily affect pricing and induce short-run movements in the green bond premium.

Figure 3: Dynamics of the Daily Green Bond Premium

This figure traces the evolution of the daily green bond premium over the sample period from 2008 to 2024. It is based on daily green bond premium estimates computed using time-specific aggregation over the full sample horizon. The figure reports the daily mean (light green), median (dark green), and interquartile range (light grey) of the green bond premium estimates. In addition, it displays the average number of matched bond pairs used to compute the daily green bond premium estimates (blue).

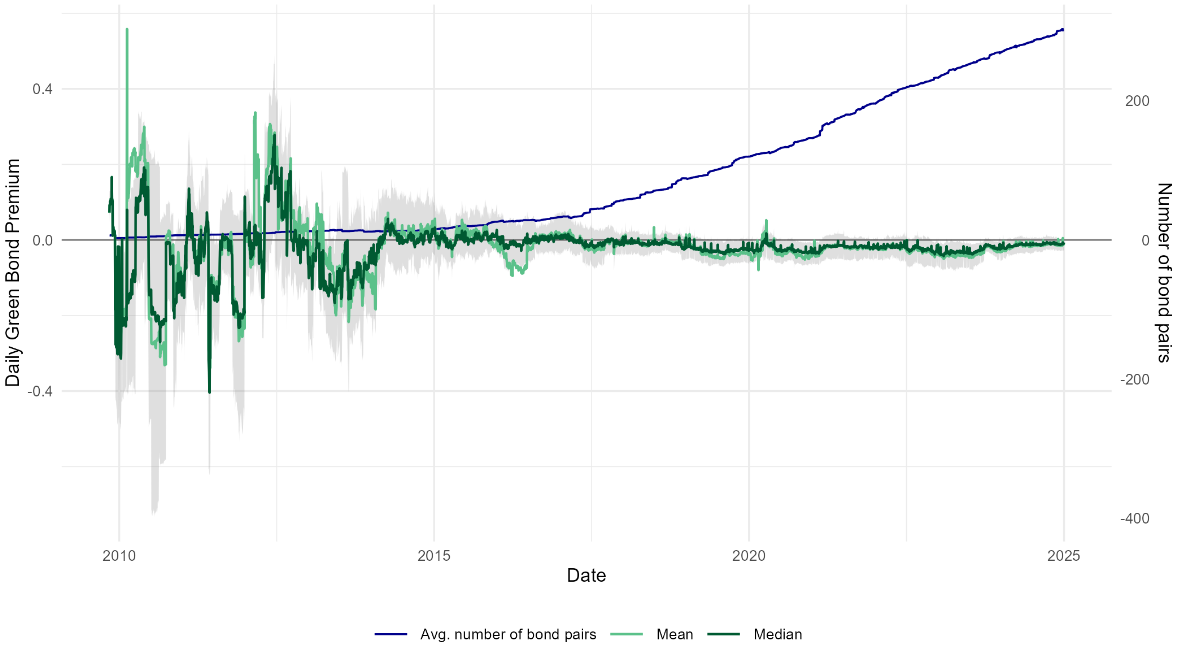

We therefore investigate whether these time-series movements are systematically related to shifts in climate-related information flow and broader measures of market risk and uncertainty. Theory predicts lower expected returns for sustainable assets either because investors derive non-pecuniary utility from holding them or because green assets hedge climate-related risks (Pástor et al. 2021, Pástor et al. 2022). While both mechanisms can rationalize lower returns in equity markets, the risk-based channel is likely limited in the context of matched green and conventional bonds. When bonds are matched on issuer, maturity, seniority, and collateral, they represent claims on the same underlying cash flows and therefore share exposure to default and systematic risk. This institutional setting provides a relatively clean environment to assess the role of investor preferences in driving time variation in the green bond premium.

To this end, we regress the daily green bond premium on several sentiment and risk proxies. By including fork-path fixed effects, we absorb all cross-specification differences arising from sample composition and methodological design, thereby isolating the relationship between these variables and within-specification time variation in the green bond premium. This identification strategy ensures that the estimated coefficients capture purely temporal fluctuations within a given fork-path, rather than differences across alternative research designs.

Our sentiment measures capture distinct dimensions of climate-related information flow. We use the Media Climate Change Coverage index (MCCC; Ardia et al. 2023), which tracks climate-change coverage in major newspapers, and the Climate Attention Index (CAI; Arteaga Garavito et al. 2025), a Twitter-based measure of global climate-related discourse. In addition, we consider broader market-based proxies for risk and uncertainty: the CBOE Volatility Index (VIX), the ICE BofA MOVE Index (MOVE) to capture volatility in the bond market, and the Equity Market-related Economic Uncertainty Index (EUI) obtained from FRED. The regressions are based on the following equation:

$$\text{green bond premium}_{p,t} = \alpha + \beta\, Z_t + \gamma_p + \varepsilon_{p,t}, \tag{6}$$

Equation (6) relates the daily green bond premium of path p on day t to a standardized climate attention or uncertainty variable Z_t (MCCC, CAI, VIX, MOVE, or EUI). The term γ_p represents fork-path fixed effects, which absorb all time-invariant differences across specification paths arising from sample composition and methodological design. The error term $\varepsilon_{p,t}$ captures residual variation. Standard errors are clustered at the fork-path and date levels. We estimate the regressions for the full sample period, for a pre-2018 period, and for a post-2017 period to account for differences in sample depth and potential heterogeneity in effects across subperiods.

Table 8 Panel A reports the regression results for the full sample period. Model (1) shows a negative and statistically significant coefficient of -0.0010 on the MCCC index, implying that a one-standard-deviation increase in newspaper-based climate attention is associated with a 0.1 basis point decline in the green bond premium. Model (2) reveals an even stronger relationship for social media attention: the coefficient of -0.0043 indicates that a one-standard-deviation increase in the CAI corresponds to a 0.43 basis point more negative green bond premium. The magnitude of this effect is economically meaningful and suggests that shifts in public climate attention are closely linked to movements in green bond premia. Models (3)-(5) show that this pattern extends to broader measures of market uncertainty. In Model (3), the coefficient on the VIX equals -0.0015, indicating that a one-standard-deviation increase in expected equity market volatility is associated with a 0.15 basis point decline in the green bond premium. Similarly, the MOVE index in Model (4) enters with a coefficient of -0.0024, corresponding to a 0.24 basis point reduction, while the EUI in Model (5) yields a coefficient of -0.0015, or a 0.15 basis point decline. When all explanatory variables are included jointly, only the coefficient on CAI remains negative and statistically significant. This attenuation of the other coefficients likely reflects multicollinearity among the attention and uncertainty measures.

