DEUTSCHE BUNDESBANK
EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 41/2024

Who pays the greenium and why? A decomposition

Daniel Fricke
Christoph Meinerding

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet http://www.bundesbank.de

ISBN 978–3–98848–016–3
ISSN 2941–7503

Non-technical summary

Research question

The transition towards a climate-neutral economy will require massive investments, both from the public and the private sector. At the same time, the market for sustainable investments is growing fast, with an increasing number of investors explicitly seeking to finance green projects. A tentative consensus has emerged that sustainable (clean, green, etc.) assets tend to have higher prices and thus lower expected returns than otherwise equivalent non-sustainable assets. This difference in yields or expected returns is often referred to as the ''greenium''. Our paper takes the research about the greenium one step further and addresses the open questions which investors pay the greenium and why they do so. To this end, we zoom in on one particularly relevant market segment that provides a unique laboratory for this purpose, namely the market for green bonds.

Contribution

We construct a unique dataset that combines matched green and conventional bond yields from the same issuer with granular information on the ownership structure of these bonds. For the latter, we draw upon the Eurosystem's confidential Securities Holdings Statistics (SHS). To single out which investors pay the greenium and why they do so, we propose a methodology that fully decomposes the greenium along the ownership structure and the cross-section of green bonds. Our study contributes to the ongoing debate about the peculiarities of the market for sustainable investments. We exploit the rich cross-sectional and time variation in both the greenium and bond ownership in order to better understand who pays for financing the green transition and why.

Results

The average yield differential between a green and a matched conventional bond (``greenium'') amounts to minus 3 basis points, in line with prior literature. Applying the first step of our decomposition methodology along the bond ownership structure, we find that banks, investment funds and insurance companies pay the bulk of the greenium. The second step of our decomposition then provides evidence for at least two distinct channels being at play. On the one hand, the greenium paid by investment funds (and their clients) is mostly explained by an average level effect, confirming the narrative that these investors have non-pecuniary sustainability preferences. They overweight green bonds relative to conventional bonds on average, across the entire green bond market. On the other hand, the greenium paid by banks is markedly different. Banks display a tilt towards a very specific subset of green bonds with a sizeable greenium, They seem to pay a significant greenium because they hold specific green bonds for motives other than general green preferences. Further analyses point towards an interaction between the greenium and bank-specific financial frictions. Overall, our results indicate that prices and quantities of sustainable assets are driven by more than the non-pecuniary green preferences of investors.

Nichttechnische Zusammenfassung

Fragestellung

Die Transition zu einer klimaneutralen Wirtschaft erfordert massive Investitionen sowohl des öffentlichen als auch des privaten Sektors. Gleichzeitig wächst der Markt für nachhaltige Investitionen und immer mehr Investoren möchten explizit zur Finanzierung grüner Projekte beitragen. In der Wissenschaft verfestigt sich ein Konsens darüber, dass nachhaltige (saubere, grüne usw.) Wertpapiere tendenziell höhere Preise und damit niedrigere erwartete Renditen haben als äquivalente nicht nachhaltige Vermögenswerte. Dieser Unterschied in der Verzinsung oder der erwarteten Rendite wird oft als „Greenium" bezeichnet. Unser Paper geht einen Schritt weiter und analysiert, welche Anleger das Greenium zahlen und warum sie dies tun. Dazu fokussieren wir uns auf ein spezifisches und relevantes Marktsegment, das für diese Art der Untersuchung besonders geeignet ist: den Markt für grüne Anleihen („Green Bonds").

Beitrag

Wir stellen einen außergewöhnlichen Datensatz zusammen, der die Zinssätze auf grüne und konventionelle Anleihen desselben Emittenten, die bezüglich einer Reihe weiterer Matching-Kriterien übereinstimmen, mit detaillierten Informationen über die Eigentümerstruktur dieser Anleihen kombiniert. Für Letzteres greifen wir auf die vertraulichen Securities Holdings Statistics (SHS) des Eurosystems zurück. Um herauszufinden, welche Investoren das Greenium zahlen und warum sie dies tun, schlagen wir eine Methodik vor, die das Greenium vollständig entlang der Eigentümerstruktur und des Querschnitts grüner Anleihen zerlegt. Unsere Studie leistet einen Beitrag zur aktuellen Debatte über die Besonderheiten des Marktes für nachhaltige Investments. Wir nutzen die vielfältige Querschnitts- und Zeitvariation sowohl beim Greenium als auch bei der Eigentümerstruktur der Anleihen, um besser zu verstehen, welche Anleger für die Finanzierung der grünen Transition zahlen und warum.

Ergebnisse

Die durchschnittliche Zinsdifferenz zwischen einer grünen und einer äquivalenten konventionellen Anleihe („Greenium") beträgt, im Einklang mit der bisherigen Literatur, minus 3 Basispunkte. Indem wir den ersten Schritt unserer Zerlegungsmethode entlang der Eigentümerstruktur der Anleihen anwenden, stellen wir fest, dass Banken, Investmentfonds und Versicherungsgesellschaften den Großteil des Greeniums zahlen. Der zweite Schritt unserer Zerlegung liefert dann Evidenz für mindestens zwei unterschiedliche Wirkungskanäle. Auf der einen Seite lässt sich das von Investmentfonds (und ihren Kunden) gezahlte Greenium größtenteils durch einen Durchschnittseffekt erklären, der das Narrativ bestätigt, dass diese Anleger nicht-monetäre Präferenzen für Nachhaltigkeit haben. Sie übergewichten grüne Anleihen im Vergleich zu konventionellen Anleihen im Durchschnitt über den gesamten Markt für grüne Anleihen hinweg. Auf der anderen Seite unterscheidet sich das von den Banken gezahlte Greenium von diesem Durchschnittseffekt deutlich. Banken sind stark investiert in ausgewählten grünen Anleihen mit einem ausgeprägten Greenium. Sie scheinen also ein Greenium zu zahlen, weil sie selektiv bestimmte grüne Anleihen aus anderen Gründen als ihren allgemein grünen Präferenzen halten. Weitere Analysen suggerieren eine Wechselwirkung zwischen dem Greenium und bankspezifischen Finanzfriktionen. Insgesamt deuten unsere Ergebnisse darauf hin, dass die Preise und Mengen nachhaltiger Wertpapiere nicht ausschließlich von den nicht-monetären grünen Präferenzen der Anleger bestimmt werden.

Deutsche Bundesbank Discussion Paper No 41/2024

Who pays the greenium and why? A decomposition

Daniel Fricke*   Christoph Meinerding*

October 17, 2024

Abstract: The average yield differential between a green and a matched conventional bond ("greenium") amounts to minus 3 basis points. We decompose this greenium along the bonds' ownership structure and document that investment funds, banks and insurance companies pay most of it. Dissecting further, the greenium paid by investment funds (and their clients) is mostly explained by an average level effect, confirming the narrative that these investors have non-pecuniary sustainability preferences. The greenium paid by banks is markedly different and cannot be explained by such preferences. Rather banks overweight specific green bonds with a sizeable greenium, pointing towards an interaction between the greenium and bank-related financial frictions.

Keywords: Green bonds, sustainable investment, greenium, ownership structure, securities holdings

JEL: G11, Q01

*Deutsche Bundesbank, e-mails: daniel.fricke|christoph.meinerding@bundesbank.de

We would like to thank Tobias Bauckloh (discussant), Martijn Boermans (discussant), Caroline Flammer, Bruno Gerard, Refet Gürkaynak, Daniel Huber (discussant), Emanuel Moench, Karol Paludkiewicz, Toan Phan, Glenn Rudebusch, Jiri Woschitz, as well as conference and seminar participants at the Bundesbank, the BI Norwegian Business School, the ESCB Research Cluster Climate Change, the Riksbank conference on Climate Change and the Financial System, the DGF conference 2023, the Conference on Climate and Energy Finance 2023 in Hannover, and the Green Finance Research Advances 2023, for valuable comments and suggestions. We also thank Stephan Jank for his contributions in early stages of this project. Preliminary versions of this paper were entitled "Who pays the greenium?" and "Who finances the green transition? Evidence from the green bond market". A supplementary online appendix is available at `https://www.bundesbank.de/dp-en/2024-41-appendix`. This work represents the authors' personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

1 Introduction

The transition towards a climate-neutral economy is one of the biggest challenges of our times and will require massive investments, both from the public and the private sector. Hence, it is no coincidence that the market for sustainable investments is growing fast, with an increasing number of investors explicitly seeking to finance green projects. A tentative consensus has emerged that sustainable (clean, green, etc.) assets tend to have higher prices and thus lower expected returns than otherwise equivalent non-sustainable assets.[1] This difference in yields or expected returns is often referred to as the "greenium" and we follow this convention. Our paper takes the research about the greenium one step further and addresses the open question which investors pay the greenium and why they do so.

To this end, we zoom in on one particularly relevant market segment that provides a unique laboratory for this purpose, namely the green bond market. A green bond is characterized by the issuer committing to use its proceeds to finance specific environmental and climate-friendly projects, and this commitment is mostly also verified externally. Focusing on green bonds has the major advantage that we can construct a unique dataset that combines matched green and conventional bond yields from the same issuer with granular information on the ownership structure of these bonds. For the latter, we draw upon the Eurosystem's confidential Securities Holdings Statistics (SHS). Green bonds tend to trade at a premium relative to otherwise comparable conventional bonds, also in our specific sample. Concerning this negative yield spread between green and conventional bonds, the green bond market is arguably representative for the larger universe of sustainable investments.

To address our two main research questions, we propose a methodology that fully decomposes the greenium along two dimensions: the ownership structure and the cross-section of green bonds. We start by splitting the greenium across different investor groups based on their green bond ownership shares, which allows us to answer the question who pays the greenium. For the second question, why different investor groups do so, we then further decompose the sector-specific greenium into three distinct subcomponents which we label as (a) the *benchmark greenium*, (b) the *green preference component*, and (c) the *bond-specific over-/underweighting component*. The first component (benchmark greenium) is the hypothetical greenium that would be paid/earned if the shares of each investor group in each green bond were exactly identical to the shares in the respective conventional twin. We view this benchmark as representing a simple null hypothesis, quantifying a sector's greenium that would be observed if the sector treated green and conventional bonds equally and thus simply replicated the ownership shares. The second component (green preferences) is the part of the greenium which is paid/earned because an investor group over-/underweights green bonds relative to conventional bonds *on average*, across the entire green bond market. The residual third component reflects the part of the greenium which is paid/earned through *bond-specific over-/underweighting* of green bonds whose greenium deviates from the average. Mathematically, it captures the covariance between the greenium and the investor group shares.

[1] See, e.g., Pedersen, Fitzgibbons, and Pomorski (2021), Pastor, Stambaugh, and Taylor (2022), Bolton and Kacperczyk (2021), Ilhan, Sautner, and Vilkov (2021).

