The Low Frequency Trading Arms Race: Machines Versus Delays *

Alexander Dickerson† Yoshio Nozawa‡ Cesare Robotti§

April 30, 2024

Abstract

We propose a novel framework to compute transaction costs of trading strategies using infrequently traded assets. The method explicitly accounts for the trade-off between bid-ask spreads and execution delays. The benefit of waiting for a better trading opportunity with lower bid-ask spreads is partly offset by the opportunity cost of delayed or missed execution. Applying this method to corporate bonds that trade infrequently, we show that even the latest machine-learning-based trading strategies earn zero or negative bond CAPM alphas after transaction costs. Consequently, our results raise doubts about the realistic outperformance capabilities of active bond trading strategies relative to the bond market factor.

JEL Classification: G12, G13

Keywords: Corporate Bonds, Liquidity, Machine-Learning, Market Efficiency, Fixed-Income Securities, Credit Risk

*We thank Andrew Chen, Oleg Chuprinin, Christian Julliard, Savina Risova, Sam Wang, Alex Zhou, and the seminar participants at Atlanta Fed, Temple University, the University of Georgia, the University of New South Wales, the University of North Carolina Chapel Hill, the University of Technology Sydney and the University of Toronto for helpful comments. We also thank Julien Vincent for cloud computing support.

†UNSW Business School, `alexander.dickerson1@unsw.edu.au`.

‡University of Toronto, `yoshio.nozawa@rotman.utoronto.ca`.

§Warwick Business School, `cesare.robotti@wbs.ac.uk`

1 Introduction

Academic research on the cross-section of equity returns has been extremely successful, and has fundamentally changed the way practitioners invest in stocks. Against this backdrop, there is a growing trend in the literature to apply identical portfolio formation methods originally developed for stocks to less liquid, infrequently traded assets such as corporate bonds. Does the verbatim application of these portfolio construction methodologies lead to an accurate description of the performance of corporate bond investment strategies and factors? Our answer is no.

With illiquid assets, an investor cannot immediately execute buy and sell orders to build a portfolio of securities after observing a set of investment signals. Instead, she must wait for her order to be executed due to search costs, dealer inventory constraints and bargaining frictions. This creates delays and drags down the performance of her portfolio as the investment signal becomes outdated. Even worse, the order may not be executed over the period for which the signal was intended (and valid) for, in which case she misses the investment opportunity and incurs the opportunity cost of capital. In addition, the delay in one leg of a long-short strategy relative to another creates basis risk and reduces the intended hedging benefit. Therefore, ignoring these costs severely distorts the assessment of the profitability of factor investing in illiquid assets. In essence, the immediate order execution assumption implicit in equity-based portfolio construction does not apply to corporate bonds or any asset which is infrequently traded. This key friction has been overlooked within the context of forming realistic corporate bond factors and portfolios.

In this paper, we impose empirical realism to the construction of corporate bond portfolios by explicitly taking into account the nuanced relationship between trading costs and delays. Our strategy considers an investor's preference for early order execution. An impatient investor is willing to pay higher bid-ask spreads in exchange for quick execution, while a patient investor waits for a trading opportunity with a tight bid-ask spread. To implement this idea, we exploit a key feature of corporate bonds pointed out by Edwards, Harris, and Piwowar (2007), where observed bid-ask spreads are a decreasing function of trade size. While we do not attempt to explain why bid-ask spreads depend negatively on size, we take this empirical fact as given and describe the key trade-off between delays and bid-ask spreads.

Consider an investor who receives a buy signal in a month. Given her portfolio size, she needs to buy $2 million of the bond. She has the choice of placing a large $2 million order and waiting for the execution, which could take a month or more. Or she can break the order into smaller pieces and execute it more quickly. In the latter case, unlike in the equity market, she will have to pay a *higher* price because of the costs charged by a dealer.[1] This fundamental tension, between trade size and the cost of delaying the trade has yet to be explored within the context of corporate bonds or other assets that trade infrequently.

To render this idea operational, we no longer assume that the trade occurs at the end of the month and compute the exact return of a bond from the day it is bought to the day it is sold. For each date, we compute transaction prices using only transactions above the trade size threshold. As the threshold increases, the bid-ask spreads tighten while the number of eligible transactions declines, delaying the trade.

The key to our method is to allow a monthly return to exist even if there is no transaction for a bond in that month. Instead of treating such an observation as missing, our method treats it as a trade execution failure. Consider a case in which an investor intends to buy a bond but there are no transactions above her size threshold in a month. Since the investor does not know when or if her order will be executed, the capital tied up in this long position cannot be used to buy other bonds and thus earns the risk-free rate of return. The difference between the corporate bond returns she would have earned by buying other bonds and the risk-free rate contributes to the cost of delay. If, on the other hand, she wants to unwind the existing position but is unable to do so due to delays or lack of trades, she will earn a mark-to-market return based on quoted prices, but will pay the cost of carry to finance the additional unwanted positions.

We show that the cost of delay is substantial. Consider a simple example in which the investor

[1] This analysis is reminiscent of the trade-off in the stock market. When trading stocks, a trader must consider the benefit of breaking large trades into smaller pieces that are executed over a longer period of time. The key question there is how to reduce the price impact by swallowing longer delays. Since equity trades are anonymous, a liquidity provider learns the informativeness of the order by its size and charges a high spread for a large trade. The key problem for the investor is how to overcome this adverse selection problem. Therefore, even though the size-cost relationship is the opposite in the equity market, there is still a trade-off. Jacobsen and Venkataraman (2023) argues that in the bond market, investors do not necessarily split a large trade because dealers knows their identity and thus splitting does not help hide private information.

Figure 1: Effect of Transaction Costs: Example of Credit Spread-Sorted Portfolio

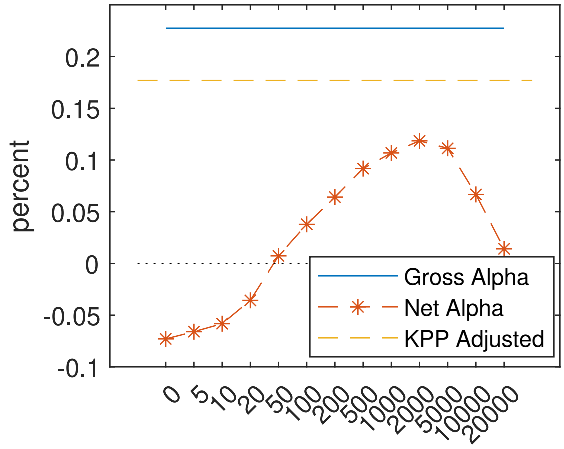

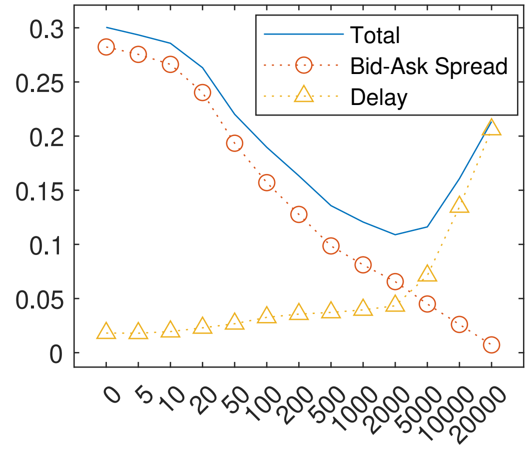

This figure plots the bond CAPM alphas of the long-short strategies based on corporate bonds' credit spreads before and after accounting for transaction costs (left panel). The transaction costs are decomposed into the bid-ask spread costs and delay costs (right panel). Values on the x-axis are the trade size in thousand dollars.

purchases bonds with the top 20% highest credit spreads and sells short those with the bottom 20% lowest spreads. The left panel of Figure 1 plots the bond CAPM alpha on this long-short strategy before and after transaction costs as a function of trade size. We observe a hump-shaped pattern in net returns, implying our cost estimates are a U-shaped function of transaction size. This transaction cost can be decomposed into half spreads and delay costs.

The right panel plots the cost of half spreads, capturing both the portfolio turnover rate and the difference between bid and ask prices, for each trade size. Consistent with Edwards, Harris, and Piwowar (2007), there is a strong negative relationship between half spreads and size, indicating a significant benefit to being patient and trading in large volumes. However, insisting on trading in large volumes causes delays in order execution. As a result, the cost of execution delays increases as the trade size increases. In this example, as the trade size becomes larger than \$2 million, the increased cost of the trade delay *outweighs* the reduced half-spread. Therefore, the optimal trade size that maximizes net profit is \$2 million.

To quantify the importance of delay costs, we use the latest machine learning (ML) algorithms to generate trading signals with 200 bond and equity-based characteristics. We use ML-based strategies for two reasons: First, the cost of delay becomes more important the more valuable the signal is. Thus, to emphasize our point about the importance of delay, it is appropriate to use strategies that perform optimally before transaction costs are taken into account. Second, there is a growing literature on how to assess the profitability of factor investing in corporate bonds. By using ML-based strategies, we can directly provide a method to adjust for realistic transaction costs for the most popular strategies today.

Our ML algorithms reflect the state-of-the-art models tested in the recent literature (e.g., Gu, Kelly, and Xiu 2020). We estimate a large set of models using a wide array of bond and stock characteristics.[2] These encompass linear models with penalization, regression tree ensembles (including extreme randomized trees and random forests) and feed forward neural networks. We use the machine learning implied model predictions of bond returns to form long-short portfolios that purchases bonds with high expected returns and short-sells bonds with low expected returns. All of the long-short ML strategy portfolios generate out-of-sample gross returns that are economically large and statistically significant (Newey-West adjusted t-statistics greater than 3). Importantly, the alphas of these strategies computed with the single-factor bond CAPM (CAPMB) remain large and significant. Individually, only a handful ($\sim$10 %) of the stock and bond characteristics generate meaningful high-low gross return spreads, which highlights the importance of combining the characteristics to form predictions through the various ML methods we employ.