The results for the post-2017 period in Panel C largely confirm those observed for the full sample period. However, in the pre-2018 period shown in Panel B, only the uncertainty measures behave similarly to the full sample results, while the coefficients for the climate attention measures reverse sign. This pattern is more plausibly attributable to the smaller sample size and limited number of successful green-conventional bond matches in the pre-2018 period than to an underlying economic effect. These findings underscore the need for cautious interpretation of results from the early subsample.

Taken together, these findings indicate that periods of heightened climate attention and elevated financial market uncertainty are systematically associated with a more negative green bond premium. One plausible interpretation is that heightened climate attention strengthens investor demand for green-labeled securities through preference-based or mandate-driven channels. As climate-related topics become more salient, ESG-oriented investors and institutions may reallocate capital toward green bonds, while supply remains relatively fixed in the short run. This demand shift raises green bond prices relative to otherwise identical conventional bonds, leading to lower relative yields and a more negative green bond premium. Similarly, periods of elevated financial market uncertainty may amplify the role of non-pecuniary preferences and institutional constraints in portfolio allocation. In such environments, investors may place greater weight on sustainability considerations, reinforcing relative demand for green bonds and widening the yield differential.

Table 9 assesses the robustness of the previous findings using an aggregation approach. Specifically, we first average the green bond premium estimates across all fork-paths within each day to obtain a single aggregate green bond premium time series. We then regress this daily series on the sentiment and market-wide variables. Because this procedure yields only one observation per day, it does not permit the inclusion of fork-path fixed effects and therefore abstracts from specification-level heterogeneity. Despite this more restrictive setup, the results closely mirror our baseline estimates. The coefficients remain negative

Table 8: Time Variation - Daily Green Bond Premium

This table reports regression results relating the green bond premium to sentiment and market-wide risk variables. We only consider fork-paths with time-specific aggregation. **Panel A** reports results for the full sample, while **Panels B** and **C** report results for the pre-2018 (including 2017) and post-2017 (starting in 2018) subsamples, respectively. MCCC is a newspaper-based climate-change attention index (Ardia et al. 2023), CAI is a Twitter-based climate-attention index (Arteaga Garavito et al. 2025). VIX and MOVE capture expected volatility in equity and bond markets, respectively, and EUI denotes the Equity Market-related Economic Uncertainty Index from FRED. All explanatory variables are z-standardized. All models include fork-path fixed effects and report standard errors clustered by fork-path and date (in parentheses). Statistical significance is indicated as *p<.10, **p<.05, and ***p<.01.

	Green Bond Premium					
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: Full Sample Period			
MCCC	-0.0010***					0.00009
	(0.0004)					(0.0004)
CAI		-0.0043***				-0.0044***
		(0.0005)				(0.0007)
VIX			-0.0015**			0.0006
			(0.0006)			(0.0007)
MOVE				-0.0024***		0.00007
				(0.0003)		(0.0004)
EUI					-0.0015***	-0.0007
					(0.0005)	(0.0007)
Fork-Path FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	217,651,574	154,720,749	212,180,454	203,955,055	217,651,574	147,144,227
R^2	0.07991	0.09847	0.08016	0.08062	0.07992	0.09899
			Panel B: Pre-2018			
MCCC	0.0030**					0.0016
	(0.0013)					(0.0013)
CAI		0.0068***				0.0044*
		(0.0021)				(0.0026)
VIX			-0.0102***			-0.0224***
			(0.0028)			(0.0030)
MOVE				-0.0094***		0.0106***
				(0.0022)		(0.0018)
EUI					-0.0018	0.0037**
					(0.0017)	(0.0016)
Fork-Path FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	57,063,751	41,355,912	55,205,750	49,651,254	57,063,751	38,661,406
R^2	0.16693	0.21705	0.16677	0.17636	0.16691	0.21847
			Panel C: Post-2017			
MCCC	-0.0012***					-0.0003
	(0.0003)					(0.0005)
CAI		-0.0048***				-0.0048***
		(0.0006)				(0.0008)
VIX			0.0010*			0.0029***
			(0.0005)			(0.0007)
MOVE				-0.0014***		0.00003
				(0.0003)		(0.0004)
EUI					-0.0007	-0.0023***
					(0.0005)	(0.0008)
Fork-Path FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	160,587,823	113,364,837	156,974,704	154,303,801	160,587,823	108,482,821
R^2	0.12719	0.13911	0.12995	0.12905	0.12713	0.14246

Table 9: Time Variation - Mean Green Bond Premium

This table reports regression results relating the green bond premium to sentiment and market-wide risk variables. We only consider fork-paths with time-specific aggregation. The dependent variable is the average daily green bond premium across fork-paths. **Panel A** reports results for the full sample, while **Panels B** and **C** report results for the pre-2018 (including 2017) and post-2017 (starting in 2018) subsamples, respectively. MCCC is a newspaper-based climate-change attention index (Ardia et al. 2023), CAI is a Twitter-based climate-attention index (Arteaga Garavito et al. 2025). VIX and MOVE capture expected volatility in equity and bond markets, respectively, and EUI denotes the Equity Market-related Economic Uncertainty Index from FRED. All explanatory variables are z-standardized. All models report heteroskedasticity-robust standard errors in parentheses. Statistical significance is indicated as *p<.10, **p<.05, and ***p<.01.