In our empirical analyses, we document a statistically significant average greenium of around minus 3 basis points. The average green bond investor thus pays a greenium, i.e., forgoes a few basis points in bond yield as compared to an otherwise equivalent conventional investment.[2] Applying the first step of our decomposition methodology, we find that banks, investment funds and insurance companies pay the bulk of the greenium. The second step of our decomposition then provides evidence for at least two distinct channels being at play. On the one hand, the greenium paid by investment funds (and their clients) is mostly driven by the green preference component (b). Specifically, investment funds overweight green bonds relative to conventional bonds *on average*, across the entire green bond market. They have an average ownership share of 29% in the green bonds in our sample, as opposed to 24% in the respective conventional twins. Arguably, investment funds closely stick to an investment mandate (as delegated by their clients). So we take our evidence as pointing towards the prominent narrative that clients of investment funds have particularly strong non-pecuniary green preferences.[3] In this sense, our paper thus lends support to the seminal models of Pedersen et al. (2021) and Pastor, Stambaugh, and Taylor (2021). These authors show that a greenium arises in equilibrium if a subset of investors has non-pecuniary sustainability preferences, i.e. they prefer sustainable assets for reasons that are orthogonal to the risk-return trade-off which usually determines optimal portfolios.

On the other hand, the greenium paid by banks is markedly different and cannot be explained by such preferences. Banks' average ownership share in green bonds is 15% as opposed to 18% in conventional twins. Instead, our decomposition shows that banks pay the greenium due to the residual component (c). Specifically, the covariance between bond yields and ownership shares indicates that banks display a tilt towards a very specific subset of green bonds *with a sizeable greenium*. Moreover, the covariance term is particularly pronounced for young bonds, small bonds, and bonds issued by the financial sector. Notably, banks are the dominant investor group in the segment of financial bonds in general, as has been emphasized by Bekaert and Breckenfelder (2019). We draw the tentative conclusion that banks (and their clients) pay a significant greenium because they hold specific green bonds for motives other than general green preferences. The results point towards an interaction between the greenium and bank-specific financial frictions, for instance along the lines of intermediary asset pricing theory. Such an interaction has also been suggested by Larcker and Watts (2020). Importantly, though, an explanation of our findings through the lens of intermediary asset pricing could also imply a reverse causality, namely that the peculiarities of green bond holdings of banks introduce cross-sectional variation in the greenium. Besides, an interpretation based on bank-specific financial

[2] These numbers are economically relevant. According to statistics from the BIS, the total amount outstanding of green bonds was approximately 2 trillion USD in 2022-Q2. If the global green bond universe displayed a greenium of similar size as in our sample, these 3 basis points would thus correspond to roughly 600 million USD. More broadly, the European Commission announced in 2019, as part of its so-called Green Deal, that it aims to mobilize at least 1 trillion EUR in public and private investments over the course of 10 years to make Europe the first continent with net zero carbon emissions by 2050.

[3] While anecdotal evidence suggests that retail clients display particularly strong green preferences, our dataset does not allow us to check whether investment funds and pension funds in our sample are predominantly retail-owned. This is due to the fact that the bond ownership data in the SHS are only provided at the aggregate investor group level, not at a more granular institutional portfolio level. For the same reason, we also cannot see to what extent banks or investment funds can pass the additional costs from investing in green assets on to their customers, i.e. how the greenium is passed through.

frictions comes with a series of further caveats that we outline in more detail in Section 5. There we formulate some basic requirements that a mechanism needs to satisfy in order to explain our results and weigh some candidate mechanisms against them.

Our empirical analysis brings about several additional noteworthy results along the way. For example, we find substantial cross-sectional and time variation in the greenium, which can possibly explain conflicting results in the literature to some extent. A statistically significant greenium emerges only in the second half of our sample period. The average greenium of minus 3 basis points is estimated with a standard error of 1.6 basis points. Moreover, the greenium is more pronounced for old bonds, large bonds, and bonds from non-financial issuers. The greenium is insignificant for bonds that are not aligned with CBI standards and thus subject to moral hazard, confirming earlier papers which argue that the credibility of the issuer's commitment is a key determinant of the yield. Furthermore, descriptive statistics of the ownership structure of green bonds indicate that investment funds, banks and insurance companies are the key investor groups when it comes to the overall level of ownership. However, as sketched above, important differences emerge when comparing green bond ownership shares to the ownership shares in the conventional counterparts. Finally, we complement the empirical analysis with a series of robustness exercises, showing that our results are robust to choices concerning our matching approach, the composition of our sample as well as different ways of clustering the standard errors.

Taken together, our findings contribute to the ongoing debate about the peculiarities of the market for sustainable investments. We exploit the rich cross-sectional and time variation in both the greenium and the securities holdings data in order to better understand who pays for financing the green transition and why. More broadly, we view our results as evidence that supply and demand, and eventually prices and quantities, on the market for sustainable assets are driven by more than just the non-pecuniary green preferences of investors. Overall, this reasoning constitutes an interesting avenue for future research.

2 Literature

Our paper contributes to several strands of literature. First, there has been an active debate about the existence and the magnitude of the greenium on the green bond market. A substantial part of this debate is due to methodological differences, but also due to differences in sample periods and the underlying object of analysis (e.g., specific bond types). While Baker, Bergstresser, Serafeim, and Wurgler (2022) document a greenium of about six basis points in the primary market for U.S. municipal and corporate green bonds, Larcker and Watts (2020) find no significant greenium in the primary market for U.S. municipal bonds and conclude that the greenium may be difficult to detect in situations where investors with green preferences are not the marginal ones. Karpf and Mandel (2017) even find a green bond discount in the secondary market for U.S. municipal green bonds. In a broad global sample that ends in 2017, Zerbib (2019) finds a greenium of about two basis points in the secondary market. Using international data, Tang and Zhang (2020) again find no significant greenium for green corporate bonds in the primary market. Dorfleitner, Utz, and Zhang (2021) document a greenium of up to five basis points in the secondary market based on a broad international dataset. Kapraun, Latino, Scheins, and Schlag (2021) argue that the most important determinant for the greenium is green-credibility of the issuer and of the entity which issues the green label. In a similar vein, Fatica, Panzica, and Rancan (2021) argue that a greenium exists for supranational and corporate issuers, but not for green bonds issued by financial institutions. Zhang (2022) finds that the greenium also depends on the stringency of climate policy across countries. D'Amico, Klausmann, and Pancost (2023) fit a term structure model to specific German twin bonds and provide robust evidence for a greenium with this alternative estimation approach. Our paper contributes to this literature in multiple ways. Both the cross-sectional and time-series variation in the greenium that we document may explain why earlier papers in the literature disagree in their assessment whether a greenium exists or not. We confirm that green bonds from financial issuers have a much smaller greenium than those from non-financial corporates or sovereigns. We also confirm that only CBI-aligned/certified (i.e., credible) green bonds carry a greenium, whereas self-labeled green bonds do not.

A related strand of literature studies the effects of green bond issuance on firms' equity prices. Most prominently, Flammer (2021) and Tang and Zhang (2020) document that the issuance of green bonds may be viewed as a positive signal by investors: equity prices increase in response to it. This is confirmed by surveys that try to elicit the motives of different players in the market for sustainable assets. Krueger, Sautner, and Starks (2020) provide evidence from a survey among institutional investors and argue that green preferences indeed exist, for instance for a subset of the clients of these institutional investors. In a similar survey, Sangeorgi and Schopohl (2021) document a strong (unmet) demand of asset managers for green bonds from non-financial corporates. They argue that unclear and poor reporting is the largest obstacle. Using a survey among issuers, Sangeorgi and Schopohl (2023) document that reputational benefits and the signaling power of green bonds are the most important reasons to issue green bonds, in line with the green preference narrative that our result seem to confirm.

Our paper also contributes to the literature on the ownership structure of green bonds. Closest to ours is the paper of Pietsch and Salakhova (2022) who also rely on the SHS data and find that, among others, the demand from retail investors is a significant driver of the greenium. Using the same database, Boermans (2023) documents that mutual funds and pension funds act as preferred habitat investors in the green bond market, in the sense that their demand for green bonds is very price-inelastic, which is generally in line with the results in our paper. Baker et al. (2022) document that green bonds have more concentrated ownership and green bonds are more often held by green investors, as measured through hand-collecting socially responsible investors by the names of the respective investment funds. Bremus, Schütze, and Zaklan (2021) study the effects of central bank asset purchases, more precisely the ECB purchase programs on yields of green bonds. They find that the yields of green bonds that are eligible for these purchase programs decrease in response to the announcements of the programs. Yang (2021) studies green investor clientele effects and documents that green bonds are generally more resilient to liquidity shocks than conventional bonds because investors with green preferences alleviate search frictions in OTC markets. Brøgger and Kronies (2021) separate institutional investors into groups, taking into account whether they are constrained or unconstrained in their investment mandates. They document that the prices of high-ESG stocks go up, leading to lower expected returns going forward, when ESG-motivated investors buy these stocks. In a similar vein, Starks, Venkat, and Zhu (2019) document that long-horizon (i.e., low turnover) investors tilt their portfolios towards stocks of high-ESG firms.

Lastly, our paper is linked to the fast-growing broader literature on sustainable investing. Giving a comprehensive overview of this enormous field goes beyond the scope of this paper. Theoretical and empirical contributions have been made, among others, by Pedersen et al. (2021), Pastor et al. (2021), Oehmke and Opp (2020), Pastor et al. (2022), Bolton and Kacperczyk (2021), Zerbib (2022), Ilhan et al. (2021), and Boermans, Bun, and van der Straten (2024). Our paper also has connections to a recently growing literature on potential real effects from sustainable investing. Most likely, such real effects are shaped by differences in the ownership structure of green assets. For instance, De Haas and Popov (2023) argue that carbon emissions are generally lower in economies that are more equity-financed. Azar, Duro, Kadach, and Ormazabal (2021) observe a strong negative association between equity ownership by the Big Three (BlackRock, Vanguard and State Street Global Advisors) and firms' subsequent carbon emissions. Hartzmark and Shue (2023) study the impact elasticity of green and brown firms, i.e. the change in firms' carbon footprint due to changes in their cost of capital.

3 Data and descriptive statistics

In what follows, we describe our main datasets and report the key descriptives that were highlighted in the introduction.

3.1 Data

Our dataset combines information from different sources. First, we obtain a historical list of green bonds from issuers domiciled in developed economies (incl. supranationals and sovereigns) from Refinitiv-Eikon. Based on this list, we also extract data on all conventional bonds of the same issuers and their characteristics from the Eurosystem's Centralized Securities Database (CSDB). Following previous work, our focus is on plain vanilla bonds, that is we neither include any exotic debt instruments with embedded options nor bonds that are backed by other securities (such as covered bonds). As we explain in more detail below, we then match each green bond with the most similar conventional bond of the same issuer. For these matched bond pairs, we obtain data on monthly yields from Eikon. Lastly, we draw upon the Eurosystem's Securities Holdings Statistics (SHS), which allows us to track the dynamic ownership structure of matched bond pairs in our sample. More specifically, the SHS contains the quarterly securities holdings by investor group as reported by euro area custodians. The sectoral classification is based on the European System of Accounts (ESA) 2010 and we focus on the main investor groups of our sample bonds, namely *monetary financial institutions* (MFIs; S 122) *investment funds* (IFs; S 124), *insurance companies* (ICs; S 128), *pension funds* (PFs; S 129), the *Eurosystem* (EuSys; S 121), and *private households* (HHs; S 14 and S 15). We aggregate smaller sectors (such as non-financial companies) and unallocated investors in a category labeled as *others*. Since bonds may be held at custodians outside the euro area, we also define a residual sector labeled *foreign*. This ensures that, for any given bond, the total holdings across all investor groups will correspond to the bond's total amount outstanding. All bond yields are winsorized at the 1st/99th percentile.