Our methodology allows us to calculate transaction costs under optimal execution. We choose the trade size that maximizes the net CAPM alpha of each strategy and find that the optimum is reached between \$2 million and \$10 million per trade. For example, the ensemble ('ENS') strategy which averages the expected returns of all ML strategies generates an alpha of 0.48% before transaction costs and 0.07% after costs at the optimum. Of the 0.41% cost, 0.19% is due to delays, while 0.22% is due to bid-ask spreads paid to the dealer. Thus, quantitatively, the cost of delay is substantial, and ignoring it leads to an incorrect assessment of the profitability of ML strategies and other corporate bond anomalies. Importantly, the techniques that have high predictive power tend to move swiftly and thus incur high transaction costs. Net of costs, even when trading at the optimal volume threshold, which captures the trade-off between reduced half-spreads

[2] The characteristics comprise 27 bond characteristics and 173 equity-based characteristics.

and trading delays, *all* of the strategies generate a single-factor alpha of close to zero.

To guide future research, and in the spirit of Harvey, Liu, and Zhu (2016), we provide a set of gross alpha "cut-offs" at various levels of portfolio turnover rates that represent the level of alpha the factor should achieve to remain profitable after costs. For example, to achieve a net alpha of 0.2% per month, a strategy with a monthly turnover rate of 10%, 20%, 30% needs to earn a gross alpha of 0.36%, 0.47%, and 0.61%, respectively. These cutoffs serve as a simple heuristic, allowing researchers to quickly check whether their gross factor alpha would remain significant at various levels of turnover after accounting for transaction costs. Because we compute transaction costs under the assumption of optimal trade size, the researcher no longer has the freedom to choose the trade size to achieve the desired results. Instead, the net profit of the strategy we compute is disciplined by the realized trade size and frequency in the data.

One potential concern about our negative findings on the performance of ML strategies is that the particular algorithms and bond characteristics we use may not be the best available in practice. To address this criticism, we turn to an analysis of mutual fund returns. We obtain the actual returns earned by corporate bond mutual funds over our sample period. We show that, on average, only 8.5% of all "corporate bond" classified mutual funds (42 funds) generated an after-cost alpha that is statistically significant at the 5% nominal level. The magnitude of the statistically significant alpha is small at 0.18% per month.

Even more discouraging, from the perspective of an active bond mutual fund investor, is the dollar "value-add" of investing in active funds relative to a passive benchmark bond market portfolio. On average, bond mutual fund investors are *worse off* to the value of $396,000 per month relative to simply holding a corporate bond ETF that tracks the market. The cumulative loss that is accrued by active investors relative to simply holding the bond market portfolio is close to $55 million. These results support the validity of our assessment of ML-based corporate bond strategies and other corporate bond anomaly portfolios.

In summary, this paper contributes to the literature on two fronts: First, we introduce a novel methodology for computing portfolio returns that explicitly conditions on realized trade sizes and accounts for trading delays induced by attempting to transact in large volumes. These methods can be applied to any infrequently traded asset and allow researchers to identify the optimal execution, striking a balance between bid-ask spreads and delays. Second, we contribute to the assessment of market efficiency and the profitability of factor investing in the corporate bond market. Overall, our results suggest that, even when using state-of-the-art portfolio construction techniques, generating alpha from systematic bond strategies is an extremely challenging task once market frictions are properly accounted for.

Our paper contributes to the rapidly growing literature that evaluates (and re-evaluates) the performance of factor investing in the corporate bond market (e.g., Bali et al. 2020; Kelly et al. 2021; Sandulescu 2022; Binsbergen et al. 2023; Dickerson et al. 2023; Dick-Nielsen et al. 2023). The paper closest to ours is Ivashchenko and Kosowski (2023), who study the performance of nine factors after accounting for transaction costs. Our paper differs from Ivashchenko and Kosowski (2023) in that we highlight the novel trade-off between half spreads and delays faced by investors and employ the latest machine learning techniques in testing the performance of factor models.

This paper also relates to the extensive literature measuring illiquidity and transaction costs in the corporate bond market (e.g., Edwards et al. 2007; Chen et al. 2007; Feldhütter 2010; Bao et al. 2011; Schestag et al. 2016; Dick-Nielsen and Rossi 2018; Pinter et al. 2021; Choi et al. 2023b). In particular, Bao et al. (2018), Bessembinder et al. (2018), and Wu (2022) examine the role of post-crisis regulations on the liquidity of corporate bonds.[3] More closely related papers include O'Hara et al. (2018), who examine the market power in determining corporate bonds' half spreads, as well as Goldstein and Hotchkiss (2020) and Reichenbacher and Schuster (2022), who argue that observed transaction costs strongly depends on transaction size and dealers' strategic inventory management. However, none of these papers quantify the impact of trading delays in evaluating trading strategies.[4]

Our paper aims to provide the best practice in accounting for transaction costs in the study of

[3]More broadly, there is a strand of literature that studies the role of liquidity and dealer inventory in explaining credit spreads and bond risk premiums. This body of research includes Lin et al. (2011); Friewald and Nagler (2019); He et al. (2019); Goldberg and Nozawa (2021); Eisfeldt et al. (2023).

[4]Goldstein and Hotchkiss (2020) note the trade-off similar to the one we propose in the paper. In particular, they write "dealers will offer customers a trade-off between pricing and immediacy (liquidity). However, ... Dealers provide little immediacy when there are few trading opportunities. For example, for a bond that trades at best once a month, investors retain price risk while dealers search for a counterparty to offset their trade,...". Our results empirically support the significance of their statement from the perspective of factor investing.

the cross-section of corporate bond returns. Table 1 lists recent papers on this topic. The papers are classified into two groups: the first group of papers does not consider net returns after transaction costs and the second group does so. However, even among the papers in the second group, there is substantial heterogeneity in the transaction cost estimates. For example, Bali et al. (2020), Jostova et al. (2013), and Kelly et al. (2021) report significant trading profits arising from anomalies after accounting for transaction costs while Chordia et al. (2017), Bartram et al. (2023), and Nozawa et al. (2023) report anomalous returns largely disappear net of costs.

The discrepancy arises because there is substantial room for researchers to make judgments on how to estimate half spreads. For example, Cao et al. (2023) reports the profit from their trading strategy is significant if they assume each transaction is of size $1 million but not significant if the transaction size is smaller. To avoid the subjective selection of trade size, we provide an exogenously-specified trade size that captures the reality to discipline the estimated transaction costs.[5]

The remainder of the paper is organized as follows: Section 2 provides detailed methods for calculating portfolio returns net of transaction costs; Section 3 describes our data set; Section 4 provides the evaluation of the ML-based strategies; Section 5 examines the performance of corporate bond mutual funds; and Section 6 provides concluding remarks.

2 Trade-Offs in the Corporate Bond Transactions

In this section, we revisit the previous finding that a transaction cost declines with trade size. There can be two explanations for this pattern. First, Duffie et al. (2005) argue that trade size is a proxy for the investor's size. Since a large investor has better outside trading options, she can trade with a dealer with a lower cost. Second, the declining pattern is the reflection of bond dealers' strategic behavior which leads to bias in the data. In this explanation, dealers act as brokers for large transactions but take small transactions in their inventory with a charge. As a result, realized large transactions appear to have lower costs than small ones.

[5] The importance of transaction costs in other asset classes, such as stocks and options, are documented in Novy-Marx and Velikov (2015), Novy-Marx and Velikov (2019), Chen and Velikov (2023), Detzel et al. (2023), and Avramov et al. (2023).

To dissect into these drivers of transaction costs, we examine insurance firms' transactions in eMAXX. We follow O'Hara and Zhou (2021) to compute the half spread for each trade. Let $\bar{P}$ denote the transaction price of the latest interdealer trade that occurred before a customer-dealer trade. Then, the cost of the customer-dealer trade for bond k is

$$c_{i,k,t} = (\log P_{i,k,t} - \log \bar{P}_{k,t}) \times 1_{i,t} \tag{1}$$

where $1_{i,t}$ is 1 if the trade is customer buy and -1 if the trade is customer sell. We allow the reference trade to occur within 5-business-day window preceding the trade. If there are no interdealer trades within this window, the half spread for the trade is treated as missing. We winsorize the half spreads at the 0.5 and 99.5 percentiles to attenuate the influence of outliers. We add to the data set the information about the insurance firm's size for each trade. Using EMAXX's firm id, we aggregate all corporate bond holdings in quarter q and merge the data to all transactions by the firm in quarter $q + 1$.

We begin the analysis by reporting the average and median half spreads by trade size and investor size category. For trade size, we classify trades into four groups, including (0,100,000], (100,000,500,000], (500,000, 1M], (1M,∞]. For investor size, we use the 20, 40, 60, and 80 percentiles of the size distribution for the cutoff.

Table 2 reports the mean and median half spreads in each category. We find that the half spreads decline in both investor and trade sizes. Controlling for trade size, the average and median spread declines as we move from the small investor category to the large investor category. This finding is consistent with O'Hara et al. (2018) and confirms the prediction of Duffie et al. (2005). However, controlling for investor size, the spreads also decline from small to large trades.

The number of observations shows that investor size and trade size are positively correlated, as expected. However, there are plenty of observations in which large investors trade with a small size. For example, there are 17,489 observations where the top 20% insurance firms trade with a size less than \$100K. The smallest insurance firms still trade with a size of more than \$1 million: there are 55, 636 of such instances in our data.

To estimate the impact of trade size more precisely, we follow Pinter et al. (2021) and estimate the following trade-level panel regression:

$$\mathrm{c}_{i,k,t} = \beta_v \times \log(\text{Volume})_{i,t} + \tau_t + \alpha_i + \lambda_b + \mathbf{Y}'_{k,t}\delta_C + \varepsilon_v, \tag{2}$$

where $\mathrm{c}_{i,k,t}$ is the trading cost as computed in Equation 1 for client i trading bond k on day t, $\log(\text{Volume})_{i,k,t}$ is the natural logarithm of the given trade's notional, τ_t are trade-day fixed effects, α_i (λ_b) are client (broker) fixed effects respectively. The vector $\mathbf{Y}_{k,t}$ captures bond-level characteristics that we use as controls including bond rating, coupon, maturity and the log of the issue size (amount outstanding). The key coefficient of interest is the estimated value of β_v: if the half-spread continues to decrease after controlling for the client base, we would expect β_v to remain negative and statistically different from zero.

In columns (1) and (4), we confirm the results from prior research documenting a negative relationship between half-spreads and trade size (volume) without and with bond-level controls. In columns (2) and (5), after including client-level fixed effects, the sign and magnitude of the β_v coefficient remains unchanged, reinforcing our results from the double sort above. It would appear that the 'Size Penalty' does not exist in the U.S. corporate bond market. As an additional test, we also include broker fixed effects in columns (3) and (6). β_v remains robustly negative.