	Mean Green Bond Premium					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample Period						
Intercept	-0.0179***	-0.0142***	-0.0201***	-0.0128***	-0.0193***	-0.0145***
	(0.0015)	(0.0005)	(0.0014)	(0.0009)	(0.0012)	(0.0006)
MCCC	-0.0025**					0.0011*
	(0.0011)					(0.0006)
CAI		-0.0094***				-0.0103***
		(0.0005)				(0.0007)
VIX			-0.0067***			0.0002
			(0.0017)			(0.0007)
MOVE				-0.0057***		0.0027***
				(0.0005)		(0.0004)
EUI					-0.0042***	-0.0013*
					(0.0011)	(0.0007)
Observations	3,954	2,153	3,832	3,204	3,954	2,034
R^2	0.00104	0.15122	0.00444	0.01168	0.00160	0.16638
Panel B: Pre-2018						
Intercept	-0.0142***	-0.0097***	-0.0193***	-0.0019	-0.0169***	-0.0109***
	(0.0022)	(0.0015)	(0.0033)	(0.0028)	(0.0027)	(0.0028)
MCCC	0.0015					0.0033**
	(0.0028)					(0.0015)
CAI		-0.0060***				-0.0030
		(0.0020)				(0.0029)
VIX			-0.0156***			-0.0145***
			(0.0045)			(0.0023)
MOVE				-0.0099**		0.0205***
				(0.0042)		(0.0019)
EUI					-0.0076**	0.0017
					(0.0035)	(0.0014)
Observations	2,127	848	2,052	1,455	2,127	791
R^2	0.00011	0.00732	0.00944	0.00415	0.00137	0.08308
Panel C: Post-2017						
Intercept	-0.0226***	-0.0189***	-0.0241***	-0.0240***	-0.0237***	-0.0196***
	(0.0005)	(0.0007)	(0.0003)	(0.0003)	(0.0003)	(0.0008)
MCCC	-0.0014***					-0.0003
	(0.0003)					(0.0005)
CAI		-0.0054***				-0.0053***
		(0.0006)				(0.0008)
VIX			0.0010*			0.0029***
			(0.0005)			(0.0007)
MOVE				-0.0014***		0.00007
				(0.0003)		(0.0004)
EUI					-0.0008	-0.0023***
					(0.0005)	(0.0008)
Observations	1,827	1,305	1,780	1,749	1,827	1,243
R^2	0.01085	0.06386	0.00374	0.01650	0.00244	0.08025

and statistically significant across the climate attention and uncertainty measures for all time periods. Although the magnitudes are somewhat larger in absolute value, they remain economically modest, staying below one basis point per one-standard-deviation increase in the explanatory variables. The consistency of the results indicates that the documented relationship between climate attention, market conditions, and the green bond premium is not driven by within-path panel structure but also holds in the aggregate time series.

Taken together, these findings show that the green bond premium varies systematically with both climate-specific attention and broader measures of market uncertainty. The fact that climate attention exhibits particularly strong explanatory power supports a preference-based interpretation of the green bond premium. At the same time, the sensitivity to general uncertainty measures suggests that the green bond premium also co-moves with broader shifts in investor sentiment and risk conditions.

4 Design Uncertainty

To quantify the magnitude and sources of variation in green bond premium estimates, we examine design uncertainty in three complementary steps. First, we measure local sensitivity using mean absolute differences (MAD), which quantify how much estimates change when individual fork choices are varied in isolation. Second, we assess the global importance of forks using a Shapley decomposition, which attributes the overall variation in green bond premium estimates to individual forks. This approach accounts for interactions and correlations between fork choices. Third, we isolate sample and methodological variation by holding either the sample or the methodology fixed, allowing us to characterize how green bond premium estimates vary within each dimension separately.

4.1 MAD

To identify which fork induces the largest variation in green bond premium estimates, we calculate mean absolute differences (MAD) following Walter et al. (2024). For each fork f, we define the set S_f of matched paths that differ only in fork f as follows: $S_f = \{(i,j) \mid c_{i,m} = c_{j,m} \; \forall \; m \in \{1,\dots,14\} \setminus f, \; c_{i,f} \neq c_{j,f}\}$. For each pair, we compute the absolute difference in their green bond premium estimates. The MAD for fork f is then defined as:

$$\mathrm{MAD}_f = \frac{1}{|S_f|} \sum_{(i,j) \in S_f} |\bar{g}_i - \bar{g}_j| \,. \tag{7}$$

Here, $\bar{g}_i$ and $\bar{g}_j$ denote the green bond premium estimates from specification paths i and j. A larger MAD_f indicates that altering fork f leads to substantial shifts in the estimated green bond premium, implying that this decision materially contributes to design uncertainty. Conversely, a small MAD_f suggests that the green bond premium is relatively insensitive to changes in that fork. We compute MAD_f for each fork using all valid matched path pairs. The results are reported in Table 10.

Table 10: Mean Absolute Differences in Green Bond Premium Estimates Across Forks

This table reports the average MAD in green bond premium estimates across pairs of specification paths that differ only in a given fork. MAD is defined as in Equation (7) and given in %. The forks are arranged in descending order of MAD.

Fork	Classification	MAD
Credit Rating	Exact Matching	0.061
Ratio	Matching	0.061
Issuer Types	Sample	0.060
Coupon Rate	Threshold	0.058
Time Horizon	Sample	0.056
Currency	Sample	0.051
Aggregation	Estimation	0.044
Method	Matching	0.040
Maturity Date	Threshold	0.038
Issuance Amount	Threshold	0.037
Issuance Date	Threshold	0.030
Liquidity Adjustment	Estimation	0.023
Green Bond Definition	Sample	0.016
Yield	Estimation	0.010

Table 10 reveals substantial heterogeneity in the sensitivity of green bond premium estimates across decision forks. The largest MAD arise from matching and sample-composition forks. Credit rating matching and the matching ratio both yield MAD values of 0.061 (i.e., 6.1 basis points), followed closely by issuer type (0.060). These shifts exceed the unconditional mean green bond premium of -2.59 basis points in absolute terms, indicating that these forks are first-order determinants of the estimated premium and can materially alter its magnitude. Threshold decision forks also generate economically meaningful variation. In particular, coupon rate (0.058) and maturity date (0.038) restrictions introduce sizeable dispersion across specification paths.

In contrast, estimation choices play a more limited role. Aggregation (0.044), liquidity adjustment (0.023), and yield specification (0.010) induce comparatively smaller shifts. The green bond definition itself has only a minor impact (0.016), confirming that the estimated premium is largely insensitive to the specific classification standard. Overall, the results show that cross-specification dispersion in green bond premium estimates is driven primarily by matching design and sample composition, whereas estimation and definitional choices contribute comparatively little to design uncertainty.

4.2 Shapley decomposition

To complement the MAD analysis, we assess the relative importance of individual forks using a Shapley value decomposition based on the Lindeman-Merenda-Gold (LMG) metric. This approach decomposes the total explained variation in green bond premium estimates R^2 into additive contributions attributable to each fork, averaged across all possible orderings of variable inclusion. In contrast to the MAD, which measures the average absolute effect of pairwise changes in fork choices on the premium, the Shapley decomposition quantifies the marginal contribution of each fork to overall explanatory power in the presence of all other forks.