To make sure that the investor groups of interest are economically important for our sample bonds, we focus on bond pairs where euro area investors, on average, hold at least 50% of a respective green bond, but we relax this assumption in robustness checks. Our final sample covers 161 green-conventional bond pairs with 3,142 observations (bond pair-months) over the period January 2015 until March 2022. Reflecting the massive growth of the green bond market in recent years, the bulk of observations are in the second half of the sample period, and we have only around 150 observations prior to 2018.

3.2 Matching

Following standard practice, we apply a matching approach that pairs each green bond with its most similar conventional "twin". We only take into account conventional bonds from the same issuer that have the same (nominal) currency, the same bond type, the same coupon type and the same seniority. Arguably, these bonds display the exact same credit risk from an investor's perspective. Moreover, we allow both the issue date and the maturity date of the paired bonds to differ by at most one year, and we require the green bond's total amount outstanding to be within 50% and 200% of the amount outstanding of the conventional bond. In case there are multiple conventional bonds that match all of the above requirements for a given green bond, we pick the conventional bond with the smallest normalized (Euclidean) distance in terms of issue date, maturity date and total amount outstanding.

These strict matching criteria reduce our sample size substantially, i.e., many green bonds do not have a conventional twin according to this specification. We end up with a total number of 436 matched bond pairs for which monthly secondary market yields are available from Refinitiv-Eikon. After applying the SHS coverage filter, this number drops further to 161 bond pairs. In robustness checks, we therefore relax these relatively strict matching criteria and find that our key results are, by and large, robust. (Internet Appendix B reports our main decomposition results when loosening our bond-level matching criteria.) Table 1 reports basic summary statistics on our matched sample.

	(1) Green	(2) Conventional
Bond size (in EUR million)	457.317	545.181
Original maturity (in years)	8.861	8.907
Yield (in percent)	0.533	0.561
Observations	3,142	
Bond-pairs	161	

Table 1: Summary statistics on matched bond-pairs.

3.3 Details on the sample composition

Arguably, our final dataset could be subject to a sample and selection bias, given the combination of several data filters and a relatively restrictive matching procedure. We therefore briefly discuss further descriptive statistics that alleviate such concerns.

Our dataset is based on a matched sample of green and conventional bonds of the same issuer. As such, our sample is confined to issuers with at least one green and at least one conventional bond outstanding (which need to satisfy further matching criteria). While our matching approach has important advantages (in particular, all issuer characteristics are identical within each bond pair by construction), our sample could be biased in that green bond issuers may differ from non-green bond issuers in a systematic way. In what follows, we therefore provide a comparison of a broader sample of green bond issuers with issuers of only conventional bonds.

For this purpose, we inspect one snapshot of the Eurosystem's Centralised Securities Database (CSDB) for December 2021. Applying broadly comparable filters as in the main sample construction (e.g., focusing only on issuers in developed markets and standard non-money market debt instruments), we identify a total number of 14,236 issuers with only conventional bonds outstanding and 740 issuers with at least one green bond outstanding. For the sake of brevity, we refer to these as *conventional issuers* and *green issuers* in the following. While much smaller in number, green issuers have a total bond volume of 26.4 trillion EUR outstanding, compared with 33.3 trillion EUR for conventional issuers. Still, despite the strong growth of the green bond market over recent years, the green bonds themselves make up only around 3% of the total bond volume of the green bond issuers. For the sake of representativeness, this analysis is based only on the CSDB data and does not incorporate any SHS ownership filters. In additional analyses, however, we find that the general patterns documented here continue to hold in a smaller subsample which applies the exact same bond-level SHS filters as in our main sample (i.e., excluding all bonds where euro area investors hold less than 50%).

Table A.2 in the Internet Appendix shows the two groups' compositions in terms of economic sectors (top) and geographic location (bottom). Both financial and non-financial corporates tend to be equally prevalent among green and conventional issuers, but sovereigns are somewhat more prevalent among green issuers.[4] Moreover, green issuers are more likely to be domiciled in the EU, which arguably hosts the biggest market for ESG products worldwide. In contrast, US issuers are under-represented in the green issuer sample.

Tables A.3 and A.4 in the Internet Appendix show regression results for different bond characteristics. To ensure a reasonable comparison that takes the differences from Table A.2 into account, we include country×sector fixed effects in Table A.3, i.e., we compare issuers within the same country and within the same economic sector. Similarly, to capture the differences in issuer ratings, the analyses in Table A.4 include country×sector×rating fixed effects (hence, unrated issuers are dropped from the analysis).

[4]The numbers in Table A.2 in the Internet Appendix are displayed such that the rows sum up to 100% for each split. Note that we classify all sovereign issuers (ESA code starting with S 13) within the same country as the same entity. In other words, we do not treat federal governments, states, municipalities, etc. as separate entities. Otherwise the number of sovereign issuers in our sample would be implausibly large ($\tilde{1}$5k).

We find that the rating of green bond issuers is on average 1.4 notches better than the rating of conventional issuers. Similarly, we find that green issuers are more likely to have an investment grade rating and are less likely to be unrated. Green issuers tend to issue a larger fraction of their bonds in EUR and less in USD, which is broadly in line with the differences in issuer domiciles. Green bond issuers have a larger bond volume outstanding (around 6.1 billion EUR) and their bonds tend to have a shorter maturity (around 1.7 years). In terms of average bond size, there is no discernible difference.

While these analyses indicate that a potential selection bias cannot be completely ruled out, if anything, it rather strengthens our key results. First, the better credit ratings and the shorter maturities imply that we rather underestimate the average greenium when we restrict our sample to green issuers only. Second, there is no selection bias regarding financial versus non-financial issuers. Still, green bond issuers overall tend to display certain differences from conventional issuers, which need to be kept in mind when interpreting our main results.

Finally, as an additional check concerning our data filters, Table A.1 in the Internet Appendix compares the sample composition before and after the application of the SHS coverage filter. In line with the snapshot analyzed above, our final sample consists predominantly of EUR-denominated bonds from Euro area issuers. The sample is also dominated by financial issuers with investment-grade rating (often above BBB). Note that there are several instances in our sample, where a non-financial company owns a financial subsidiary which acts as the official issuer of the firm's corporate bonds. We therefore manually correct the sector classification from the CSDB. For the sake of completeness, we also run our empirical tests with the non-corrected sample and find that our results are qualitatively very similar.

3.4 Descriptive statistics

We introduce some notation that will be useful also in subsequent sections. For ease of notation, we drop time subscripts t in the following. Unless stated otherwise, the empirical analyses are based on the pooled sample.

We denote the yield-to-maturity for a given pair of green and conventional bonds i by y_i^G and y_i^C, respectively. The EUR amounts that sector x has invested in green or conventional bond i are denoted by $n_{i,x}^G$ and $n_{i,x}^C$, respectively. From these amounts, we can compute the share of a bond i that is held by sector x:

$$s_{i,x}^C = \frac{n_{i,x}^C}{\sum_y n_{i,y}^C}, \quad s_{i,x}^G = \frac{n_{i,x}^G}{\sum_y n_{i,y}^G} \tag{1}$$

Note that the denominator is equal to the total amount outstanding of bond i in our sample, such that the shares sum up to one:

$$\sum_x s_{i,x}^G = 1 \text{ and } \sum_x s_{i,x}^C = 1 \quad \text{for all } i. \tag{2}$$

We denote the greenium of bond pair i as $g_i = y_i^G - y_i^C$. Importantly, throughout the paper, differences are always taken as "green minus conventional" (not "conventional minus green").

Wherever confusion is unlikely, we denote both equal-weighted and value-weighted averages by a bar, e.g.

$$\bar{g} = \frac{1}{\#i} \sum_i g_i \quad \text{or} \quad \bar{g} = \sum_i \frac{n_i}{\sum_j n_j} g_i. \tag{3}$$

where $\#i$ denotes the total number of green bonds in our sample and n_i is the total amount outstanding of green bond i.

Similarly, we denote the average share of green bonds that is held by sector x by

$$\bar{s}_x^G = \frac{1}{\#i} \sum_i s_{i,x}^G \quad \text{or} \quad \bar{s}_x^G = \sum_i \frac{n_i}{\sum_j n_j} s_{i,x}^G, \tag{4}$$

respectively.

3.4.1 Greenium

We first provide some descriptive statistics of our matched sample. Of particular interest is the (average) greenium. Figure 1 depicts the development of the unconditional, equal-weighted average greenium $\bar{g}$ over time. The vertical bars denote the 95% confidence intervals (using robust standard errors) at each point in time. Similarly, Figure 2 presents the average value-weighted greenium.

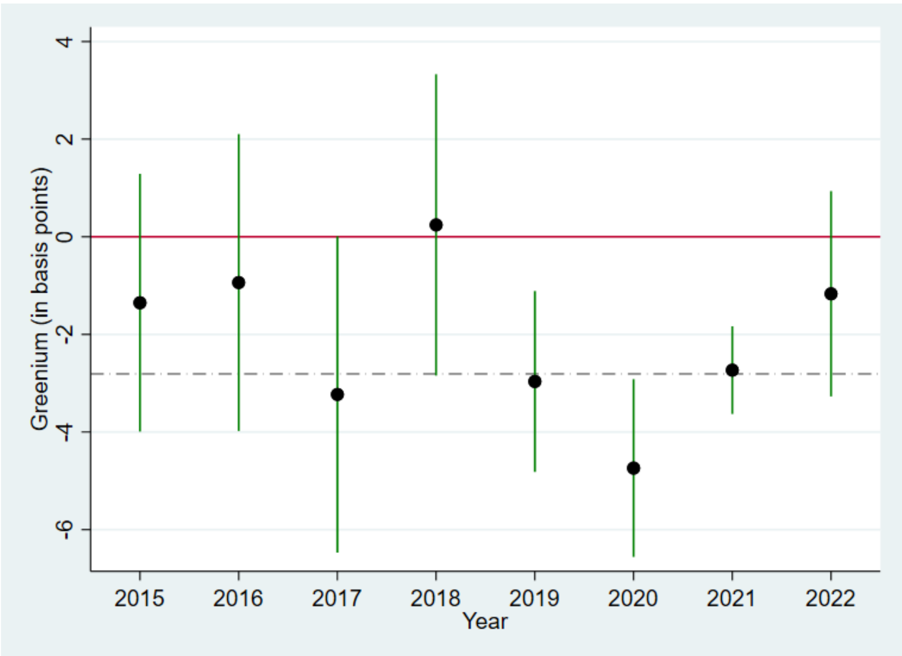

Figure 1: Greenium Over Time (Equal-Weighted).

The figure depicts the yield differential between green and matched conventional bonds over time in our sample. Black dots indicate the equal-weighted average greenium $\bar{g}_t$. The vertical lines indicate the 95% confidence intervals of the greenium in each cross-section (robust standard errors). The dashed gray line shows the full-sample equal-weighted average greenium of -2.8 basis points.