The economic significance of the size premium (captured by β_v) is sizeable. A one standard deviation increase in the log of trading volume decreases half-spreads by about 0.104% in column (1).[6] Relative to the average half-spread of 0.179%, this change is economically large. Once including client fixed effects in column (2), the effect remains sizeable and results in a 0.07% decrease in half-spreads.

The evidence from the U.S. insurance firms' transactions shows that both investors' bargaining power and dealers' strategic behavior are at work in explaining the observed transaction costs. The second channel suggests that for an investor, there is a hidden cost of trading with a large size that is missed by the observed half spread. To capture this cost, we turn to the frequency of trades in

[6] The standard deviation of the log of trading volume is 1.73.

the data. We show that large trades occur infrequently, suggesting that investors bear the cost of delayed or missed trades if they desire to buy or sell a large quantity.

3 Methodology

We develop a methodology to compute net returns for bonds that explicitly accounts for half spreads and execution delays. The core idea behind this method is that, when trade executions are delayed, an investor may end up with unintended positions or may not initiate the trade at all. We carefully treat each case by studying which positions must be financed through risk-free lending and borrowing and by keeping track of inventory positions each month.

At the end of month t, an investor receives signals and decides which bonds to go long on and which bonds to short. She tries to execute the trade as soon as a trade opportunity with her target volume arrives. At the earliest, she executes at the end of month t, but more typically she would trade in month $t+1$ or later. For computational simplicity, we separately consider delays within a month and delays beyond a month. The delay beyond a month is considered as a part of her inventory, and this affects her action at the end of $t+1$.

Delays directly affect the return computation. If she executes the trade, she pays the half spreads and starts earning returns from the position. Before she does so, her position earns the risk-free rate of returns. Our method below explicitly accounts for the delay using daily transaction prices.

3.1 Returns with Execution Delays

This section explains our return construction. Suppose an action is taken at the end of month t and we want to measure the monthly return of a strategy from month t to $t+1$. The investor's possible actions in month t for each bond are buy, hold, or sell. If she buys, then she trades on the ask side and if she sells, then she trades on the bid side. If she holds, her position is marked to market using quotes. In such cases, a return on the bonds in her long positions can be described by one of the following patterns:

- Hold-Hold (hh): R^{hh}
- Buy-Hold (bh): $R^{b(v)h}$
- Hold-Sell (hs): $R^{hs(v)}$
- Buy-Sell (bs): $R^{b(v)s(v)}$

where $b(v)$ is the buy order with a minimum volume v, $s(v)$ is the sell order with a minimum volume v, and h indicates that she holds the position. The first of two superscripts for R describes the investor's action in month t and the second one describes her intended action in month $t+1$. These actions are taken when an opportunity with the minimum volume v arrives.

If an investor already holds a bond and maintains her position throughout month $t+1$, then her return can be measured using a standard formula,

$$R_{t+1}^{hh} = \left(\frac{P_{t+1}^{h} + AI_{t+1} + C_{t+1}}{P_{t}^{h} + AI_{t}} \right) - 1, \tag{3}$$

where P_{t+1}^{h} and P_{t}^{h} are the end-of-month quotes in months $t+1$ and t, respectively. AI_t is the accrued interest at the end of month t and C_{t+1} is any coupon paid in month $t+1$. Since there is no trade, it does not take half spreads into account when measuring the return. When we compute gross returns and alphas of a strategy, we use this return for all bonds in all months.

If an investor initiates a new long position, then she has to pay an ask price. In addition, if there is a delay in a buy order, then she earns a risk-free rate on cash while waiting for her order to be executed. Suppose she wants to buy a bond in month t and hold it until month $t+1$, and she buys the bond on the d-th day in month $t+1$, then her return is

$$R_{t+1}^{b(v)h} = \left(1 + R_{t+1}^{f} \times \frac{d}{Days_{t+1}} \right) \left(\frac{P_{t+1}^{h} + AI_{t+1} + C_{t+1,d}}{P_{t+1,d}^{b(v)} + AI_{t+1,d}} \right) - 1, \tag{4}$$

where $P_{t+1,d}^{b(v)}$ is the ask price on either the last business day of month t or day d in month $t+1$. $Days_{t+1}$ is the number of days in month $t+1$ and $C_{t+1,d}$ is any coupon paid after day d in month $t+1$. We use subscript d for the observation on a specific day in a month. If a variable does not have a d subscript, then the variable is measured at the end of a month. If there are multiple daily prices for $P_{t+1,d}$, we use the first available day, capturing the idea that an investor is trying to implement the strategy as soon as possible.

Since we explicitly take into account which side of the market the investor is trading, the return in (4) measures the net return after transaction costs. This return is not only influenced by a half spread (i.e. the difference between ask price $P^{b(v)}_{t+1,d}$ and mid quote on the same day) but also by when the trade to initiate the position is executed. Until the corporate bond is bought, the cash is invested in risk-free asset, incurring an opportunity cost. If the delay becomes extreme and no transaction price is available in month $t+1$, then she cannot execute the trade and her return is the risk-free rate (i.e., $R^{b(v)h}_{t+1} = R^f_{t+1}$).

To illustrate the idea, consider an example where an investor receives a buy signal for a bond on September 30. At this point, she commits cash to the position and waits for a trading opportunity to arrive. Suppose that an opportunity of the size \$100,000 arrives on September 30, that of \$500,000 arrives on October 10, and that of \$1 million arrives on November 10. If her target trade volume v is \$100,000, then she buys the bond on September 30 and the bond's October return is the one-month return on the bond less the half spread paid to enter the position.

If, instead, her target trade size is \$500,000, then she waits for her order to be executed until October 10. Her October return is the product of the risk-free rate of return for the first ten days and the 21-day returns on the corporate bond. If her target size is \$1 million, she does not execute the trade and the October return is the risk-free rate. If the updated signal on October 31 is still a buy signal, then she would buy the bond on November 10, which contributes to the November return. If, on the other hand, the October signal is 'not buy', then she misses this buying opportunity entirely.

Similarly, when the investor unwinds the long position she already has, she executes the sell order on the d-th day of month $t+2$ or the end of month $t+1$, if possible. Then, the return is

$$R^{hs(v)}_{t+1} = \left(\frac{P^{s(v)}_{t+2,d} + AI_{t+2,d} + C_{t+1} + C_{t+2,d}}{P^h_t + AI_t} \right) \div \left(1 + R^f_{t+2} \times \frac{d}{Days_{t+2}} \right) - 1, \tag{5}$$

where $P^{s(v)}_{t+2,d}$ is the bid price on either the last business day of month $t+1$ or day d in month $t+2$, and $C_{t+2,d}$ is any coupon paid before day d in month $t+2$. We divide this return by the month $t+2$ risk-free rate because the investor must finance her extra long position by borrowing cash until she unloads it. If the sales does not occur, then her return is $R^{hs(v)}_{t+1} = R^{hh}_{t+1}$ and the bond is added to month $t+1$ inventory.

Continuing on the example, consider that the investor bought a bond in October and the October signal is 'buy' and thus she keeps the long position. Suppose further that the November signal is 'not buy'. The sell opportunity with a size of \$100,000 arrives on November 30, but that of \$500,000 arrives on December 10, and that of \$1 million arrives on January 10 next year. Then, for an investor with a target size of \$100,000, the bond's November return is a one-month return on the bond adjusted for a half spread. If her target size is \$500,000, then she earns a 40-day return on the corporate bond minus the 10-day risk-free rate in December. If her target is \$1 million, her November return is the buy-and-hold one-month return on the bond, creating an extra inventory influencing her portfolio choice at the end of December.

Finally, if an investor buys a bond in month t and sells it in month $t+1$, then her net return is

$$R^{b(v)s(v)}_{t+1} = \left(1 + R^f_{t+1} \times \frac{d_1}{Days_{t+1}}\right)\left(\frac{P^{s(v)}_{t+2,d_2} + AI_{t+2,d_2} + C_{t+1,d_1} + C_{t+2,d_2}}{P^{b(v)}_{t+1,d_1} + AI_{t+1,d_1}}\right) \div \left(1 + R^f_{t+2} \times \frac{d_2}{Days_{t+2}}\right) - 1. \tag{6}$$

If the purchase does not occur in month $t+1$ (i.e., a delay of more than a month), then $R^{b(v)s(v)}_{t+1} = R^f_{t+1}$. If the purchase occurs but the sales is delayed by more than a month, then $R^{b(v)s(v)}_{t+1} = R^{b(v)h}_{t+1}$.

In our main results, we allow an investor to short bonds. Her net return for short positions can be described similarly using $-R^{hh}_{t+1}$, $-R^{s(v)h}_{t+1}$, $-R^{hb(v)}_{t+1}$, and $-R^{s(v)b(v)}_{t+1}$.

Our method avoids the two problems that have plagued the literature studying corporate bond returns. The first is the martingale approximation of bond prices, as pointed out by Bartram et al. (2023). Previous research using TRACE data treats a transaction price near the end of a month as the month-end price. Since this is an approximation, there is no guarantee that real-time investors can trade a bond at this month-end price. Furthermore, the noise in prices tends to inflate the average returns due to Jensen's inequality (Blume and Stambaugh 1983).

Second is the censoring of returns. Typically, if there is no month $t+1$ return in TRACE due to a lack of transactions, one assumes that investors do not consider these bonds as trading targets and do not include the observation in the analysis. This creates a look-ahead bias because the real-time investor receiving the time-t signal does not know whether the bond will be traded in the next month or not. In addition, this censoring biases the sample towards liquid bonds by omitting illiquid bonds from the computation. In our framework, all bonds with a valid signal are considered for trading and the investor commits capital to take positions. If the trade does not occur, she earns or pays the risk-free rate or mark-to-market return, which allows us to closely replicate the real-time investor's trading profits.

3.2 Inventory and Round-Trip Transactions Over Months

In this section, we introduce inventory to account for the delays beyond one month. If there is no eligible trade in a month, an investor cannot execute the intended trade as suggested by a signal, creating a gap between the signal and the position. This gap, in turn, affects her actions in the next month. To illustrate the idea, we continue with the previous example of buying a bond at the end of September and selling it at the end of December. In Appendix A, we perform a formal analysis and provide a complete set of scenarios for both long and short positions.