Figure 4 shows the results of the Shapley decomposition. The matching ratio emerges as the most important contributor, accounting for 23.0% of the explained variance. This indicates that the choice between 1:1 matching, forced interpolation, and extrapolation is the single most influential methodological decision in shaping green bond premium estimates. The issuance-date threshold ranks second (20.3%), followed by currency (13.4%), issuer type (12.5%), and time horizon (10.5%). These results underscore that matching design and sample composition are the primary drivers of cross-specification variation, consistent with the MAD analysis. Exact credit-rating matching also contributes meaningfully, explaining 8.4% of the variance.

By contrast, several forks add little incremental explanatory power. Yield specification accounts for only 0.2% of the explained variance, while aggregation level (0.5%), coupon-rate threshold (0.7%), matching method (0.8%), and green bond certification (0.9%) each contribute less than 1%. These findings indicate that once the most influential forks are taken into account, these additional design choices play only a minor role in determining the magnitude of the estimated green bond premium.

4.3 Isolating Sample and Methodological Design Uncertainty

The full fork-path analysis combines variation in both sample composition and methodological choices. To complement this analysis, we next consider restricted settings in which either the sample or the methodology is held constant. By varying only the sample or only the methodological forks, we characterize how green bond premium estimates vary within each dimension separately and assess whether the patterns observed in the full analysis

Figure 4: Relative Importance of Forks

This figure shows the results of a Shapley decomposition based on LMG where we regress the green bond premium estimates on their respective fork choices. The green bars show how much each fork contributes to the total explained variation in green bond premium estimates.

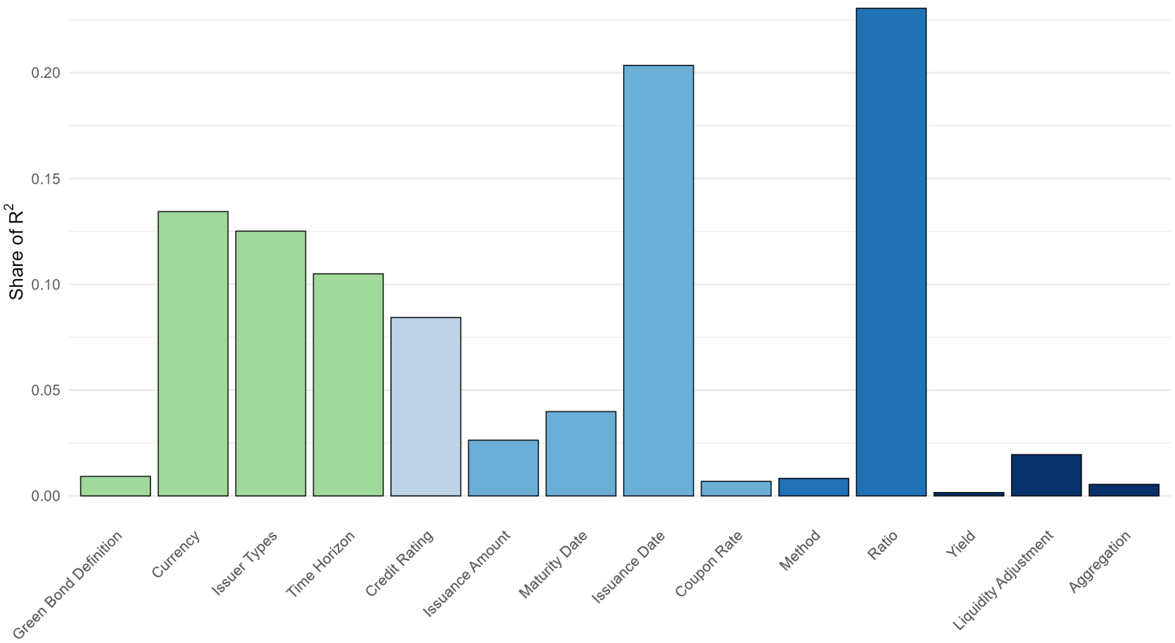

persist.

First, we adopt a single, representative methodological path, selected on the basis of the prior fork-path assessment and aligned with standard practice in the literature. The fixed specification includes no exact match on credit rating, a two-year maturity and issuance-date threshold, an issuance-amount filter allowing values between one-half and twice the green bond's amount (log(2)), no coupon restriction, simple matching based on the smallest maturity difference with forced interpolation, ask yields, and bond-level aggregation using a liquidity-adjusted green bond premium metric. Holding this methodological configuration constant across all sample variants ensures that any variation in the resulting green bond premium estimates can be attributed exclusively to differences in sample composition.

Table 11 reports the distribution of green bond premium estimates when the methodology is held constant. The distribution closely resembles the base result but shows fewer outliers. The mean is -2.65 bps and the median is -2.33 bps, indicating a small but persistent negative green bond premium.

Figure 5 displays boxplots of the green bond premium for all sample forks. The boxplots reveal large variation depending on the sample choices. Earlier time periods, specific issuer

Table 11: Summary Statistics of Green Bond Premium Estimates with a Fixed Methodology

This table reports the distribution of the mean green bond premium estimates with a fixed methodology in %. IQR represents the interquartile range.

Mean	Min	1%	25%	Median	75%	99%	Max	IQR
-0.0265	-0.252	-0.2333	-0.0397	-0.0233	-0.0112	0.0671	0.0922	0.0285

types (i.e., corporates and municipals), and USD-denominated bonds show wider interquartile ranges and longer tails. In contrast, EUR-denominated bonds, SSAs, and post-2018 samples produce tighter and more stable distributions. These differences indicate that the dispersion of the green bond premium is highly sensitive to the sample composition. Consistent with the full fork-path analysis, these findings indicate that heterogeneity in reported green bond premium estimates largely reflects differences in the underlying sample rather than sensitivity to specific methodological choices.

Figure 5: Impact of Fixing Specific Sample Forks with a Fixed Methodology

This figure shows boxplots of the distribution of the green bond premium when fixing a specific sample fork with a fixed methodology. The remaining forks each vary in their choices as defined in Table 3. The red stars indicate the mean of the distribution.