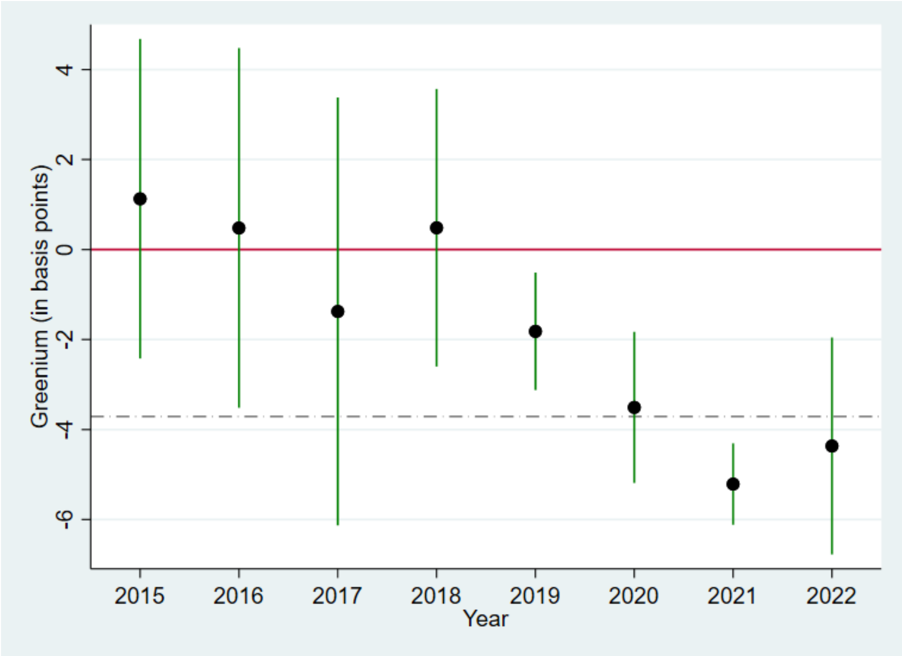

Figure 2: Greenium Over Time (Value-Weighted).

The figure depicts the yield differential between green and matched conventional bonds over time in our sample. Black dots indicate the value-weighted average greenium $\bar{g}_t$ in our sample. The vertical lines indicate the 95% confidence intervals of the greenium in each cross-section (robust standard errors). The dashed gray line shows the full-sample value-weighted average greenium of -3.7 basis points.

In line with previous work (e.g., Baker et al. (2022), Zerbib (2019)), we find evidence in favour of the existence of a greenium: on average, green bonds trade at lower yields (and higher prices) than their conventional counterparts. Moreover, the greenium has become more pronounced over time, which is in line with empirical evidence for other asset classes like equity, derivatives or corporate loans (see, e.g., Bolton and Kacperczyk (2021), Ilhan et al. (2021), Delis, de Greiff, and Ongena (2021)). This trend potentially reflects the increasing relevance of environmental concerns and climate-related risks for investors in general, which has been widely documented. As such, our findings may explain why previous literature came to varying conclusions as to whether a greenium exists or not. For instance, in an early paper, Larcker and Watts (2020) argue that green US municipal bonds do not carry a greenium, whereas D'Amico et al. (2023) provide robust evidence for a greenium using German twin bonds after 2020. In terms of economic magnitude, the greenium fluctuates around 3-4 basis points in the most recent years, where the bulk of our sample is concentrated. Both the negative time trend since 2018 and the magnitude of the greenium are more pronounced in the value-weighted analysis. Hence, larger (and potentially more liquid) bonds display a larger greenium of around 6 basis points in the most recent sample. We will return to this point below.

3.4.2 Bond ownership structure

Figures 3 and 4 illustrate the average ownership structure of the green and conventional bonds in our sample. More precisely, the top panel of Figure 3 shows the average shares $\bar{s}_x^G$ and their respective counterparts $\bar{s}_x^C$, i.e., it quantifies which fraction of the green or conventional bonds in our sample is held by which investor group. Figure 4 shows the respective value-weighted averages. The top left panels address our first research question: banks, investment funds, and insurances are the most important investors in green bonds, making up close to 50% of the ownership base.

The results are more nuanced in the case of private households: their direct investments in bonds (both green and conventional) are indicated by the green bars in Figures 3 and 4. The relative shares are quite sizeable when all bonds are equally weighted, but negligible when value-weighted. This indicates that private households' direct investment is largely focused on relatively small bonds.

The lower panels in Figures 3 and 4 depict the average ownership differentials by investor group, based on the average holdings in the top panels. Table 2 reports results from a statistical analysis based on the granular ownership shares. The average differentials indicate how much each sector over-/underweights green bonds relative to their conventional counterparts in general. Arguably, these differentials can serve as rough proxies for a given sector's general green preference. This interpretation is in line with theoretical models like the ones by Pedersen et al. (2021) or Pastor et al. (2021). Assuming that the "greenness" of all bonds in our sample is of a similar magnitude, a group of investors that is guided solely by green preferences should be expected to invest the same fraction of wealth in each available green asset, irrespective of any cross-sectional differences in the greenium across these assets. When viewed from such a perspective, our results indicate that investment funds have the largest green preference, followed by private

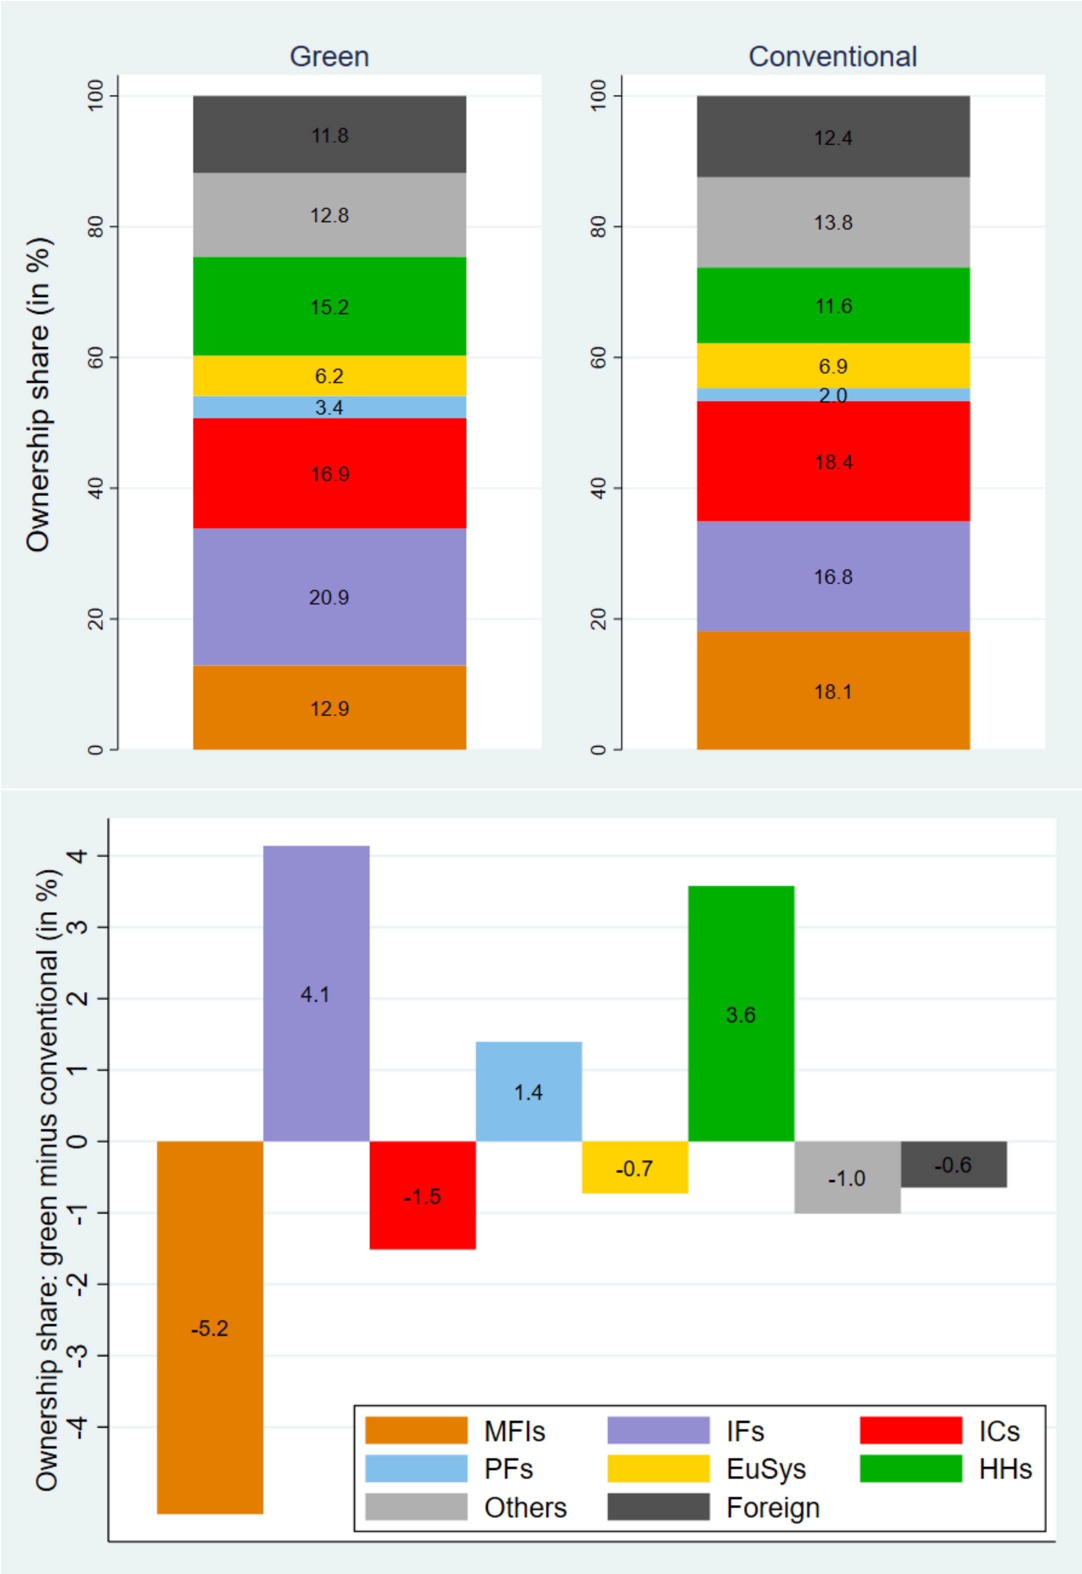

Figure 3: Ownership Structure (Equal-weighted).

The top graph displays the equal-weighted average shares of green and conventional bonds that are held by the different investor groups in our sample. The bottom graph depicts the difference between the respective numbers in the two upper graphs. The labels refer to monetary financial institutions (MFIs), investment funds (IFs), insurance companies (ICs), pension funds (PFs), the Eurosystem (EuSys), and private households (HHs).