Panel A of Figure 2 illustrates a base case with no delays. The investor earns the returns over the three consecutive months (October, November, and December), and the month-end inventory changes accordingly. The key here is that at the end of October, the investor has a bond in her inventory (the 'Y' sign for 'Inventory'). Thus, if she receives a buy signal, she can hold the bond and earn a mark-to-market return (R^{hh}) in November. Therefore, the half spread is charged only in October (when she buys) and in December (when she sells).

Now consider the case where the investor cannot buy the bond in October due to the trade failure, as shown in Panel B. Then her end-of-month inventory is 'No' (inventory), and she must buy a new bond in November. Due to the change in her action, her return in November is R^{bh}, incurring a half spread. Comparing Panels A (no delay) and B (purchase delay), there is no difference in the signals. The difference is the inventory dynamics due to the delay, and this changes the November return of the strategy.

Finally, Panel C of Figure 2 explains the case of delayed sales. In this example, the intended sale does not occur in December, creating an unwanted inventory at the end of December (the 'Y' sign). Since the signal at the end of December is still 'N', the investor sells the bond at the end of January. In this case, her return is $R^{hs} - R^f$. She pays not only a half spread to execute the sale, but also the risk-free rate to finance the additional position.

In our algorithm, we keep track of the inventory of each bond, decide the investor's action in the next month, and select the appropriate type of return. This method allows us to calculate the returns of a trading strategy without approximation.

3.3 Portfolio Formation

To measure the performance of a trading strategy net of costs, we must explicitly account for changing compositions in a portfolio. This is a challenging task because a standard portfolio construction prescribes a constantly changing portfolio weight for each security. To see this, consider three bonds as potential buy targets: A, B, and C. They have market values of $80 million, $40 million, and $20 million, respectively. Suppose in a month a signal suggests that an investor should buy A and B. Then her portfolio weight is 66.7% for A and 33.3% for B. Suppose also that the next month, the signal suggests that she should instead buy B and C. Then the weight for B increases to 66.7% from 33.3% the previous month. Thus, even though the signal for B has not changed, she must buy a fraction of B to increase its weight, incurring a transaction cost. This adjustment results in different cost-adjusted returns for the same bond in the same month, because some of the positions in B incur zero transaction costs, while others require her to pay a half spread when she enters a new position.

One way to overcome this problem is to set a constant fraction of the market value, rather than a constant number of bonds, for the long and short positions in each month. Typically, one would divide N bonds into P portfolios so that each portfolio has approximately the same number of bonds, N/P. This is achieved by categorizing bonds based on their signal percentile rankings.

To avoid trading a fraction of bonds, one can instead define a strategy by dividing bonds into P portfolios so that each portfolio has the same total market value. In this case, the cutoff is set by the value-weighted percentile rankings. In this way, each bond in the long position always has the same weight as long as it is in the portfolio. For example, if the total market value of the long position is set to $1 billion and the investor receives a buy signal for Bond A, then the bond will always have a portfolio weight of 8% in the long position, regardless of the other bonds in the portfolio. This method allows us to describe a position on each bond as a simple binary choice between 'Y' and 'N', obviating the need to adjust the existing position by a small amount. It not only reduces the complexity in portfolio return computation but matches the reality that bond investors won't adjust their positions from, say, $1 million to $1.05 million simply because it is too costly to make such adjustments.

4 Data

4.1 Data for the Machine Learning Return Predictions

Our datasets include daily bond data from Enhanced TRACE (TRACE) and the constituent bonds from the Bank of America (BAML) Investment Grade and High Yield indices as made available via the Intercontinental Exchange (ICE). We source equity and accounting data from CRSP and COMPUSTAT. We filter the data using standard approaches as prescribed by the literature which is explicitly described in Internet Appendix A. To train the machine learning models, we construct commonly used bond and equity variables used in the literature and then merge these to the equity-based characteristics from Chen and Zimmermann (2022).[7] This data combines several monthly bond and stock-based characteristics that have been shown the predict one-month ahead future corporate bond excess returns. Our data includes 200 characteristics of which 27 are bond-based characteristics and 173 are equity-based.

Detailed descriptions of the construction of these variables are provided in Table A3 of the Appendix. Missing characteristic data is set to its cross-sectional median at each month t. All

[7] openassetpricing.com. We thank Andrew Chen and Tom Zimmermann for making their data publicly available.

characteristics are cross-sectionally rank demeaned to lie in the interval $[-1,1]$. Overall, the data used to train the ML models with the 200 stock and bond characteristics comprises 17,815 bonds issued by 1,913 firms over the sample period from January 1998 to December 2022 ($T = 288$).

4.2 Data for Net Returns

To compute the net returns of the strategies, we combine daily data from both the TRACE and ICE data sets. We use dealer-customer trades in the TRACE data, filtered as described in Internet Appendix A. We then compute the simple average of transaction prices separately for bids ($P_{t,d}^{s(v)}$) and asks ($P_{t,d}^{b(v)}$) on a day, using only transactions with volume above the cutoff, v. We use the size cutoff of \$0, \$5,000, \$10,000, \$20,000, \$50,000, \$100,000, \$200,000, \$500,000, \$1 million, \$2 million, \$5 million, \$10 million, and \$20 million. The eleven cutoff values between \$5,000 and \$10 million are set following Edwards et al. (2007) and we add \$0 and \$20 million as the minimum and the maximum. Since v is a lower bound, a higher value of v leads to a smaller number of transactions included in the averages.

We then merge the daily transaction prices in TRACE to the quote prices in ICE. If there is no observation for a bond in a month in TRACE but there is one in ICE, then we treat it as the trade not happening in that month and still compute returns as described in Section 3.1. In the end, by combining TRACE and ICE, we calculate six types of net returns ($R^{b(v)h}$,$R^{hs(v)}$,$R^{b(v)s(v)}$,$R^{s(v)h}$,$R^{hb(v)}$,$R^{s(v)b(v)}$) and gross returns (R^{hh}).

Using both TRACE and ICE databases allows us to pin down the effect of half spreads and delays, but forces us to use a smaller sample to compute the net returns on the strategies than estimating the ML models. Focusing on this intersection between the two databases, we have 854,216 bond-month observations from August 2002 to November 2022 ($T = 244$).

Table 4 reports the summary statistics of the panel data for selected transaction sizes of \$10,000, \$100,000, \$1 million, and \$10 million. The average returns on the six types of net returns and gross returns are quite different from each other. For example, for the volume of \$100,000, the average returns for the long positions are -0.29%, 0.33%, and -0.31% for R^{bh}, R^{hs}, and R^{bs}, respectively. The average for short positions are higher and 0.15%, 0.77%, and 0.52% for R^{sh}, R^{hb}, and R^{sb}, respectively.[8] The difference among various net returns reflects bid-ask spreads and delays.

ICE provides bid quotes for all prices, not mid-prices. However, when we mark to market, we use these quotes. As a result, R^{bh}, where the investor pays an ask price and marks the bond at bids, tends to be low on average. In contrast, R^{hb} tends to be high because the position starts at a bid quote and ends at an ask, with R^{hs} and R^{sh} in between. The gap in average returns is more pronounced for small transactions, as their bid-ask spreads are larger.

To understand the sample across trade sizes, we plot the mean and median returns for by trade size in Panel A, Figure 3. As the volume threshold increases, R^{bh}, R^{bs}, R^{sh}, and R^{sb} converge to the risk-free rate because if there is no trade, investors do not initiate the position and earn the risk-free rate. In contrast, R^{hs} and R^{hb} converge to the mark-to-market return, R^{hh}, because delays prevent investors from unwinding the existing positions. Panel B shows that the percentage of the observations with no trade increases significantly with trade size. As size increases, the bid-ask spreads shrink, but the proportion of trades which are not executed increases. This explains the net return's convergence to either the gross return or the risk-free rate.

Table 5 reports the same statistics using duration-adjusted corporate bond returns, which subtract duration-matched Treasury returns from the corporate bond returns. We later study the model performance using these alternative measures of excess returns.

Figure 4 shows the distribution of trade sizes in the corporate bond market over time. Throughout the sample period, more than 50% of realized transactions are \$50,000 or less, and trades above \$1 million account for less than 20% of the number of trades. Interestingly, the share of small trades increases during the 2008 financial crisis, suggesting the increasing importance of adverse selection. In the post-Volcker periods, dealers use their inventory capacity less frequently and increase the share of pre-arranged trades (Bessembinder et al. 2018; Wu 2022), leading to the declining share of small trades in the 2020 pandemic crisis. This pattern confirms the increasing importance of trade delays.[9]

[8]The returns for short positions are higher as the trade starts from a bid price (as an investor sells) and concludes with an ask price (as an investor buys).

[9]In Internet Appendix D, we study institutional investors' quarterly position changes in eMAXX to confirm the validity of trade size estimates. In addition, in Figure 4 Panel B, we plot the distribution of insurance company trade sizes using NAIC data. As expected, their trade size is generally larger than that shown in Panel A, with a cross-sectional median of around \$1 million throughout the sample period. More than 10% of trades are above \$5

5 Performance of the Machine Learning Models

5.1 Estimating the Machine Learning Models

Following the notation in Gu, Kelly, and Xiu (2020), we describe a corporate bond's return in excess of T-bill rates as an additive prediction error model:

$$R_{i,t+1} = E_t(R_{i,t+1}) + \epsilon_{i,t+1}, \tag{7}$$

where,

$$E_t(R_{i,t+1}) = g^*(z_{i,t}). \tag{8}$$

Bonds are indexed as $i = 1, \ldots, N$ and months by $t = 1, \ldots, T$. Our objective is to isolate a representation of $E_t(R_{i,t+1})$ as a function of predictor variables that maximizes the out-of-sample explanatory power for realized $R_{i,t+1}$. We denote those predictors as the K-dimensional vector $z_{i,t}$, and assume the conditional expected return $g^*(\cdot)$ is a flexible function of these predictors. All of our model estimates minimize the mean squared prediction errors (MSE). In total, we consider six linear and non-linear machine learning models including penalized linear regression techniques: Lasso (LASSO), Ridge (RIDGE) and Elastic Net (ENET); non-linear regression tree ensemble methods including random forests (RF), and extreme trees (XT); and feed forward neural networks (NN). In addition, we form the linear ensemble model (LENS), the nonlinear ensemble model (NENS) and the ensemble across all models (ENS), which is the equally-weighted average across the respective models one-month ahead predictions (Rapach, Strauss, and Zhou, 2010).