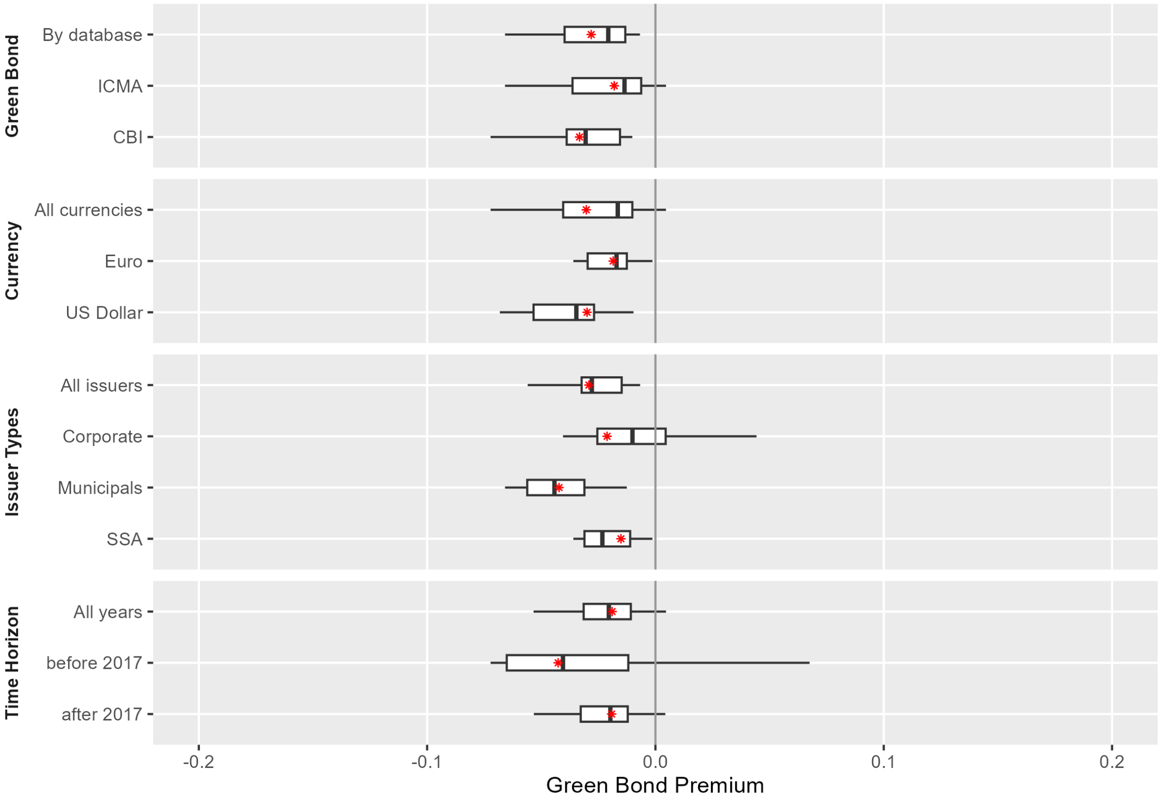

Having isolated the influence of sample choices in the previous analysis, we now examine the complementary question: how sensitive is the estimated green bond premium to variation in methodology when the sample is fixed? We fix the sample using the broadest available inclusion criteria. The time horizon is set to the full sample period, and green bonds are identified according to the LSEG Workspace classification. All currencies and all issuer types are included. We further ensure that every green bond can be matched to at least one conventional (synthetic) bond in every fork-path, meaning that all selected green bonds have suitable conventional counterparts even under the most restrictive matching specifications. This guarantees that the underlying set of bonds remains constant across all paths, allowing us to attribute any variation in the resulting green bond premium estimates solely to differences in methodological choices.

Table 12 reports the distribution of green bond premium estimates when the sample is held constant. The mean and median green bond premium, -5.45 bps and -5.57 bps, indicate that the distribution is shifted toward more negative values compared to the full fork-path analysis. The interquartile range is 2.87 bps, showing that the estimates remain tightly clustered despite variation in methodological choices. Overall, the distribution remains narrow but distinctly negative, confirming that all feasible methodological specifications within this fixed sample produce consistently negative green bond premium estimates.

Table 12: Summary Statistics of Green Bond Premium Estimates with a Fixed Sample

This table reports the distribution of the mean green bond premium estimates in a fixed sample in %. IQR represents the interquartile range.

Mean	Min	1%	25%	Median	75%	99%	Max	IQR
-0.0545	-0.6442	-0.1562	-0.0674	-0.0557	-0.0387	0.0297	0.6438	0.0287

Figure 6 summarizes the influence of the methodology forks on the estimated green bond premium. Across all threshold forks, the boxplots are tightly clustered and centered around negative values, indicating that the green bond premium remains consistently negative regardless of the specific threshold choice. Differences in medians between alternative thresholds within a fork are generally small, and the interquartile ranges largely overlap. An exception is the issuance-date threshold: applying the strictest criterion of two years produces a tighter and more negative distribution compared to the more permissive alternatives, consistent with the regression and Shapley results. Matching with or without credit ratings produces very similar distributions, indicating that rating alignment has little influence on the results. In contrast, the choice of matching method matters more. Propensity score matching generates a noticeably wider interquartile range and overall dispersion, whereas simple closest-maturity matching leads to more concentrated and consistent estimates. The matching ratio also affects the distribution: 1:2 matching with extrapolation produces a broader spread, while both 1:1 matching and 1:2 matching with interpolation yield tighter and more stable green bond premium estimates. The choice of ask, bid, or mid yields makes almost no difference. Liquidity adjustment changes the distribution only slightly, even though the literature often stresses its importance. Within the estimation forks, the largest variation arises when fixing the aggregation fork. The median green bond premium is less negative and the dispersion of the estimates is greater when the green bond premium is computed as a time-varying measure. These results indicate that the negative green bond premium is not driven by any single methodological choice. When the sample is held constant, the estimated green bond premium remains negative and tightly distributed across a wide range of estimation approaches.

Figure 6: Impact of Fixing Specific Methodology Forks with a Fixed Sample

This figure shows boxplots of the distribution of the green bond premium when fixing a specific methodology fork in a fixed sample. The remaining forks each vary in their choices as defined in Table 3. The red stars indicate the mean of the distribution.