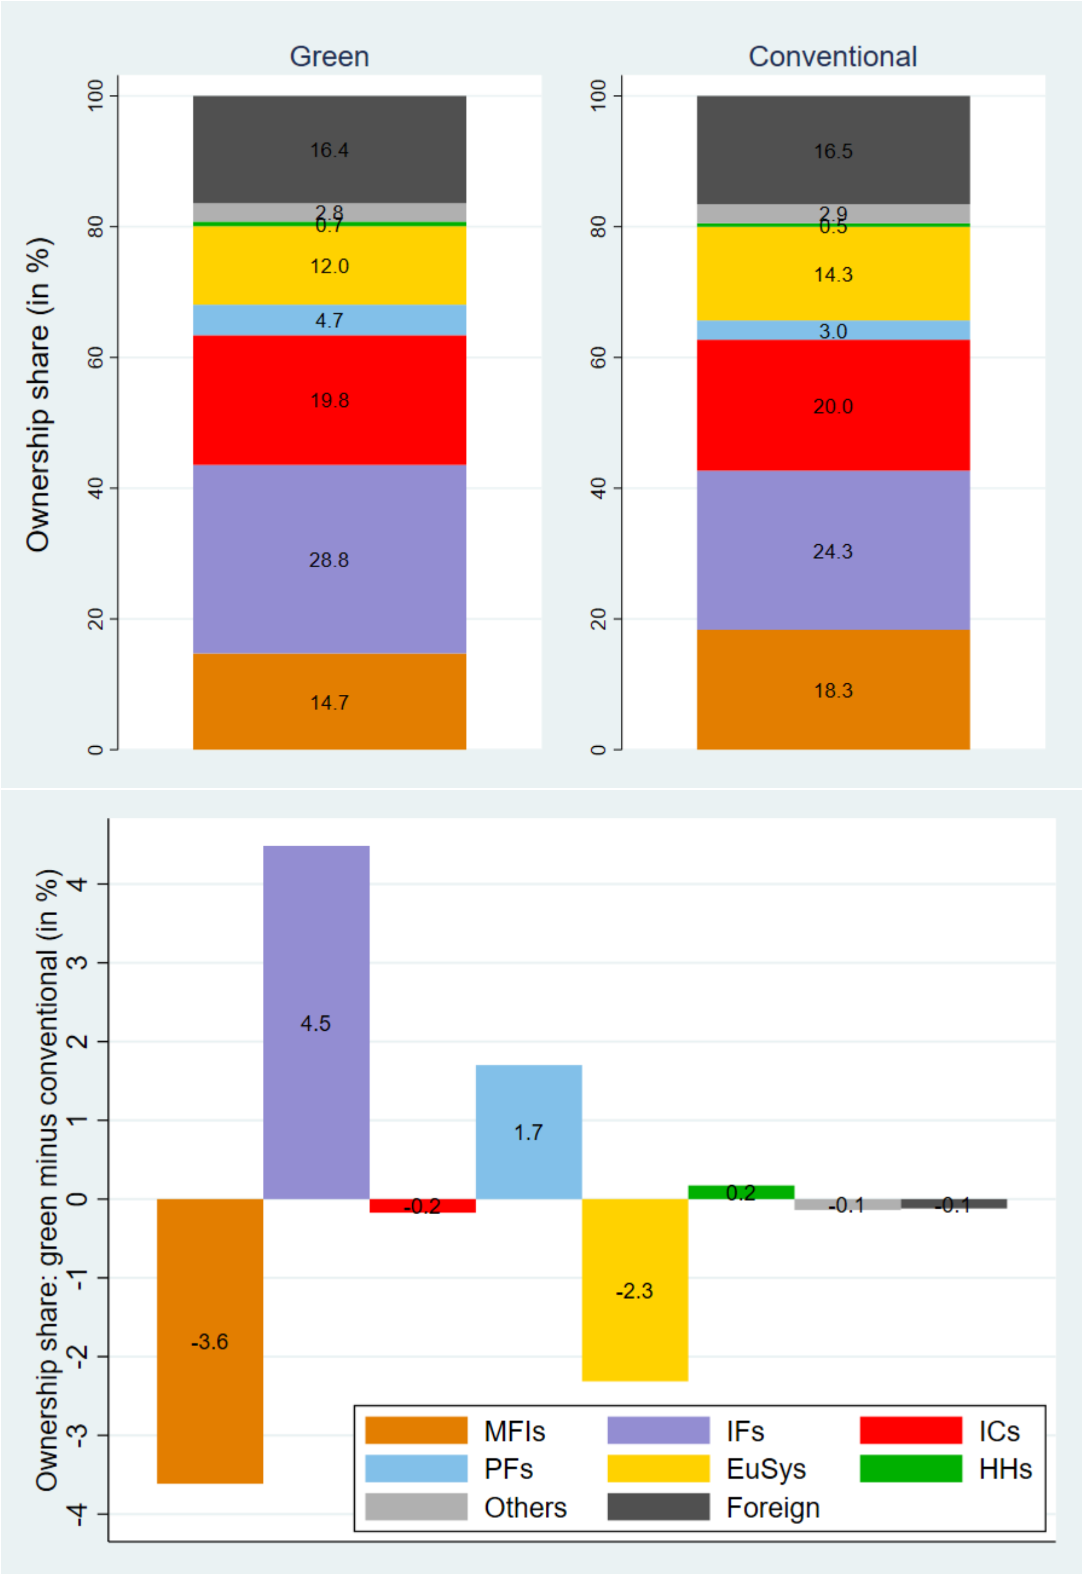

Figure 4: Ownership Structure (Value-weighted).

The top graph displays the value-weighted average shares of green and conventional bonds that are held by the different investor groups in our sample. The bottom graph depicts the difference between the respective numbers in the two upper graphs. The labels refer to monetary financial institutions (MFIs), investment funds (IFs), insurance companies (ICs), pension funds (PFs), the Eurosystem (EuSys), and private households (HHs).

households and pension funds. This is in line with anecdotal and survey evidence (see, e.g., Krueger et al. (2020) or Sangeorgi and Schopohl (2021)). Investment and pension funds manage portfolios on behalf of their clients, such that the green preferences of these clients arguably carry over to the portfolios of the funds in which they invest.

Ownership delta $s^G_{i,x}$ - $s^C_{i,x}$	MFIs	IFs	ICs	PFs	EuSys	HHs	Others	Foreign
Panel A								
Equal-weighted	-5.219***	4.140***	-1.512*	1.397***	-0.728	3.578*	-1.011	-0.646
	(-3.881)	(3.982)	(-1.707)	(3.077)	(-1.203)	(1.987)	(-0.521)	(-0.587)
Panel B								
Value-weighted	-3.262**	4.530***	-2.495*	1.633***	-1.003	0.081	-0.441	0.956
	(-2.436)	(3.207)	(-1.993)	(3.317)	(-1.096)	(0.736)	(-0.925)	(0.908)
Observations	1,076							

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Ownership delta - Green minus conventional.
The table reports the average bond-pair level ownership differential between green and conventional bonds, separately for each investor group. Panel A reports equal-weighted estimates and Panel B value-weighted ones based on the corresponding green bond's total amount outstanding. Here we use quarterly data given that the ownership shares do not vary within each quarter. Standard errors clustered by bond and date (t-statistics in parentheses). The labels refer to monetary financial institutions (MFIs), investment funds (IFs), insurance companies (ICs), pension funds (PFs), the Eurosystem (EuSys), and private households (HHs).

The ownership delta is significantly negative for banks. Hence, they generally underweight green bonds as compared to their conventional counterparts. Importantly, however, we refrain from overinterpreting this as evidence for a "brown preference" of banks. Instead, it is more likely that banks use other vehicles for "greening" their asset portfolios. One has to keep in mind that banks and insurance companies often invest in a broader range of asset classes as compared to mutual funds and pension funds.[5] For insurers, we see a smaller negative ownership delta when all bonds are equally weighted, but the sign switches in the value-weighted case. This implies that insurers tend to underweight smaller green bonds relative to their conventional counterparts. For the Eurosystem, we see a slight underweighting in both cases but these deviations are not statistically significant. We view this pattern as being in line with the principle of market neutrality which guided the Eurosystem's asset purchase programs over the course of our sample period. Foreign investors consistently display a negative ownership delta as well, i.e., are more strongly invested in conventional bonds.

Figure A.1 in Internet Appendix A depicts annual time series of the share differentials. After 2018, as the sample grows larger, the graphs indicate little variation over time in general, with one notable exception, namely the late arrival of private households. For them, the average ownership delta shows a substantial increase after 2020 (at least in the equal-weighted case), which is in line with the findings of Pietsch and Salakhova (2022).

[5]In a recent paper, Darmouni and Papoutsi (2022) merge confidential Euro Area credit register and securities holdings data and study the joint cross-section of bonds and bank loans in great detail. The authors do not analyze sustainable investments specifically, though.

For the sake of completeness, Figures A.2-A.6 in Internet Appendix A also contain ownership decompositions for a set of sample splits. The sample split between financial and non-financial issuers deserves some attention. Consistent with Bekaert and Breckenfelder (2019), banks hold a much larger share in bonds issued by financials as opposed to bonds issued by non-financials. Importantly, we observe this pattern among both green and conventional bonds in our sample. We draw the tentative conclusion that banks' motives for investing in green financial bonds are likely similar to their motives for investing in conventional financial bonds. Still, the tilt towards financial bonds is more pronounced among conventional bonds than among green bonds. We also find some striking patterns with regard to the size of the bonds. Banks have a general tilt towards small bonds and this tilt is much more pronounced among conventional bonds than among green bonds. As a result, banks have a very large share in small, conventional bonds. Banks also have a tilt towards short maturity bonds, and this tilt is equally pronounced among green and conventional bonds.

4 Who pays the greenium (and why)?

In the previous section, we have presented descriptive statistics for green bond yields and green bond holdings in isolation. In the following, we complete the picture with a joint analysis of yields and holdings and take a closer look at the relationship between the greenium and a bond's ownership structure. For this purpose, we develop our novel methodology to decompose the greenium across both the ownership structure and the cross-section of bonds.

4.1 Who pays the greenium on average?

To this end, we first decompose the total greenium into investor group-specific contributions. Drawing on the notation introduced above, we can rewrite $\bar{g}$ as

$$\bar{g} = \frac{1}{\#i} \sum_i \underbrace{\left(\sum_x s_{i,x}^G \right)}_{=1} g_i = \sum_x g_x \tag{5}$$

where g_x is the greenium that sector x earns/pays for its green bond portfolio

$$g_x = \frac{1}{\#i} \sum_i s_{i,x}^G g_i. \tag{6}$$

Similarly, we also define a value-weighted version of g_x:

$$g_x = \sum_i \frac{n_i}{\sum_j n_j} s_{i,x}^G g_i. \tag{7}$$

The first row in Table 3 reports g_x across sectors, i.e. it quantifies which sector pays how much of the overall greenium on average. By construction, the numbers in the first row add up to the total greenium given in the first column. They allow us to answer the main research question posed in the title of our paper: Out of the total greenium of roughly minus three basis points, insurances bear more than half, investment funds bear one third and banks bear another quarter. Thus, these three groups of institutional investors (or their clients) pay the bulk of the greenium, answering our first research question.[6] Panel B confirms this finding for a value-weighted decomposition. In this case, both banks and investment funds bear an even larger share of the overall greenium. Given the pronounced cross-sectional heterogeneity, the point estimates for insurers are however not statistically significant in both cases.