For the first estimation as of July 2002, we source the last 55 months of data back to January 1998, and estimate the respective ML model. We measure excess returns at t and the 200-dimensional vector of bond characteristics at $t-1$. We perform cross-validation using a 70:30 training-validation split which preserves the temporal ordering of the panel data. We then use the vector of characteristics available at time t to produce a forecast of bond excess returns for $t+1$. These forecasts (expected returns) are available to the bond portfolio manager at time t, mean-

million.

ing they can trade on them at the end of the month. Thereafter, all models are re-trained every 12-months and cross-validated every 5-years with an expanding window.[10] We provide additional details related to the cross-validation and training of the respective models in Section B of the Internet Appendix.

5.2 Portfolio Performance Before Transaction Costs

Before considering the machine learning-based long-short portfolios, we pin down which anomaly characteristics are individually useful in forming profitable long-short bond portfolios. For each one of the 200 characteristics, we form quintile portfolios and initiate a long position in the fifth quintile and a short position in the first quintile. We use the ICE data and perform a preliminary analysis to create value-weighted quintiles to see if the long-short strategy has positive or negative bond CAPM alphas. We then sign the raw characteristics so that the long position has higher alphas than the short position.

We estimate the bond CAPM (CAPMB) alpha by running time series regressions of the strategies' returns on the corporate bond market factor:

$$R_t^e = \alpha + \beta MKTB_{Net,t} + \varepsilon_t, \tag{9}$$

where $MKTB_{Net,t}$ is the excess returns of BlackRock's corporate bond exchange-traded funds (ETFs), averaged between the investment-grade ETF (Ticker: LQD) and the high-yield ETF (Ticker: HYG) using the total market value of corporate bonds in each respective rating category as the weights. We use the ETF returns because they reflect the cost of buying and holding the bond market portfolios. Therefore, ETF returns provide a fair benchmark to evaluate the performance of trading strategies net of costs. The detailed construction method of the market factor is provided in Appendix B. We find that the average excess returns on our ETF-based market factor is 0.32% per month, while the corresponding value for the value-weighted market bond portfolio of Dickerson et al. (2023) is 0.36% over the same period. The lower value of the ETF returns suggests that even

[10]This gives the models an advantage in that they are re-trained and re-cross-validated multiple times over our sample.

holding the market portfolio is somewhat costly for investors. To account for autocorrelation in the returns, we adjust the standard errors using Newey and West (1987) 12 lags.

We first examine the CAPMB alphas before transaction costs, shown in Table 6. In Panel A, we sort the characteristics by the information ratio, which is the ratio of alpha to the standard deviation of the residual term in (9) multiplied by the square root of 12. We report the top 15 characteristics in the table and relegate the rest to the Internet Appendix Table A.1.

We find that quite a few characteristics generate significant CAPMB alphas before transaction costs. In terms of the information ratio, the best performers are the past one-month stock returns (strev) of Chordia et al. (2017), stock returns on earnings announcement days (pead) of Nozawa et al. (2023), and 6-month increase in credit spreads (mom6mspread)[11] of Kelly et al. (2021), which have ratios of 1.38, 1.23, and 0.88, respectively. Most of the significant characteristics are based on past returns for bonds, stocks, and options.

Next, we combine the information in each signal and examine the performance of the machine learning algorithms. Each month, we sort corporate bonds into value-weighted quintiles based on the month-end expected returns generated by the machine learning algorithms. We then take a long position on the top quintile and a short position on the bottom quintile, calculate the excess return of the long-short strategy, and estimate the CAPMB alphas.

The second and fourth columns of Table 7 report the average excess returns and CAPMB alphas of the machine learning-based strategies. We find that the ML algorithms are successful in combining signals into return predictions that work well out of sample. All nine strategies generate significant CAPMB alphas, led by Ridge (0.50%, t=3.77), the average of linear models (LENS; 0.48%, t=3.18), and the average of all models (ENS; 0.48%, t=2.90). They also exhibit a high annualized information ratio, ranging from 0.60 to 1.37. These values are economically significant given that the Sharpe ratio for the bond market portfolio over the same period is 0.65 before transaction costs.

Binsbergen et al. (2023) find that adjusting for corporate bond duration significantly affects

[11] Kelly et al. (2021) report that mom6mspread, defined as a decrease in credit spreads over the previous 6 months, negatively predicts bond returns. We find the same as they do and thus reverse the sign of the characteristic to represent an increase in credit spreads. In effect, the results show that corporate bond returns exhibit a return reversal rather than momentum.

the test of asset pricing models. Thus, we replace corporate bond returns with the difference between corporate bond returns and duration-matched Treasury returns (computed by ICE) and compute the duration-adjusted returns of the strategies. Table 8 shows the average duration-adjusted returns and the CAPMB alphas. Using duration-adjusted returns slightly improves the performance of the model after adjusting for market exposure. The information ratio ranges from 0.83 to 1.38, confirming the value of combining multiple signals to generate reliable return forecasts.

5.3 Impact of Transaction Costs

In this section, we evaluate the impact of transaction costs on the performance of the ML strategies. Although transaction costs depend on trade size, we first focus on the optimal value that maximizes the net CAPMB alpha.

The ninth column of Table 7 reports the optimal trade size for each ML algorithm. We find that the optimal size is mostly \$2 million except for NN (\$5 million), LASSO (\$5 million), and RF (\$10 million). These values are larger than the typical "institutional trade" size of \$100,000 used in the literature (e.g., Bessembinder et al. 2008), reflecting the lower bid-ask spreads for larger trades.

The third and fifth columns of Table 7 show the net returns and alphas at the optimal trade size. At the optimum, the CAPMB alpha ranges from -0.03% to 0.12%. After costs, none of the algorithms produce economically and statistically significant net alphas. Figure 5 visualizes these estimates. The panels on the left, which show the performance before costs, show a striking contrast with those in the middle, which show the performance after costs: The average returns and alphas are well below the counterparts before costs, suggesting that it is a challenge for the ML algorithm to perform well net of costs. Using duration-adjusted returns, reported in Table 8, leads to a similar conclusion.

The right-most panels of Figure 5 show the net returns and alphas for the long-short strategies based on ML algorithms with the \$100,000 threshold. Clearly, the net alphas are negative for all strategies, and seven of them have significantly negative net alphas. The difference between the middle and right panels of Figure 5 suggests that the choice of trade size is crucial in evaluating the performance of trading strategies.

Transaction costs are also significant for long-short portfolios based on individual signals. In Table 6 Panel B, we sort the characteristics by the net information ratio with the optimal trade size. When we cannot not find the optimum because the gross alpha is low, we use \$2 million as the trade size. Now, the rankings change dramatically from Panel A because characteristics with high gross alphas tend to have high turnover and their net alphas are much lower. For example, the net information ratio for the previous one-month stock return (strev) is 0.34, down from 1.38 before costs.

Figure 6 plots the number of significant characteristics before and after transaction costs. Transaction costs reduce the number of significant variables significantly, from 26 to 7 before adjusting for the CAPMB (the left panel), and from 20 to 4 after adjusting for the CAPMB (the right panel). Thus, without transaction cost adjustments, one reaches dramatically different conclusions as to the profitability of factor investing in the bond market.

The key to the above results is the inclusion of delay costs in the calculation of transaction costs. Without it, the half-spreads shrink to zero as trade size increases, and we would incorrectly conclude that the ML algorithm generates profitable strategies after costs. To illustrate the key mechanism, Figure 7 plots the CAPMB alphas of the long-short strategies before and after transaction costs as a function of trade size. For example, the left panel of Panel A plots gross and net CAPMB alphas using Neural Networks. Before costs, this signal generates an alpha of 0.43%. The net alphas, on the other hand, are a hump-shaped function of the target trading volume, with a maximum at \$2 million.

We decompose the difference between gross and net alphas into the component explained by half spreads and the component explained by delays. To do this, we compute alternative net returns using ICE's quotes on the transaction dates provided by TRACE. Specifically, for each bond in each month, we compute hypothetical net returns by replacing $P_{t,d}^{b(v)}$ and $P_{t,d}^{s(v)}$ with $P_{t,d}^{h}$ in Equations (4), (5), and (6) as well as the corresponding returns for short positions. Thus, this net return reflects the cost of delays but not half spreads. We then compute the delay-only net returns on portfolios using these hypothetical returns and the associated CAPM alpha, denoted $\alpha^{NetDelay}$. The difference between the gross alphas and delay-only alphas gives us the pure effect of delays, and the remainder is accounted for by the half spreads between ICE quotes and TRACE transaction prices on the same day. More formally, the total cost is decomposed into:

$$\alpha^{Gross} - \alpha^{Net} = \underbrace{\alpha^{Gross} - \alpha^{NetDelay}}_{\text{=Delay Cost}} + \underbrace{\alpha^{NetDelay} - \alpha^{Net}}_{\text{=Half-Spread Cost}} \tag{10}$$

The right panels of Figure 7 show the decomposition of costs. Continuing with the Neural Network example, the effect of half-spreads falls from 1.2% at the volume of $0 to near zero at the maximum trade size of $20 million, reflecting the standard spread-volume relationship. Note that our half-spread cost takes into account the wedge between transaction prices and quotes as well as portfolio turnover. For example, if the price wedge is 1% and portfolio turnover is 30%, then our half-spread cost is approximately 0.3%.

On the other hand, the delay effect increases as the target volume increases from near zero at the $0 volume threshold to 0.40% at the $20 million threshold, reflecting the cost of missing trading opportunities. As a result, the sum of the two costs exhibits a U-shaped pattern with respect to volume. As trade size increases beyond $2 million, the increase in delay costs dominates the decrease in half-spread costs. Thus, it is impossible to argue that ML strategies provide profitable trading opportunities when investors trade with very large volumes. We observe the U-shaped transaction costs for other ML strategies in other panels of Figure 7.