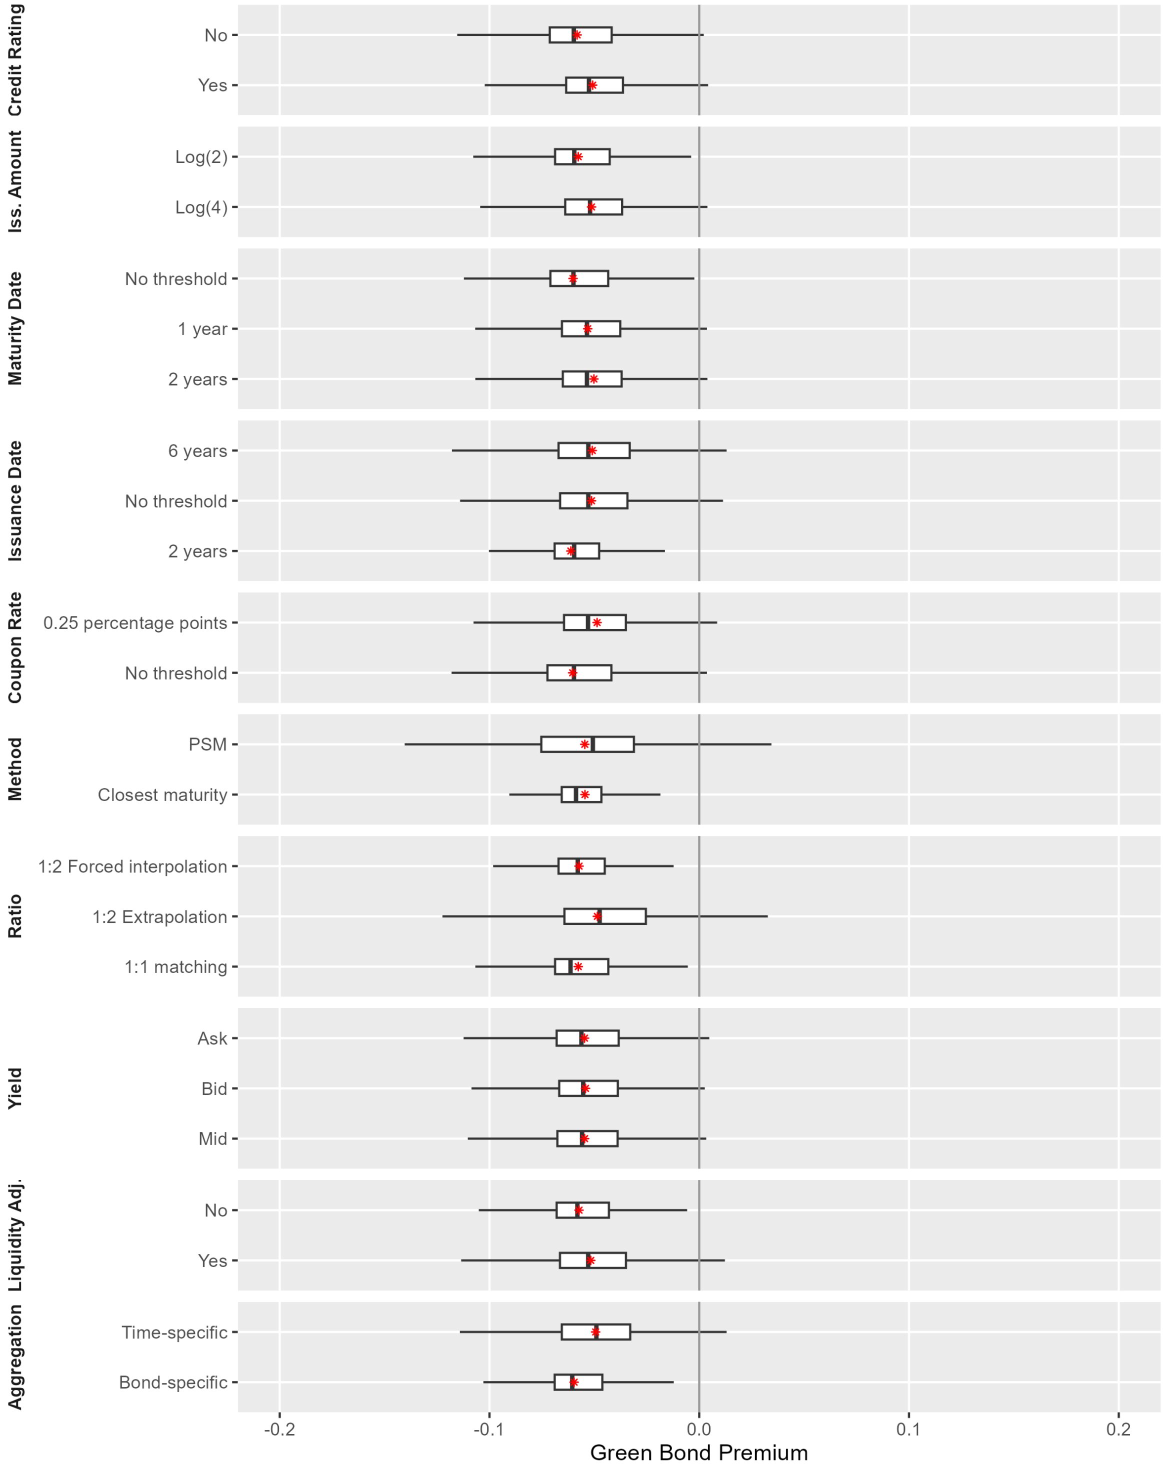

5 Conclusion

This paper examines the pricing of green bonds in secondary markets using the matched bond pairs approach and asks whether the green bond premium reflects a robust market phenomenon or is primarily driven by empirical design choices. To address this question, we estimate the green bond premium across more than 500,000 empirical designs that span the key sample and methodological choices in the literature.

Across this extensive design space, the green bond premium is on average negative but economically modest. The distribution of estimates is centered around -2.59 basis points, although individual estimates vary depending on sample composition and matching choices. The size of the mean premium is especially dependent on issuer type, currency, and the sample period. We further document pronounced time variation in the green bond premium. Daily movements are closely linked to fluctuations in climate-related attention, with both newspaper-based and social-media-based measures exhibiting explanatory power. These results suggest that the green bond premium behaves as a dynamic pricing component that responds to shifts in investor attention and sustainable investment demand rather than as a fixed structural spread.

When analyzing the influence of methodological choices on the magnitude and sign of the premium, we find that the matching ratio is a key driver of dispersion, whereas some commonly debated methodological choices such as liquidity adjustment contribute comparatively little to overall variation. When the sample is held constant, the premium remains consistently negative across a wide range of estimation approaches, indicating that it is not an artifact of specific empirical implementations.