As we will discuss in Section 4.4, the finding that banks, investment funds and insurers pay most of the greenium is robust to various changes concerning how we filter the data to put together our final sample. There we also present and discuss results for different standard error specifications. (In the main part, the reported t-statistics are based on standard errors clustered by bond and date.)

Panel A: Equal-weighted	Total	MFIs	IFs	ICs	PFs	EuSys	HHs	Others	Foreign
Greenium (g_x)	-2.811*	-0.529**	-0.918**	-1.570	-0.090	-0.030	0.370*	0.380	-0.423**
	(-1.723)	(-2.109)	(-2.388)	(-1.329)	(-1.102)	(-0.353)	(1.869)	(1.528)	(-2.443)
I: Benchmark Greenium		0.185	-0.899***	-1.812	-0.067	-0.069	0.038	0.185	-0.372*
		(0.487)	(-2.645)	(-1.507)	(-1.661)	(-0.550)	(0.232)	(0.966)	(-1.709)
II: General green preference		0.139***	-0.120***	0.043*	-0.039***	0.018	-0.100*	0.033	0.025
		(3.759)	(-4.194)	(1.794)	(-3.122)	(1.053)	(-1.927)	(0.579)	(0.856)
III: Bond-specific deviations		-0.854***	0.102	0.198	0.017	0.021	0.432**	0.162	-0.077
		(-2.706)	(0.805)	(1.493)	(0.244)	(0.287)	(2.411)	(0.628)	(-0.711)
Observations	3,133								

Panel B: Value-weighted	Total	MFIs	IFs	ICs	PFs	EuSys	HHs	Others	Foreign
Greenium (g_x)	-3.714**	-0.929**	-1.355**	-0.592	-0.185*	-0.104	-0.025*	-0.110*	-0.412**
	(-2.278)	(-2.219)	(-2.455)	(-1.218)	(-1.731)	(-0.847)	(-1.727)	(-1.780)	(-2.206)
I: Benchmark Greenium		-0.525	-1.367***	-1.013*	-0.120	-0.151	-0.018	-0.015	-0.505*
		(-1.413)	(-2.726)	(-1.687)	(-1.539)	(-0.937)	(-1.625)	(-0.258)	(-1.953)
II: General green preference		0.062	-0.124***	0.059	-0.059***	0.024	-0.003	0.011	0.031
		(1.644)	(-3.038)	(1.617)	(-3.144)	(0.950)	(-0.890)	(0.816)	(1.071)
III: Bond-specific deviations		-0.467**	0.136	0.363	-0.006	0.022	-0.004	-0.106**	0.062
		(-2.081)	(0.746)	(1.554)	(-0.077)	(0.253)	(-0.389)	(-2.032)	(0.463)
Observations	3,133								

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Greenium decomposition.

The table reports our decomposition of the total greenium. Panel A shows results for the equal-weighted case, Panel B for the value-weighted case. The first row in each panel reports the sectoral greenium g_x, as defined in Section 4.1. The numbers in this row add up to the total greenium in the first column. The subsequent rows in each panel report the three components defined in Section 4.2. The numbers in each column add up to the greenium g_x in first row. Standard errors are clustered by bond and date (t-statistics in parentheses). The labels refer to monetary financial institutions (MFIs), investment funds (IFs), insurance companies (ICs), pension funds (PFs), the Eurosystem (EuSys), and private households (HHs).

[6] Our dataset does not allow us to "look through" the different investor groups in our sample and find out the ultimate investors behind each group, in order to check whether, e.g., investment funds and pension funds are predominantly retail-owned. The bond ownership data from the SHS are only provided at the aggregate investor group level, but not at a more granular institutional portfolio level.

4.2 Further decomposing the sector-specific greenium

Having split the total greenium into sector-specific shares above, we now seek to understand the underlying reasons why the main investor groups pay the greenium. To this end, we decompose the sector-specific greenium further. Specifically, we rewrite g_x as follows:

$$g_x = \underbrace{g_x^C}_{\text{I: Benchmark Greenium}} + \underbrace{\frac{1}{\#i} \cdot \bar{g} \cdot \sum_i \Delta s_{i,x}}_{\text{II: General green preference}} + \underbrace{\frac{1}{\#i} \cdot \sum_i \Delta s_{i,x}(g_i - \bar{g})}_{\text{III: Bond-specific deviations}}, \tag{8}$$

with the share differential $\Delta s_{i,x} := (s_{i,x}^G - s_{i,x}^C)$ and $g_x^C := \frac{1}{\#i} \sum_i s_{i,x}^C g_i$. (For the sake of brevity, we omit the straightforward, complementary definitions for value-weighted averages.) The three components have straightforward economic interpretations. The first component, g_x^C, is a benchmark greenium that would be paid/earned if the shares of each sector in each green bond were identical to the shares in the respective conventional bond. We view this benchmark as representing a simple null hypothesis, quantifying a sector's greenium if the sector would treat green and conventional bonds identically and thus displayed the same ownership shares. Under this simple null hypothesis, we would have $\Delta s_{i,x} = 0$, such that the other two components would disappear.

The second component, $\frac{1}{\#i} \cdot \bar{g} \cdot \sum_i \Delta s_{i,x}$, is the part of the greenium which is paid/earned because a sector over-/underweights green bonds relative to conventional bonds on average, across the entire green bond market. As the formula shows, it is the product of the average greenium $\bar{g}$ that we have discussed in Section 3.4.1 and the average of the ownership share differential, $\frac{1}{\#i} \cdot \sum_i \Delta s_{i,x}$, discussed in Section 3.4.2. In a broader sense, we may interpret this second component as the greenium that a sector x pays/earns because of its *general green preference*. However, one has to keep in mind that certain investor sectors mainly act on behalf of their clients, such that this general green preference would rather reflect their clients' green preferences. Moreover, there are many alternative ways how investors can make their portfolios greener besides holding green bonds, which turns this component into a noisy and potentially biased proxy of green preferences.

The third component, $\frac{1}{\#i} \cdot \sum_i \Delta s_{i,x}(g_i - \bar{g})$, reflects the part of the greenium which results from *bond-specific over-/underweighting* of green bonds whose greenium deviates from the average. Technically speaking, this term captures the covariance between $\Delta s_{i,x}$ and $(g_i - \bar{g})$. It is negative when sector x overweights particularly those green bonds which have a very high (i.e. very negative) greenium. Economically, this residual term subsumes the fraction of the greenium which is paid for reasons other than (I) replicating the equivalent conventional bond portfolio or (II) general green preferences.

The results from this decomposition are also reported in Table 3. Most importantly, we find that the total greenium cannot fully be explained by the benchmark greenium (I). While the benchmark greenium tends to be relatively close to the total greenium for both investment funds and pension funds, the total greenium is more pronounced (i.e., more negative) compared to the simple benchmark in both cases. The deviations between the total and the benchmark greenium are the largest for banks and insurances. In particular, banks would *earn* a greenium of 0.19 basis points (rather than *pay* 0.53 basis points) if their green bond portfolio had the same structure as their conventional bond portfolio. Interestingly, in line with the concept of market neutrality that guides the asset purchase programs of the Eurosystem throughout our sample period, we find that neither of the three components are significantly different from zero for the Eurosystem.

We therefore conclude that components II and III must play a role across the four main investor groups. Indeed, consistent with the descriptives in Figure 3, the second component is negative and significant for investment funds, pension funds and private households. We interpret this as evidence that these groups pay a substantial part of the greenium due to their strong general green preferences. This finding is in line with recently developed theories that try to explain the price differential between sustainable and non-sustainable assets through a preference channel. In contrast, the general preference term is positive and significant for banks and insurances.

The third component is insignificant for all major sectors with one big exception: for banks it is strongly negative at minus 0.85 basis points. In fact, the negative contribution from bond-specific over-/underweighting exceeds the positive general preference term, so that banks pay an average greenium of 0.53 basis points overall.[7] Besides, we find a strongly significant positive third component of 0.43 basis points for private households, but this term disappears in the value-weighted case. For insurances and investment funds we also find a positive, albeit statistically insignificant third component of 0.10 and 0.18 basis points, respectively. The positive estimates indicate that these three sectors tend to slightly underweight bonds with a large greenium.

4.3 Sample splits

The significant and strongly negative covariance between the greenium and the ownership delta for banks stands out from the decomposition results presented above. In order to gain further intuition concerning the drivers behind it, we conduct several sample splits. The first exercise is a median split based on bond size, as measured by total amount outstanding. Second, we perform a median split based on bond age, as measured by the time since issuance. This exercise potentially also captures liquidity differences between on-the-run (young) and off-the-run (old) bonds. Third, to capture potential market segmentation effects regarding bonds in different maturity buckets, we also conduct a sample split based on the remaining time until maturity.

Another sample split differentiates between financial and non-financial issuers. Green bonds issued by banks are somewhat special since banks do not execute the respective green projects themselves, but rather invest the proceeds from the green bonds into green loans (Flammer (2021)). In addition, banks may hold other banks' bonds for reasons unrelated to sustainability, e.g., when they provide underwriting or market making services. For example, Bekaert and Breckenfelder (2019) document that European banks hold 28% of all bank-issued securities. This is much higher than their share in the overall securities holdings, which amounts to 18%. The authors find this pattern to be particularly pronounced for long-term debt instruments.

[7] In additional unreported analyses we find that the MFI results are largely due to German banks. This could be due to the fact that Germany hosts one of the largest banking systems in the euro area (market share of 28% as of 2022-Q2). More granular datasets would be needed to dig deeper into institution-specific bond holdings.

The last sample split distinguishes self-labeled green bonds and CBI-aligned/-certified green bonds. Previous literature (see, e.g., Kapraun et al. (2021)) has documented that green bonds only carry a greenium if the commitment to use the proceeds sustainability is reasonably credible and, for instance, confirmed by the CBI.[8]

Table 4 reports the average covariance term in Eq. (8) for the different sample splits. The t-statistics for the differences in means refer to a two-sample t-test of the null hypothesis that the two respective subsamples have the same means, allowing the variances to differ between subsamples. The first column documents that, for banks, the covariance term is significantly more pronounced (i.e. more negative) for small bonds, young bonds, bonds with a longer remaining maturity, and bonds issued by financials. In contrast, we find only a slightly more pronounced value for CBI-aligned/certified bonds, but the difference is insignificant. The opposite pattern is observed for private households, who effectively profit from holding specifically those small or young or financial green bonds with a relatively small (or even positive) greenium. Altogether, these findings also confirm the value-weighted results in Panel B of Table 3. The corresponding estimate for bond-specific deviations for banks in the last row is about half as large as in Panel A (-0.47 bps versus -0.87 bps), suggesting that the bond-specific over-/underweighting effect is much more pronounced for small green bonds as opposed to large green bonds.

For the sake of completeness, Table 5 reports the average total (sector-specific) greenium g_x across the different sample splits. Remarkably, the sample split results for the covariance term are very different from the ones for the total greenium. In contrast to the bond-specific over-/underweighting component, the total greenium is significantly more pronounced for (1) larger bonds, (2) older bonds, (3) bonds issued by non-financials, and (5) CBI-aligned/-certified bonds. Hence, for our sample we confirm that only CBI-certified/-aligned green bonds carry a greenium, whereas self-labeled green bonds (for which the green commitment may not be credible) do not.