Table 9 reports the decomposition of the trading costs for a trade size of $100,000 and the optimal size for each ML strategy. With the transaction size of $100,000, the cost due to half spread ranges from 0.51% to 0.62% while the delay cost ranges from 0.06% to 0.11%. With the optimal trade size, the half-spread cost is lower, ranging from 0.06% to 0.22%, reflecting the cost savings for large transactions. On the other hand, the delay cost is now higher, ranging from 0.16% to 0.25%. This pattern highlights the key trade-off between half spreads and delays.

Next, we compare our cost estimates with those provided in other papers. In the standard setup, the trading costs and net returns for portfolio p are calculated by

$$c_{p,t+1} = \sum_{i \in N_t} \left| w_{i,p,t+1} - \frac{1 + R^c_{i,t+1}}{1 + R^c_{p,t+1}} w_{i,p,t} \right| \mathbb{s}_{i,t+1}, \tag{11}$$

$$R^{Net}_{p,t+1} = R^{Gross}_{p,t+1} + \mathbb{1}_p c_{p,t+1} \tag{12}$$

where R^c is a clean price return, $\mathbb{s}_{i,t+1}$ is a half spread for bond i in month $t+1$, $w_{i,p,t}$ is its weight in portfolio p, $\mathbb{1}_p$ is an indicator function that equals one if p is a short position and minus one if p is a long position. Essentially, the position changes are derived from changes in portfolio weights that are not due to the bond's or portfolio's returns, and the transaction costs are the product of the position change and an estimate of the half spread. For example, Kelly et al. (2021) use the constant spread by setting $\mathbb{s}_{i,t+1} = 0.19\%$ for all i and t.[12]

The last column in Table 9 reports this standard calculation of transaction costs. To be specific, we compute the CAPMB alphas using the gross and net returns in (12) and report the difference. We find that this alternative method underestimates the transaction costs when compared to the costs with optimal trade size. For example, for the ENS strategy, the alternative method leads to a cost estimate of 0.23%, which is lower than the total cost at the optimal trade size of 0.41%. This alternative cost is close to the bid-ask spread component of the total cost (0.22%) and thus Kelly et al. (2021)'s method correctly captures the cost of half spreads. However, they do not account for the cost of delays in implementing the strategy, which explains the substantial gap between their cost estimates and our approach.

The dashed line in the left column of Figure 7 shows the net alphas based on the alternative cost adjustments. In all panels, our net alphas are always below this alternative, further confirming the significance of the cost of delays.

One might ask whether it is realistic to always trade in the fixed dollar amount or whether the optimal trade size is constant over time. To take a first look at the importance of time-varying trade size, we split our sample in half, one period from August 2002 to December 2012 and the

[12]Kelly et al. (2021) use the lagged weight change $|w_{i,t} - w_{i,t-1}|$ times 19 bps instead of (11), which does not appear to account for organic changes in portfolio weights due to bond and portfolio returns. This procedure slightly overestimates cost and turnover, as all position changes are considered sales and purchases of a bond.

other from January 2013 to November 2022. Figure 8 shows the cost decomposition for the ENS strategy for these two subperiods. In this case, the optimal trade size remains unchanged at $2 million. This is because two forces cancel each other out. On the one hand, the lower average gross returns in the second period make it optimal to wait longer, thus increasing the optimal trade size. On the other hand, the lower bid-ask spreads in the second period make it less costly to trade a small quantity. As a result, investors would not benefit from changing the target trade size between these two periods.

In the Internet Appendix Section F, we investigate the prospect of dividing large target trade sizes into smaller portions to determine if this strategy can help minimize transaction costs. Our findings indicate that splitting a large trade does not reduce trading costs because it increases the bid-ask spreads, offsetting any benefits from faster execution.

5.4 Are Our Cost Estimates Biased?

We compute the cost of delay by assuming that if we do not observe trades of size at least v in a month on TRACE, the investor's order remains unfilled. Although this estimate shows the actual return net of cost for a real-time investor, it still reflects an equilibrium outcome. The interpretation of our cost estimates varies depending on how one views off-equilibrium trades. Our perspective is that liquidity supply (i.e., the need for dealers to pre-arrange a round trip) is the key driver for the observed trade size and frequency. However, an alternative viewpoint is that liquidity demand determines them.

For example, one could argue that very large investors always enjoy near-zero half spreads but choose not to trade often because they do not need to. According to this argument, the rarity of large trades in TRACE reflects the lack of liquidity demand rather than liquidity supply. However, the evidence in the prior literature favors the liquidity supply-based explanation. Goldstein and Hotchkiss (2020) find that for bonds rated BBB and below, large client orders tend to have a shorter time in dealer inventory.[13] This provides direct evidence that dealers pre-arrange trades for larger trades to avoid inventory risk. If they do not do this, then large trades should remain in

[13]See their Table 7.

their inventory longer, not shorter, than small trades. Consistent with this finding, Kargar et al. (2023) presents direct evidence that customers experience delays using the order-level data on the electronic trading platform.[14]

Furthermore, the explanation based on liquidity demand suggests that large investors have a clear advantage in terms of costs compared to small investors, which leads to their outperformance and faster growth. However, our analysis of corporate bond mutual funds in Internet Appendix Figure A.3 and Table A.2 shows no evidence supporting this idea.

If a limit on liquidity supply is the reason for the rarity of large trades, then our cost estimates are the lower bound of true costs, as deviating from the observed equilibrium would increase the cost even more.

5.5 Time Variation in Transaction Costs

Unlike Kelly et al. (2021), we do not assume that half spreads are constant over time. Since we back out transaction costs from differences between gross and net returns, the costs vary every month. This allows us to study the time-series behavior of costs and check whether the recent technological development in corporate bond trading (as pointed out by O'Hara and Zhou 2021) mitigates the cost of delays.

Figure 7 plots the transaction costs due to half spreads and delays for the ENS strategy as an example. We use the optimal trade size of $2 million. Since the cost of delay is volatile, we also plot its twelve-month moving averages.[15]

The plot reveals a familiar counter-cyclical pattern in transaction costs for both components. Despite the introduction of electronic trading platforms, we do not see a clear decline in transaction costs in the most recent sample, possibly because the market share of transactions through these new platforms is still small.

[14] O'Hara and Zhou (2021) show that the dependence of transaction costs on trade size is weaker on the electronic platform. However, in Kargar et al. (2023), the market share of the electronic platform is less than 20% of the total volume and lower for bonds with high transaction costs: i.e., those with large size and higher credit risk.

[15] The delay cost is volatile because the cost of delay can be positive or negative, depending on the performance of the signal in the month. For example, if the signal mispredicts a return in a month and the strategy's gross return is negative, then delaying the execution to start the trade can be beneficial, leading to a negative delay cost. If the signal produces a profitable strategy on average, then the cost of delay is positive on average. This does not mean that the cost of delay is positive every month.

5.6 Determinants of Optimal Trade Size

Across the 200 characteristics and 9 ML strategies, there is a significant variation in the optimal trade size. In this section, we take advantage of the observed difference across strategies and investigate the determinants of optimal trade size. As we have seen in the previous section, the optimal size depends on how profitable the strategy is, as measured by gross alpha. In addition, it may depend on how frequently an investor must trade, which is measured by the average turnover rate of strategy s:

$$Turn_p = \frac{1}{T}\sum_{t=1}^{T}\sum_{i\in N_t}\left| w_{i,p,t+1} - \frac{1+R^c_{i,t+1}}{1+R^c_{p,t+1}} w_{i,p,t}\right| \quad \text{where } p \in \{long, short\}, \tag{13}$$

$$Turn_s = 0.5(Turn_{long} + Turn_{short}). \tag{14}$$

In this exercise, we use the turnover rate when the minimum trade size $v = 0$ so that it captures the persistence of the signal. A higher size would artificially reduce turnover due to implementation delays.

To describe optimal v, we classify 114 strategies with optimal trade size into three categories based independently on their turnover rate and gross CAPM alpha, creating nine bins. From a total of 209 candidates, we remove 95 strategies for which we cannot find the optimal trade size.[16] The cutoff values for alphas are 0.1% and 0.2%, while those for turnover rate are 10% and 19%. For each of the nine categories, we compute the average across strategies within a bin for optimal trading volume, total transaction cost (at optimal volume), half-spread cost, and delay cost.

Table 10 shows the averages for nine bins. The value for the high gross alpha/low turnover bin is missing because no strategy falls into this category. In Panel A, we report the average of the optimal trade size. For medium and high turnover strategies, the optimal trade size is decreasing in gross alpha. For the medium turnover category, the optimal trade size is \$6.1 million for low alpha strategies and \$2.6 million for high alpha strategies. This is because when the signal is profitable, it

[16] In Internet Appendix Figure A.4, we plot the net alpha of bond age as an example of the characteristics in which we do not find an optimal value. Intuitively, when the gross alpha is so low, it is better not to trade on the signal at all, and it is optimal to choose the infinitely large trade size. This leads to the failure to find an interior solution.

is better to execute trades as soon as possible and avoid missing the trading opportunity generated by the signal. On the other hand, the pattern for turnover rate is less clear.

Turning to half spread costs (Panel C), they increase in gross alpha. This is because a highly profitable signal optimally sets a small trade size, leading to higher bid-ask spreads. For the middle and high gross alpha strategies, they also increase in turnover rate. The positive relationship between half spread cost and turnover rate is somewhat mechanical, as the cost increases as investors trade more frequently. (For the low alpha strategies, the optimal trade size increases significantly as we increase the turnover rate, offsetting the increase in costs.)

Delay costs are increasing in gross alpha. When gross alpha is high, it is costly to miss a trading opportunity, and thus delay costs are high. This cost is mitigated by the fact that trading volume is optimally chosen to reduce delay when gross alpha is high. However, when the turnover rate is also high (i.e., the signal is moving quickly), this cost mitigation is not as effective, resulting in the high delay cost. For example, for the bin with the highest gross alpha and turnover rate, the delay cost is 0.15% on average, while the delay cost is close to zero for the strategies with a low alpha. The total cost, shown in Panel B, is the sum of the half-spread cost and the delay cost.

5.7 How Much Gross Alpha Do We Need?

We have emphasized the role of transaction costs in evaluating the performance of trading strategies. In this section, we provide a guide for future research that explores the new signals that predict corporate bond returns. The goal of this exercise is to present the target level of gross alpha that achieves the desired level of net alpha under the assumption that trade size is optimally chosen. This will allow other researchers to calculate the gross alpha of their strategies and quickly check whether they also generate a positive alpha net of costs.