Our findings carry implications for issuers, investors, and researchers. For issuers, green labeling is associated with a small but consistent yield advantage in secondary markets that may lower funding costs at issuance if primary and secondary markets are integrated. This advantage, however, is neither large nor stable: it concentrates in municipal and USD-denominated segments and amplifies when climate attention is elevated. Issuers who expect a reliable financing benefit from green certification may be disappointed outside of high-attention regimes. For investors, the greenium implies that green bonds carry a return cost relative to otherwise identical conventional bonds that fluctuates with market sentiment and is largest when sustainability is most salient. For researchers, our central finding is that the wide dispersion in published greenium estimates is driven overwhelmingly by sample composition rather than by methodological choices that have received the most attention in the literature. Future work should treat sample heterogeneity as a first-order concern.

A Appendix

Table A1: Variable Definitions

Variable	**Definition**	**Data Source**
Issuer	This variable identifies the issuing entity of the bond, represented by its ticker symbol.	LSEG Workspace
Issuer Type	This variable classifies issuers into Corporate, Municipals, and SSA (Supranational, Sovereign, and Agency) categories.	LSEG Workspace
Currency	This variable denotes the currency in which the bond's principal is denominated.	LSEG Workspace
Issuance Amount	This variable measures the total issuance size of the bond, reported in both the original currency and USD equivalents.	LSEG Workspace
Coupon Rate	This variable reports the bond's coupon rate expressed in percentage terms.	LSEG Workspace
Green Bond Definition	This variable identifies green bonds based on three alternative classifications: (i) the LSEG Workspace green bond flag, (ii) compliance with ICMA Green Bond Principles, and (iii) certification by the Climate Bonds Initiative.	LSEG Workspace
Credit Rating	This variable captures long-term credit ratings. Ratings from Moody's and Fitch are harmonized by mapping all observations to the Moody's scale. When available, Moody's ratings are used directly; otherwise, Fitch ratings are converted to their Moody's-equivalent.	LSEG Workspace
Maturity Date	This variable records the contractual maturity date of the bond.	LSEG Workspace
Issuance Date	This variable records the original issuance date of the bond.	LSEG Workspace
Bond Structure	This variable indicates the structural features of the bond. The sample is restricted to plain-vanilla bonds.	LSEG Workspace
Coupon Type	This variable identifies the coupon structure. The sample is restricted to bonds with fixed coupon payments.	LSEG Workspace

Variable	**Definition**	**Data Source**
Seniority	This variable describes the seniority ranking of the bond within the issuer's capital structure.	LSEG Workspace
Collateral	This variable specifies the type of collateral of the bond.	LSEG Workspace
Asset Status	This variable reflects the current status of the bond. Bonds in default are excluded from the sample.	LSEG Workspace
Yield	This variable reports bid, ask, and mid yields. The mid-yield is calculated as the simple average of bid and ask yields.	LSEG Workspace
Bid-Ask Spread	This variable measures market liquidity as the difference between ask and bid prices, computed as the mid-point spread.	LSEG Workspace
Green Bond Premium	This variable measures the estimated green bond premium for each fork-path specification. It is defined as the average premium obtained from the corresponding estimation (cross-sectional or time-series) for a given specification path.	Calculated
Mean Green Bond Premium	This variable reports the daily average of all estimated green bond premiums across specifications, obtained from time-specific aggregations.	Calculated
MCCC	This variable captures media-based climate attention using the Media Climate Change Coverage index by Ardia et al. (2023).	`https://sentometrics-research.com/download/mccc/`
CAI	This variable captures social media-based climate attention using the Climate Attention Index by Arteaga Garavito et al. (2025).	`https://sites.google.com/view/internationalclimatenews/download`
VIX	This variable measures expected volatility in equity markets using the CBOE Volatility Index.	FRED: `https://fred.stlouisfed.org/series/VIXCLS`

Variable	**Definition**	**Data Source**
MOVE	This variable measures expected volatility in bond markets using the ICE BofA MOVE Index.	LSEG Workspace
EUI	This variable captures economic policy uncertainty related to equity markets using the Equity Market-related Economic Uncertainty Index.	FRED: https://fred.stlouisfed.org/series/WLEMUINDXD

References

Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht (2023): “Climate change concerns and the performance of green vs. brown stocks,” *Management Science*, 69, 7607–7632.

Arteaga Garavito, M. J., R. Colacito, M. M. M. Croce, and B. Yang (2025): “International climate news,” *NBER Working Paper*.

Bachelet, M. J., L. Becchetti, and S. Manfredonia (2019): “The green bonds premium puzzle: The role of issuer characteristics and third-party verification,” *Sustainability*, 11, 1098.

Chen, A. Y. and T. Zimmermann (2022): “Open Source Cross-Sectional Asset Pricing,” *Critical Finance Review*, 11, 207–264.

Dick-Nielsen, J., P. Feldhütter, and D. Lando (2012): “Corporate bond liquidity before and after the onset of the subprime crisis,” *Journal of Financial Economics*, 103, 471–492.

Dorfleitner, G., S. Utz, and R. Zhang (2022): “The pricing of green bonds: external reviews and the shades of green,” *Review of Managerial Science*, 16, 797–834.

Dragotto, M., A. Dufour, and S. Varotto (2025): “Greenium fluctuations and climate awareness in the corporate bond market,” *International Review of Financial Analysis*, 104281.

Hachenberg, B. and D. Schiereck (2018): “Are green bonds priced differently from conventional bonds?” *Journal of Asset Management*, 19, 371–383.

Hasler, M. (2023): “Looking Under the Hood of Data-Mining,” *Working Paper*.

Hou, K., C. Xue, and L. Zhang (2020): “Replicating Anomalies,” *The Review of Financial Studies*, 33, 2019–2133.

Immel, M., B. Hachenberg, F. Kiesel, and D. Schiereck (2021): “Green bonds: shades of green and brown,” *Journal of Asset Management*, 22, 96–109.

Jensen, T. I., B. Kelly, and L. H. Pedersen (2023): “Is There a Replication Crisis in Finance?” *The Journal of Finance*, 78, 2465–2518.

Kapraun, J., C. Latino, C. Scheins, and C. Schlag (2021): “(In)-credibly green: Which bonds trade at a green bond premium?” *Working Paper*.