In fact, the results for the total greenium are rather mirrored both in the first ("benchmark greenium") and second component ("general green preferences"), when we again decompose the greenium in these subsamples. (Tables A.5 and A.6 in the Internet Appendix report the sample split results for components I and II.) It seems reasonable to assume that the attention of sustainability-oriented investors to green bonds, which supposedly drives component II of the greenium, focuses on large bonds issued by non-financial corporations, where the presumed impact of sustainable investing is more directly observable.

[8]The Climate Bond Initiative (CBI) is an international organization that monitors the green bond market in order to help mobilise capital for the green transition. It screens the entire universe of green bonds. If a bond satisfies certain criteria (most importantly that the funded activities are in line with the CBI taxonomy and that the proceeds of the green bonds are fully used for these intended purposes), then it is added to the CBI Green Bond Database and is labeled as "CBI-aligned". On top of that, an issuer can also actively file for certification by the CBI in collaboration with external auditors. After passing this process, the bond is labeled as "CBI-certified". All other green bonds are considered as "self-labeled" by the CBI.

III: Bond-specific deviations Sample splits	MFIs	IFs	ICs	PFs	EuSys	HHs	Others	Foreign
Size (total amount outstanding)								
Small	-0.970	0.042	0.115	0.033	0.040	0.628	0.347	-0.234
Large	-0.563	0.223	0.296	-0.023	-0.018	-0.001	-0.147	0.234
(Diff)	-0.41**	-0.18***	-0.18**	0.06**	0.06*	0.63***	0.49***	-0.47***
	(-2.26)	(-2.76)	(-2.52)	(2.39)	(1.78)	(7.25)	(3.39)	(-10.24)
Age (time since issuance)								
Young	-1.220	0.028	0.251	-0.003	-0.005	0.536	0.459	-0.046
Old	-0.463	0.195	0.117	0.070	0.048	0.197	-0.092	-0.072
(Diff)	-0.76***	-0.17***	0.13**	-0.07***	-0.05*	0.34***	0.55***	0.03
	(-3.50)	(-2.61)	(2.37)	(-2.74)	(-1.87)	(3.03)	(3.04)	(0.55)
Residual maturity								
Low	-0.657	0.137	0.319	0.041	0.046	0.094	0.138	-0.117
High	-0.862	0.139	0.200	0.028	0.023	0.385	0.169	-0.082
(Diff)	0.20*	0.00	0.12***	0.01*	0.02*	-0.29***	-0.03	-0.03*
	(1.76)	(1.13)	(4.29)	(1.90)	(1.84)	(-5.88)	(-0.26)	(-1.76)
Issuer industry								
Non-Financials	0.052	-0.044	-0.074	-0.036	0.078	-0.013	-0.021	0.058
Financials	-1.471	0.277	0.316	0.050	-0.011	0.819	0.215	-0.195
(Diff)	1.52***	-0.32***	-0.39***	-0.09***	0.09**	-0.83***	-0.24	0.25***
	(8.67)	(-5.69)	(-7.05)	(-3.36)	(2.23)	(-8.27)	(-1.55)	(5.25)
ESG Flag								
CBI Aligned/Certified	-0.862	0.139	0.200	0.028	0.023	0.385	0.169	-0.082
Self-labelled	-0.756	-0.077	0.207	-0.003	-0.020	0.216	0.467	-0.034
(Diff)	-0.11	0.22***	-0.01	0.03*	0.04**	0.17	-0.30*	-0.05
	(-0.53)	(4.03)	(-0.11)	(1.77)	(2.17)	(1.41)	(-1.65)	(-1.37)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Sample-splits for bond-specific deviations (equal-weighted).

This table reports Component III (bond-specific deviations) from our main decomposition (as shown in Panel A of Table 3) for several sample splits. The rows labeled "Diff" report the differences between the two sub-samples and the t-statistic from a two-sample t-test of the null hypothesis that the two respective subsamples have the same means, allowing the variances to differ. The labels refer to monetary financial institutions (MFIs), investment funds (IFs), insurance companies (ICs), pension funds (PFs), the Eurosystem (EuSys), and private households (HHs).

Greenium (g_x) Sample splits	Total	MFIs	IFs	ICs	PFs	EuSys	HHs	Others	Foreign
Size (total amount outstanding)									
Small	-1.867	-0.240	-0.529	-2.140	-0.014	0.017	0.598	0.682	-0.242
Large	-4.471	-1.038	-1.602	-0.569	-0.223	-0.112	-0.030	-0.153	-0.743
(Diff)	2.60***	0.80***	1.07***	-1.57***	0.21***	0.13***	0.63***	0.84***	0.50***
	(3.95)	(5.47)	(5.61)	(-4.30)	(6.82)	(3.01)	(6.19)	(5.37)	(6.11)
Age (time since issuance)									
Young	-0.843	-0.560	-0.790	-0.293	-0.050	-0.010	0.514	0.663	-0.317
Old	-4.836	-0.497	-1.050	-2.885	-0.131	-0.051	0.222	0.088	-0.533
(Diff)	3.99***	-0.06	0.26	2.59***	0.08**	0.04	0.29**	0.58***	0.22***
	(5.71)	(-0.42)	(1.40)	(5.79)	(2.55)	(1.06)	(2.27)	(2.93)	(2.78)
Residual maturity									
Low	-2.885	-0.779	-1.149	-0.321	-0.103	0.073	-0.052	0.080	-0.633
High	-3.394	-0.620	-1.006	-1.883	-0.102	0.000	0.444	0.273	-0.499
(Diff)	0.51	-0.16***	-0.14**	1.56***	0.00	0.07***	-0.50***	-0.19***	-0.13***
	(-0.21)	(-3.40)	(-2.51)	(5.66)	(-0.89)	(5.51)	(-6.58)	(-3.05)	(-5.53)
Issuer industry									
Non-Financials	-6.089	-0.131	-0.916	-4.394	-0.149	-0.235	0.064	-0.014	-0.314
Financials	-1.185	-0.727	-0.919	-0.170	-0.060	0.072	0.522	0.575	-0.478
(Diff)	-4.90***	0.60***	0.00	-4.22***	-0.09***	-0.31***	-0.46***	-0.59***	0.16**
	(-6.04)	(5.26)	(0.02)	(-6.46)	(-2.98)	(-6.16)	(-4.47)	(-3.97)	(2.30)
ESG Flag									
CBI Aligned/Certified	-3.394	-0.620	-1.006	-1.883	-0.102	0.000	0.444	0.273	-0.499
Self-labelled	0.250	-0.052	-0.456	0.073	-0.022	-0.189	-0.017	0.942	-0.028
(Diff)	-3.64***	-0.57***	-0.55***	-1.96***	-0.08***	0.19***	0.46***	-0.67**	-0.47***
	(-5.73)	(-5.96)	(-2.98)	(-6.42)	(-3.55)	(3.79)	(3.02)	(-2.16)	(-6.17)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Sector-specific greenium - Sample-splits (Equal-weighted).

This table reports the equal-weighted average greenium g_x (as shown in Panel A of Table 3) for several sample splits. The rows labeled "Diff" report the differences between the two sub-samples and the t-statistic from a two-sample t-test of the null hypothesis that the two respective subsamples have the same means, allowing the variances to differ. The labels refer to monetary financial institutions (MFIs), investment funds (IFs), insurance companies (ICs), pension funds (PFs), the Eurosystem (EuSys), and private households (HHs).

4.4 Robustness

We complement our analysis with a set of robustness checks. The detailed results are reported in Internet Appendix B.

Our final sample of green-conventional bond pairs is based on a series of data filters. In our baseline setup, we only include bond pairs where euro area investors hold at least 50% of a green bond's total amount outstanding. Internet Appendix B.1 documents that our results continue to hold when we lower this threshold to 10%.

Our empirical setup imposes a set of assumptions concerning the matching strategy. Internet Appendix B.2 addresses our matching criteria. In our baseline setup, we allow the issue date and the maturity date of green and conventional bonds to differ by at most 1 year. Allowing for 2 years difference, the data gets noisier, but our results are, by and large, robust. The same is true when we allow the amount outstanding of matched green and conventional bonds to differ by a factor of 4 instead of 2.

Finally, our statistical tests in Table 3 are based on t-statistics with standard errors clustered by bond and date. Table B.3 in Internet Appendix B.3 provides results for alternative clustering approaches (by time, bonds, and/or issuers). The first row in the table contains the numbers from the baseline case discussed in the main text. For brevity, we only present the sectoral greenium, not the decomposition. Our findings are very robust for the value-weighted case. For the equal-weighted case, the results concerning the sectoral greenium also remain essentially unchanged. Only the total greenium turns insignificant (t-statistic of 1.5) when we cluster by issuers and time.

5 Discussion

We have documented that investment funds, banks and insurance companies pay most of the greenium in our sample of European green bonds. But the decomposition of the investor group-specific greenium into the three subcomponents suggests that different mechanisms are at play for the different investor groups. Investment funds, on the one hand, overweight green bonds relative to conventional bonds on average, across our entire sample, irrespective of the individual greenium of a single bond. Such a pattern supports the narrative that investment funds (or their clients) have non-pecuniary sustainability preferences, i.e. they prefer sustainable assets for reasons that are orthogonal to the risk-return trade-off which usually determines optimal portfolios.

The greenium paid by banks, on the other hand, is markedly different. For one, banks overweight specific green bonds with a sizeable greenium, giving rise to a large covariance term (the third component in our decomposition). Moreover, banks underweight green bonds relative to conventional bonds on average in our sample. In the following, we want to discuss potential economic mechanisms behind this intriguing finding. Unfortunately, our unique dataset, which matches bond yields with confidential securities holdings, has limitations which do not allow us to dig much deeper into potential mechanisms quantitatively. Still, at the very least, we can explicitly formulate a few requirements that any candidate mechanism would need to satisfy, also keeping an eye on the sample splits in the previous section, in order to guide future research on this topic.

Broadly speaking, a theoretical channel that is supposed to rationalize the large covariance term for banks needs to satisfy three conditions. (a) It needs to rationalize why banks are overly invested in some green bonds, but not in others. That is, the mechanism needs to explain cross-sectional variation in $\Delta s_{i,\mathrm{MFI}}$. Our findings indicate that banks are *on average* underinvested in green bonds, but they do invest in green bonds very *selectively.* (b) A candidate mechanism needs to pick up cross-sectional variation in the greenium g_i across green-conventional bond pairs, beyond just giving an explanation for why a negative greenium exists on average. (c) It needs to rationalize why the cross-sectional variation in $\Delta s_{i,x}$ and in g_i should be aligned for banks, but not aligned for other sectors. Ideally, a mechanism should also be able to explain why this alignment is more pronounced in specific market segments (young, small, and financial bonds) and less pronounced in others, and why the credibility of the issuer commitment, as measured by the CBI alignment flag, does not appear to play a role. In the following, we weigh a few hypotheses against these requirements.