Since the relationship between gross and net alpha is affected by the persistence of the signal, we use the 114 strategies and estimate multivariate regressions of net alpha and the associated

t-statistics on the turnover rate and gross alpha:

$$\alpha_{Net,s} = -0.079 + 0.040 \log Turn_s + 1.099\alpha_{Gross,s} - 0.254 \log Turn_s \times \alpha_{Gross,s} + \varepsilon_s, \quad (15)$$

$$t(\alpha_{Net,s}) = -0.220 + 0.179 \log Turn_s + 13.176\alpha_{Gross,s} - 2.976 \log Turn_s \times \alpha_{Gross,s} + \varepsilon_s. \quad (16)$$

Let $\hat{\alpha}_{Net}(Turn, \alpha_{Gross})$, $t(\hat{\alpha}_{Net})(Turn, \alpha_{Gross})$ be the fitted value of the regression evaluated at ($Turn$, α_{Gross}). Then, we plot the combination of $Turn$ and α_{Gross} that satisfies the minimum level of α_{Net} or $t(\alpha_{Net})$.

The left panel of Figure 10 shows the combination of $Turn$ and α_{Gross} required to achieve net alpha of 0%, 0.1%, and 0.2% per month. In the figure, a strategy in the northwest region of the graph generates higher net alpha, while a strategy in the southeast region generates lower net alpha. The dashed line is the break-even point needed to match the passive ETF returns. To outperform the ETF by a modest 0.1% per month, a strategy must be above the dotted line. As in the figure, the cutoff is strongly increasing in turnover rate. To achieve a monthly alpha of 0.1%, it is essential to keep the turnover rate below 70%.

The right panel plots the bound using the t statistic of the net alpha. To achieve the t value above the 10% level, a strategy must be above the dashed line, while the 5% significance requires being above the dotted line. The strongly upward-sloping curve in this figure reinforces the message that for a strategy to work, it is essential to keep the turnover rate low.

Achieving statistically significant net CAPM alpha requires relatively high values of gross alpha. For example, consider a hypothetical strategy with a turnover rate of 20%. Then, it must generate gross alpha of 0.31% and 0.39% to achieve a t-statistic of 1.65 and 1.96, respectively. If the turnover rate is 30%, the corresponding required gross alpha's are 0.41% and 0.52%, respectively. If a strategy's turnover is higher, then it requires higher levels of gross alpha to be useful in practice. For future reference, we tabulate the net CAPM alpha as a function of gross alpha and turnover rate in Table A4 in Appendix.

5.8 Robustness

In this section, we perform several robustness checks. First, we consider long-only strategies instead of the long-short strategies used in the main analysis. This analysis is important because shorting corporate bonds can be quite costly for some investors although Asquith et al. (2013) show that the cost of borrowing corporate bonds is comparable to that of stocks.

Using the expected returns generated by the ML models, we take a long position in the top 20% bonds and calculate their gross and net returns over T-bill rates. Table 11 reports the performance of the long-only strategies. We find that the average gross and net excess returns are higher than those of the long-short strategies in Table 7. This is to be expected because corporate bond returns are generally higher than T-bill rates. Once we account for market risk, the CAPM alphas of the long-only strategies are similar to those of the long-short strategies. For example, using ENS, the gross and net alphas for the long-short strategy are 0.48% and 0.07%, respectively, while the corresponding values for the long-only strategy are 0.25% and 0.09%. Eight of the nine long-only strategies fail to generate significant alphas after transaction costs even at the optimal trade size with an exception of Ridge, which is marginally significant.

Second, we consider a potential boost for the ML strategies by selecting a smaller number of bonds for the long and short positions instead of buying and selling all the bonds in the top and bottom 20%. In our main results, the average number of bonds in the long and short positions for the ENS strategy is 905 and 776, respectively. In practice, investors may sample a smaller number of bonds for ease of implementation. For example, Choi, Cremers, and Riley (2023a) find that in their sample of actively managed corporate bond mutual funds, the average number of bonds held is 533. In line with this approach, we construct long-short strategies by selecting the top 2% and bottom 2% of the cross-section of corporate bonds, thereby narrowing the number of bonds in each position.

Table 12 shows the gross and net returns/alphas for each ML strategy. By construction, the number of bonds in the long and short positions is smaller. Since we are using the bonds with extreme signal values, the average gross excess returns are higher than the main results in Table 7. For example, for ENS, the gross alpha is now 1.21% (t=3.75), higher than the main result using quintile portfolios (0.48%, t=2.90). The benefit of higher average returns is partly offset by the cost of higher volatility, which attenuates the statistical significance (and the Sharpe ratio). In addition, the monthly turnover of all strategies increases to 69% from 49% in the main results. The high turnover inflates the transaction costs, which leads to a barely significant net CAPM alpha for ENS (0.45%, t=1.96). Looking across strategies, four of the nine ML strategies now have significant net alphas. Nonetheless, the main message of the paper remains unchanged: accounting for transaction costs significantly reduces the strategy's alpha. Even for the best performer (NENS), the net information ratio of 0.64 is less than half of the gross ratio of 1.31.

Lastly, we consider a strategy to reduce the cost of delay by trading a subsample of liquid bonds. Following Goldstein and Hotchkiss (2020), we classify bonds into two categories based on the trade counts and the number of non-zero trade days in month t, when we form portfolios. Then, within a group of liquid bonds (i.e., those with above-median trade frequency), we form a quintile portfolios based on the signals generated from ML models.

Table 13 reports the performance of the ML strategies using liquid subsample of bonds. The results are very similar to the main results. Because we use liquid bonds, the strategies' turnover rate increases from the main results: For example, for ENS, it is now about 60%, higher than the main results of 49%. However, the resulting net CAPMB alpha is only 0.11% using bonds with high trade counts (Panel A) and 0.09% using bonds with high trade days (Panel B). None of the nine strategies generate significant net CAPMB alphas. Therefore, using a subsample of liquid bonds does not significantly reduce the transaction costs.

6 Do Corporate Bond Mutual Funds 'Beat the Market'?

In order to identify whether our net of cost machine learning strategy returns are 'underperforming' what is being achieved in reality, we investigate the performance of corporate bond mutual funds over an identical sample period. If many bond mutual funds are indeed beating the market across a wide variety of fund styles, it would indicate that our strategies could be refined. However, if most funds underperform relative to a simple passive benchmark, it would corroborate our main findings showing that generating net of cost alpha is an immensely challenging task.

6.1 Distribution of Corporate Bond Mutual Fund Alphas

We examine the CRSP mutual fund database. The sample is from July 2002 to November 2022 where the start and end date is set to be the same as our main results. We identify corporate bond mutual funds by CRSP's fund classification. In particular, we choose the subcategory 'Corporate' among 'Fixed Income' funds. Funds with less than 36 monthly observations and total net assets (TNA) less than $10 million are removed from the sample. We also remove all funds that track an index or are passively managed, i.e., we focus on actively managed bond funds.

After filtering, we are left with a sample of 485 mutual funds that invest in corporate bonds. To pin down which funds exhibit alpha, for each fund we estimate a single-factor model of each fund's net return in excess of the one-month risk-free rate on the $MKTB_{Net}$ factor.

We present summary statistics in Table 14. Panel A presents fund summary statistics and B reports cross-sectional fund performance statistics. On average a representative mutual fund remains in the sample for 109 months, with TNA of US$ 575 million, average annual expense ratios of 0.90% and annual turnover of 119%. In the cross-section, gross (net) fund alphas are 0.05% (0.03%) on average, with an average gross (net) return of 0.33% (0.26%) per month. The passive net of costs bond market factor explains over 70% of the time-series variation of fund returns with a beta of close to 0.70.

We present the distribution of the funds monthly net alphas and associated t-statistics in Panels A and B of Figure 11. Of the 485 mutual funds we consider, only 42 of them (8.65% of the sample) generate risk-adjusted net returns relative to the passive net of costs $MKTB_{Net}$ benchmark. The average net alpha of these funds is economically small at 0.18% per month. Of the 42 funds that do generate alpha, 33% invest in investment grade bonds with higher yields (bonds rated closer to BBB-) and 17% invest in high rated investment grade bonds.[17] Noninvestment grade bond funds do not generate any alpha. For active corporate bond mutual fund investors, these preliminary results are somewhat discouraging, but are supportive of our findings related to the poor net of cost alphas

[17]The remaining funds cannot be classified due to a lack of information.

generated by our machine learning strategies. Relative to the average gross alpha generated by the mutual funds (0.05%), the machine learning based portfolios perform admirably (average alpha across the strategies is 0.20%).

Very few mutual funds offer incremental risk-adjusted performance in excess of simply holding the passive net of cost bond market portfolio. Of funds that do outperform, the economic magnitude of the outperformance is small. What is perhaps more disheartening, is that 37% of the funds (over one third of the sample) generate net of cost alphas that are less than zero. The distribution of the corporate bond mutual fund alphas is not unsurprising given that active portfolio management is considered a zero (or negative) sum game (Fama and French 2010 and Sharpe 1991). If some active bond bonds generate alpha, it comes at the expense of other bond funds. However, relying on alpha as a measurement of skill can be misleading. We now turn to identifying whether active bond mutual funds 'add value' through skillful management to further corroborate our findings that outperforming a simple passive bond benchmark is a tall order.

6.2 Skill and Manager Value Added

Berk and Binsbergen (2015) show that gross alpha does not measure mutual fund manager skill, and it also need not be positively correlated with skill. We examine a proxy for skill which directly measures the ability of the fund manager to extract money from the markets. To do this we compute the value that the fund offers to an investor over and above a gross return passive benchmark. Following Berk and Binsbergen (2015), we measure a funds added value ($V_{i,t}$) by multiplying the benchmark adjusted realized gross mutual fund return, $R^g_{i,t} - R^g_{MKTB,t}$, by the real size of the fund (assets under management scaled by inflation) at the end of the previous month,

$$V_{i,t} = TNA_{i,t-1} \cdot (R^g_{i,t} - R^g_{MKTB,t}),$$

where $TNA_{i,t-1}$ is the total assets of the fund in the prior month, $R^g_{i,t}$ is the gross return of fund i in month t computed as the funds net return plus the monthly management fee, and $R^g_{MKTB,t}$ is the gross return of the bond market factor. In this equation, $V_{i,t}$ represents the monthly 'value-add' from fund manager skill in US$ millions.