Karpf, A. and A. Mandel (2018): “The changing value of the ‘green’label on the US municipal bond market,” *Nature Climate Change*, 8, 161–165.

Menkveld, A. J., A. Dreber, F. Holzmeister, J. Huber, M. Johannesson, M. Kirchler, and et al. (2024): “Non-Standard Errors,” *The Journal of Finance*, 79, 2339–2390.

MERLI, M., J. PETEY, AND T. ROGER (2025): "Climate concerns, salient events, and green preferences," *Economics Letters*, 112534.

NANAYAKKARA, M. AND S. COLOMBAGE (2019): "Do investors in green bond market pay a premium? Global evidence," *Applied Economics*, 51, 4425–4437.

PARTRIDGE, C. AND F. R. MEDDA (2020): "Green bond pricing: The search for greenium," *Journal of Alternative Investments*, 23, 49–56.

PÁSTOR, L., R. F. STAMBAUGH, AND L. A. TAYLOR (2021): "Sustainable investing in equilibrium," *Journal of Financial Economics*, 142, 550–571.

——— (2022): "Dissecting green returns," *Journal of Financial Economics*, 146, 403–424.

WALTER, D., R. WEBER, AND P. WEISS (2024): "Methodological uncertainty in portfolio sorts," *Working Paper*.

ZERBIB, O. D. (2019): "The effect of pro-environmental preferences on bond prices: Evidence from green bonds," *Journal of Banking & Finance*, 98, 39–60.

cfr working paper series

centre for financial research
cologne

cfr working papers are available for download from www.cfr-cologne.de.

2026

No.	Author(s)	Title
26-06	T. Bauckloh, P. Kirsch	The Green Bond Premium: Evidence From a Multiverse Analysis
26-05	S. Weibels	Hard to Process: Atypical Firms and the Cross-Section of Expected Stock Returns
26-04	B. Ballensiefen, F. Somogyi, H. Winterberg	Demand for Dollars: Evidence from Survey Expectations
26-03	J. Fausch, M. Frigg, S. Ruenzi, F. Weigert	Machine Learning Mutual Fund Flows
26-02	T. G. Bali, A. Goyal, M. Mörke, F. Weigert	In Search of Seasonality in Intraday and Overnight Option Returns
26-01	S. Müller, N. Pugachyov, F. Weigert	Forecasting Mutual Fund Performance – Combining Return-Based with Portfolio Holdings-Based Predictors

2025

No.	Author(s)	Title
25-11	D. Aobdia, G. Köchling, P. Limbach, A. Yoon	Emissions Restatements after the SEC's Request for Public Input on Climate-Related Disclosures: Evidence from Carbon Disclosure Project Filings
25-10	V. Agarwal, J.-P. Gomez, K. Hosseini, M. Jha	ESG Metrics in Executive Compensation: A Multitasking Approach
25-09	V. Agarwal, P. Ghosh, N. Prabhala, H. Zhao	Animal Spirits on Steroids: Evidence from Retail Options Trading in India
25-08	V. Agarwal, S. Cao, S. Huang, M. Kim	Incentive Realignment: Mutual Funds' Influence on Executive Compensation Contracts
25-07	V. Agarwal, Y. E. Arisoy, T. Trinh	Eponymous Hedge Funds
25-06	D. Hess, F. Simon, S. Weibels	Interpretable Machine Learning for Earnings Forecasts: Leveraging High-Dimensional Financial Statement Data
25-05	B. F. Ballensiefen	Collateral Choice
25-04	C. Andres, F. Brochet, P. Limbach, N. Schumacher	Sell-Side Analysts with Accounting Experience
25-03	W. Bazley, G. Cici, J. Liao	Conflicts of Interest among Affiliated Financial Advisors in 401(k) Plans: Implications for Plan Participants

25-02	A. T. Maître, N. Pugachyov, F. Weigert	Twitter-Based Attention and the Cross-Section of Cryptocurrency Returns
25-01	N. Käfer, M. Mörke, F. Weigert, T. Wiest	A Bayesian Stochastic Discount Factor for the Cross-Section of Individual Equity Options

2024

No.	Author(s)	Title
24-06	V. Beyer, T. Bauckloh	Non-Standard Errors in Carbon Premia
24-05	C. Achilles, P. Limbach, M. Wolff, A. Yoon	Inside the Blackbox of Firm Environmental Efforts: Evidence from Emissions Reduction Initiatives
24-04	Ivan T. Ivanov, T. Zimmermann	The “Privatization” of Municipal Debt
24-03	T. Dyer, G. Köchling, P. Limbach	Traditional Investment Research and Social Networks: Evidence from Facebook Connections
24-02	A. Y. Chen, A. Lopez-Lira, T. Zimmermann	Does Peer-Reviewed Research Help Predict Stock Returns?
24-01	G. Cici, P. Schuster, F. Weishaupt	Once a Trader, Always a Trader: The Role of Traders in Fund Management

2023

No.	Author(s)	Title
23-08	A. Braun, J. Braun, F. Weigert	Extreme Weather Risk and the Cost of Equity
23-07	A. G. Huang, R. Wermers, J. Xue	“Buy the Rumor, Sell the News”: Liquidity Provision by Bond Funds Following Corporate News Events
23-06	J. Dörries, O. Korn, G. J. Power	How Should the Long-term Investor Harvest Variance Risk Premiums?
23-05	V. Agarwal, W. Jiang, Y. Luo, H. Zou	The Real Effect of Sociopolitical Racial Animus: Mutual Fund Manager Performance During the AAPI Hate
23-04	V. Agarwal, B. Barber, S. Cheng, A. Hameed, H. Shanker, A. Yasuda	Do Investors Overvalue Startups? Evidence from the Junior Stakes of Mutual Funds
23-03	A. Höck, T. Bauckloh, M. Dumrose, C. Klein	ESG Criteria and the Credit Risk of Corporate Bond Portfolios
23-02	T. Bauckloh, J. Dobrick, A. Höck, S. Utz, M. Wagner	In partnership for the goals? The level of agreement between SDG ratings
23-01	F. Simon, S. Weibels, T. Zimmermann	Deep Parametric Portfolio Policies

This document only covers the most recent cfr working papers. A full list can be found at www.cfr-cologne.de.

centre for financial research

cfr/university of cologne
Albertus-Magnus-Platz
50923 cologne
Germany
fon +49 221 470-91310
info@cfr-cologne.de
www.cfr-cologne.de