The intermediary asset pricing paradigm has been developed to explain the empirical link between intermediaries' financial frictions and the dynamics of asset prices.[9] Assets for which financial intermediaries act as marginal investors trade at a discount relative to others when intermediaries are financially constrained. Being a marginal investor can, for instance, result from bond underwriting, market making or other transformation services offered by banks (see, e.g., Larcker and Watts (2020)). Intermediary asset pricing could satisfy some of the requirements above, but not in a straightforward way. First and foremost, by construction, it can explain cross-sectional differences in bond prices between various types of green and conventional bonds. However, it is a priori unclear why financial institutions should be the marginal investors for some green/conventional bonds, but not for others – requirement (a) above. Moreover, it is unclear why financial institutions should be marginal investors for a few selected financial (or young or small) green bonds, but not for other financial (or young or small) green bonds. Still, even if intermediary asset pricing may not be sufficient to explain the large covariance term of the greenium decomposition for banks, it could still be a rationale for the existence of an average greenium. Notably, financial intermediaries would then need to be marginal investors in conventional bonds, rather than in green bonds, because the average greenium is negative in our sample. Importantly, the potential link to intermediary asset pricing raises questions about reverse causality for our results. Obviously, we cannot claim that any investor group in our sample invests into specific green (or conventional) bonds *because* these bonds carry a greenium. We also cannot make the reverse statement, namely that the greenium of a particular green bond is large because some investor groups hold larger shares than others. The intermediary asset pricing paradigm would tend to favor the latter implication, rather than the former.

The market segmentation or preferred habitat narrative has gained popularity in the literature about yield curves. The idea is that certain investor groups (market makers, arbitrageurs, pension funds) predominantly operate in specific segments of the market only, for instance the long or the short end of the yield curve. Supply and demand of these different investor groups, together with group-specific frictions, then shape the dynamics of the yield curve. Boermans (2023) documents – also on the basis of SHS data – that mutual funds and pension funds have a green bond preferred habitat. However, such a green bond preferred habitat does not square with requirement (a) above because it rather relates to the average level of the greenium, not to its cross-sectional variation. Consequently, it may rather serve as an explanation for the second term of our greenium decomposition which we labeled and interpreted as "general green preference" before. In this spirit, our findings regarding mutual funds and pension funds are in line with Boermans (2023).

The fact that the covariance term of the greenium decomposition is particularly pronounced among financial bonds points towards specifics of the interbank bond market as another candidate mechanism. In fact, Bekaert and Breckenfelder (2019) show that banks have a very large share in the market for financial bonds. Jasova, Laeven, Mendicino, Peydro, and Supera (2023) trace this back to incentives arising from lender of last resort policy and central bank haircut

[9]See, e.g., He and Krishnamurthy (2013), Adrian, Ettula, and Muir (2014), or He, Kelly, and Manela (2017) for seminal contributions to this literature.

gaps. However, our summary statistics in Internet Appendix B confirm that the interconnectedness of banks is pronounced both for green and conventional financial bonds. Our results also suggest that the specificities of financial bonds essentially carry over from conventional to green bonds. If they were supposed to serve as an explanation for our findings, there should be cross-sectional variation in these specificities. In this regard, it is important to note that there is no differential treatment of green and conventional bonds with regard to central bank haircuts in our sample period. Finally, it is worth being mentioned that the average greenium g_x is statistically insignificant for green financial bonds. This makes our finding that banks pay a large part of their greenium for their holdings of very few selected financial green bonds even more striking.

This list of candidate mechanisms to explain our results is, of course, by no means exhaustive. For instance, banks may hold more or fewer green bonds because they care about their reputation as sustainable institutions, because they want to improve their portfolio diversification, or because green bonds are less liquid than conventional bonds. While all these motives may indeed play a role, we are skeptical that they can explain the results presented here, for the reasons that we have outlined above. Our findings are rather related to cross-sectional heterogeneity among green bonds, not to average differences between green and conventional bonds. Regulatory reasons can also be largely ruled out as a rationale behind our results. This is because our matching approach makes sure that credit risk (and thus the regulatory treatment) is identical within each green-conventional bond pair. Moreover, to date, there are no official regulatory requirements that prefer green over conventional bonds. In addition, for regulation to explain our findings, we would need a differential regulatory treatment across green bonds, which would then – on top – need to align with differentials in prices and security holdings.

6 Conclusion

Mitigating climate change as well as adapting to its potential consequences requires massive investments, both from the public and the private sector. At the same time, a tentative consensus has emerged that sustainable assets have lower expected returns than equivalent non-sustainable counterparts. This paper is concerned with the broad questions which investors pay the costs arising from greening investment portfolios – the so-called "greenium" – and why they do so. We elaborate on these questions in the context of green bonds using a unique dataset that contains both pricing and ownership information on green-conventional bond pairs.

Our main methodological contribution is a decomposition of the greenium along two dimensions: the ownership structure and the cross-section of green bonds. We start by establishing that the average greenium in our sample amounts to minus 3 basis points. Decomposing this average greenium along the bonds' ownership, we then document that it is largely borne by banks, investment funds and insurance companies (or their clients). Dissecting further along the cross-sectional dimension of bonds, we find that the greenium paid by investment funds is largely explained by an average level effect. Investment funds generally overweight green over matched conventional bonds, potentially reflecting strong non-pecuniary green preferences of their clients. On the other hand, the greenium paid by banks is markedly different. More precisely, banks display a tilt towards specific green bonds with a relatively pronounced greenium. This tilt is particularly sizeable when the sample is restricted to young bonds, small bonds, bonds with a long residual maturity, or bonds issued by the financial sector.

We draw the tentative conclusion that banks pay a significant greenium because they hold specific green bonds for motives other than green preferences. For instance, intermediaries may be marginal investors for certain green bonds because of market making, underwriting or liquidity management activities. The trade-off between offering such financial intermediation services and avoiding to pay the greenium may be particularly relevant, for instance, when operating in the segments of small, young or financial bonds.

Overall, our findings contribute to the broader debate about potential sources of price differentials between green and conventional financial assets. Importantly, they point towards a prominent role of drivers that go beyond investors' green preferences, such as financial intermediary frictions. But, as we discuss in our paper, these alternative channels also come with a series of caveats which need to guide future research. Our focus on green bonds is partly driven by data availability. Looking ahead, it would thus be interesting to conduct similar analyses for other segments of the market for sustainable assets. Moreover, regarding the interaction between financial frictions and the greenium, further analysis is needed to understand which way the causality runs. Lastly, we believe that the relationship between bond ownership and price differentials in general deserves further attention.

References

Adrian, T., E. Ettula, and T. Muir (2014). Financial intermediaries and the cross-section of asset returns. *Journal of Finance 69*(6), 2557–2596.

Azar, J., M. Duro, I. Kadach, and G. Ormazabal (2021). The big three and corporate carbon emissions around the world. *Journal of Financial Economics 142*(2), 674–696.

Baker, M., D. Bergstresser, G. Serafeim, and J. Wurgler (2022). The pricing and ownership of U.S. green bonds. *Annual Review of Financial Economics 14*, 415–437.

Bekaert, G. and J. Breckenfelder (2019). The (re)allocation of bank risk. Working Paper.

Boermans, M. (2023). Preferred habitat investors in the green bond market. *Journal of Cleaner Production 421*(1).

Boermans, M., M. Bun, and Y. van der Straten (2024). Funding the fittest? Pricing of climate transition risk in the corporate bond market. *DNB Working Paper No. 797*.

Bolton, P. and M. Kacperczyk (2021). Do investors care about carbon risk? *Journal of Financial Economics 142*(2), 517–549.

Bremus, F., F. Schütze, and A. Zaklan (2021). The impact of ECB corporate sector purchases on european green bonds. *DIW Discussion Paper No. 1938*.

Brøgger, A. and A. Kronies (2021). Skills and sentiment in sustainable investing. *Working Paper*.

D'Amico, S., J. Klausmann, and N. A. Pancost (2023). The benchmark greenium. *Working Paper*.

Darmouni, O. and M. Papoutsi (2022). The rise of bond financing in europe. *ECB Working Paper No. 2022/2663*.

De Haas, R. and A. Popov (2023). Finance and green growth. *The Economic Journal 133*, 637–668.

Delis, M., K. de Greiff, and S. Ongena (2021). Being Stranded with Fossil Fuel Reserves? Climate Policy Risk and the Pricing of Bank Loans. *Working Paper*.

Dorfleitner, G., S. Utz, and R. Zhang (2021). The pricing of green bonds: external reviews and the shades of green. *Review of Managerial Science 16*, 797–834.

Fatica, S., R. Panzica, and M. Rancan (2021). The pricing of green bonds: Are financial institutions special? *Journal of Financial Stability 54*.

Flammer, C. (2021). Corporate green bonds. *Journal of Financial Economics 142*(2), 499–516.

Hartzmark, S. and K. Shue (2023). Counterproductive sustainable investing: The impact elasticity of brown and green firms. *Working Paper*.

He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics 126*(1), 1–35.

He, Z. and A. Krishnamurthy (2013, April). Intermediary asset pricing. *American Economic Review 103*(2), 732–70.

Ilhan, E., Z. Sautner, and G. Vilkov (2021). Carbon tail risk. *Review of Financial Studies 34*(3), 1540–1571.

Jasova, M., L. Laeven, C. Mendicino, J. L. Peydro, and D. Supera (2023). Systemic risk and monetary policy: The haircut gap channel of the lender of last resort. *Review of Financial Studies: forthcoming*.

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 (2017). Does it pay to be green? *Working Paper*.

Krueger, P., Z. Sautner, and L. Starks (2020). The importance of climate risk for institutional investors. *Review of Financial Studies 33*(3), 1067–1111.

Larcker, D. and E. Watts (2020). Where's the greenium? *Journal of Accounting and Economics 69*.

Oehmke, M. and M. Opp (2020). A theory of socially responsible investment. *Working Paper*.

Pastor, L., R. Stambaugh, and L. Taylor (2021). Sustainable investing in equilibrium. *Journal of Financial Economics 142*(2), 550–571.

Pastor, L., R. Stambaugh, and L. Taylor (2022). Dissecting green returns. *Journal of Financial Economics 146*, 403–424.

Pedersen, L. H., S. Fitzgibbons, and L. Pomorski (2021). Responsible investing: the ESG-efficient frontier. *Journal of Financial Economics 142*(2), 572–597.

Pietsch, A. and D. Salakhova (2022). Pricing of green bonds - drivers and dynamics of the greenium. *ECB Working Paper*.

Sangeorgi, I. and L. Schopohl (2021). Why do Institutional Investors Buy Green Bonds: Evidence from a Survey of European Asset Managers. *International Review of Financial Analysis 75*.

Sangeorgi, I. and L. Schopohl (2023). Explaining green bond issuance using survey evidence: Beyond the greenium. *British Accounting Review 55*(1).

Starks, L., P. Venkat, and Q. Zhu (2019). Corporate ESG profiles and investor horizons. *Working Paper*.

Tang, D. and Y. Zhang (2020). Do shareholders benefit from green bonds? *Journal of Corporate Finance 61*.

Yang, A. (2021). Green investor clientele effects. *Working Paper*.

Zerbib, O. (2019). The effect of pro-environmental preferences on bond prices: Evidence from green bonds. *Journal of Banking and Finance 98*(1), 39–60.

Zerbib, O. (2022). A sustainable capital asset pricing model (S-CAPM): Evidence from green investing and sin stock exclusion. *Review of Finance*.

Zhang, H. (2022). Climate policy and corporate green bond issuances. *Working Paper*.