The measure of 'skill', S_i for each fund is the time-series average of each funds value-add. We then compute the cross-sectional average of S_i, using (i) equal-weights (the weights are equal for each fund in the cross-section), (ii) time-weights (the weights are the number of months each fund is present in the sample) and (iii) expense ratio weights (the weights are the average fund expense ratios).

We report the respective cross-sectional averages of V_i in Panel A Table 15 and cross-sectional percentiles in Panel B. Strikingly, the equally-weighted average monthly value-add of a given fund is negative \$396,000 per month, or negative \$4.75 million annually. This value is economically large in absolute value, and highly statistically significant at the 1% nominal level. The time weighted and expense weighted estimates are similar in magnitude (negative) and also statistically significant. In contrast to results presented in Berk and Binsbergen (2015), as opposed to adding value on average, we show that bond mutual fund managers are value extractors, implying active bond investors are paying for relatively adverse performance with respect to the passive benchmark. In Panel B, the variation in value-add is large. Bond funds at the 1st (99th) percentile generated a negative (positive) value-add of \$9.12 (\$2.63) million per month. The median fund lost investors an average of \$60,000 per month relative to the passive benchmark, and only 75% of the mutual funds we consider generated a positive value-add.

6.3 Luck vs. Skill?

Given that a few mutual funds do generate net of costs alpha, we follow the methodology of Barras, Scaillet, and Wermers (2010) to partition the proportion of funds that exhibit significant alphas by luck and skill. We first estimate mutual fund alphas and their associated p-values individually, using net of fee returns and the $MKTB_{Net}$ passive benchmark portfolio. Funds can be classified as either 'Unskilled', implying they have a net alpha shortfall ($\alpha < 0$), 'Zero-alpha', which means managers have enough skill which is just sufficient to recover trading costs ($\alpha = 0$), and 'Skilled funds' meaning managers are skilled enough to generate an alpha surplus after costs ($\alpha > 0$). Given we cannot observe the true alphas of each fund in the population, we infer the prevalence of each of the above skill groups by using the false discovery rate (FDR) as a methodology for separating skill from luck (See Benjamini and Hochberg 1995 and Barras et al. 2010 for the estimation details).

We present the results in Panel A and B of Table 16. Of the 485 corporate bond specific mutual funds we consider, 76.45% (371 funds) are estimated to be zero-alpha funds.[18] This implies that, confirming prior results in the literature, the majority of the funds we consider are run by managers with enough ability to generate a net alpha that roughly covers their management fees. In other words, the economic rents extracted from these managers from their clients are about enough to cover their fees and trading costs. Funds that generate a non-zero alpha amount to 23.55% of our mutual fund sample (114 funds). Of these funds, and in contrast to results for *all* mutual funds as in Barras, Scaillet, and Wermers (2010) and others, only 8.07% of these funds are truly unskilled with a true alpha less than zero. Skilled funds with true alpha greater than zero comprise 15.48% of the proportion of non-zero alpha funds.[19] In Panel B, we present the proportion of the significant alphas in the left and right tails of the distribution (denoted as $\widehat{S}_{\gamma}^{-}$ and $\widehat{S}_{\gamma}^{+}$) at four significance levels ($\gamma = 0.05, 0.10, 0.15, 0.20$). Focusing first on the right tail, when $\gamma = 0.20$, 14.02% (68) funds generate a positive alpha with a two-sided p-value below 20%. However, of these funds, more than half (37) of the funds are merely lucky, i.e., the positive alpha is not due to manager skill in a statistical sense. As we decrease the level of γ (increase the level of significance), this phenomenon reverses, i.e., fund alphas that have a greater degree of statistical significance are earned by a greater proportion of skilled managers. The proportion of corporate bond mutual fund managers who generate statistically significant alpha in the right tail at the 5% nominal level is 6.39% (31 funds). Of these managers, 4.48% (1.91%) are skilled (lucky). Unfortunately (for active bond mutual fund investors), this result broadly confirms those presented in the prior section on value. Of the 31 funds that generate positive alpha, only 22 funds (out of 485) generate the positive alpha through skillful management. Only a tiny fraction of very top performing mutual funds appear to

[18]Given the critique of the FDR method when applied to mutual funds by Andrikogiannopoulou and Papakonstantinou (2019), our results are robust to changing the FDR parameters which generates the 'Zero-alpha' fund percentage.

[19]This is in contrast to estimates from the CRSP Mutual Fund database that uses *all* funds. In Barras et al. (2010) the percentage of skilled funds is estimated to be 0.60% (statistically indifferent to zero).

outperform a passive bond market ETF net of costs.

Overall, when synthesizing the results from both of the methods we use to analyze corporate bond mutual fund returns, two salient results are worth emphasizing. First, a representative investor is, on average, better off simply purchasing a portfolio of low cost, passive bond market ETFs. Second, the probability of selecting an active bond portfolio manager who is able to generate statistically significant net of fees alpha through skill is extremely unlikely.

7 Conclusion

In this paper, we present delayed trade execution as a key cost in the evaluation of trading strategies using illiquid assets. When transactions are infrequent, the standard portfolio approach of Fama and French (1992) no longer provides a realistic performance benchmark for trading strategies, even after adjusting for bid-ask spreads. The cost of missing trading opportunities is particularly severe when the signal contains valuable information and moves quickly.

In our framework, investors face a trade-off between tighter bid-ask spreads and execution speeds. As a result, total transaction costs are a U-shaped function of trade size, as opposed to the monotonically decreasing function described in Edwards et al. (2007). This allows us to identify an optimal trade size and ties our hands in selecting a trade size for net return calculations. We show that the optimal size decreases as the gross alpha of the strategy increases.

Our methodology applies to a broader set of illiquid assets other than corporate bonds. The key is to find a proxy for the bid-ask spreads on which investors can condition their orders. In the stock market, the relationship between trade size and bid-ask spreads is positive. However, the basic tension remains: the trading opportunity at tight bid-ask spreads is limited, and thus one has to wait longer for order execution if one insists on a tight spread. In the corporate bond market, trade size is negatively correlated with bid-ask spreads and serves as an excellent proxy for trading opportunity, but we can use different proxies in different markets.

To underscore the importance of delay costs, we estimate the ML models to generate out-of-sample forecasts of corporate bond returns. Consistent with previous research, the long-short strategy based on these forecasts generates significant CAPM alphas before transaction costs. However, after adjusting for transaction costs and trading delays, the net alphas are essentially zero.

We confirm the difficulty of developing profitable strategies after transaction costs by examining the returns of corporate bond mutual funds. Consistent with the unimpressive performance of ML strategies, most corporate bond mutual funds have insignificant alphas relative to passive net of cost ETF returns. Taken together, these results suggest that generating factor investing strategies in corporate bonds is a challenge for researchers and practitioners alike.

References

Andrikogiannopoulou, Angie, and Filippos Papakonstantinou, 2019, Reassessing false discoveries in mutual fund performance: Skill, luck, or lack of power?, *Journal of Finance* 74, 2667–2688.

Asquith, Paul, Andrea S. Au, Thomas Covert, and Parag A. Pathak, 2013, The market for borrowing corporate bonds, *Journal of Financial Economics* 107, 155–182.

Avramov, Doron, Si Cheng, and Lior Metzker, 2023, Machine learning vs. economic restrictions: Evidence from stock return predictability, *Management Science* 69, 2587–2619.

Bai, Jennie, Turan G Bali, and Quan Wen, 2019, RETRACTED: Common risk factors in the cross-section of corporate bond returns, *Journal of Financial Economics* 131, 619–642.

Bai, Jennie, Turan G. Bali, and Quan Wen, 2021, Is there a risk-return tradeoff in the corporate bond market? time-series and cross-sectional evidence, *Journal of Financial Economics* 142, 1017–1037.

Bali, Turan G., Heiner Beckmeyer, and Amit Goyal, 2023, A joint factor model for bonds, stocks, and options, Working Paper.

Bali, Turan G, Amit Goyal, Dashan Huang, Fuwei Jiang, and Quan Wen, 2020, Predicting corporate bond returns: Merton meets machine learning, *Georgetown McDonough School of Business Research Paper* 20–110.

Bali, Turan G, Avanidhar Subrahmanyam, and Quan Wen, 2021a, Long-term reversals in the corporate bond market, *Journal of Financial Economics* 139, 656–677.

Bali, Turan G., Avanidhar Subrahmanyam, and Quan Wen, 2021b, The macroeconomic uncertainty premium in the corporate bond market, *Journal of Financial and Quantitative Analysis* 56, 1653–1678.

Bao, Jack, Maureen O'Hara, and Xing (Alex) Zhou, 2018, The volcker rule and corporate bond market making in times of stress, *Journal of Financial Economics* 130, 95–113.

Bao, Jack, Jun Pan, and Jiang Wang, 2011, The illiquidity of corporate bonds, *Journal of Finance* 66, 911–946.

Barras, Laurent, Olivier Scaillet, and Russ Wermers, 2010, False discoveries in mutual fund performance: Measuring luck in estimated alphas, *Journal of Finance* 65, 179–216.

Bartram, Söhnke M., Mark Grinblatt, and Yoshio Nozawa, 2023, Book-to-market, mispricing, and the cross-section of corporate bond returns, *Journal of Financial and Quantitative Analysis* forthcoming.

Benjamini, Yoav, and Yosef Hochberg, 1995, Controlling the false discovery rate: a practical and powerful approach to multiple testing, *Journal of the Royal statistical society: series B (Methodological)* 57, 289–300.

Berk, Jonathan B., and Jules H. van Binsbergen, 2015, Measuring skill in the mutual fund industry, *Journal of Financial Economics* 118, 1–20.

Bessembinder, Hendrik, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman, 2018, Capital commitment and illiquidity in corporate bonds, *Journal of Finance* 73, 1615–1661.

Bessembinder, Hendrik, Kathleen M. Kahle, William F. Maxwell, and Danielle Xu, 2008, Measuring Abnormal Bond Performance, *Review of Financial Studies* 22, 4219–4258.

Binsbergen, Jules H. van, Yoshio Nozawa, and Michael Schwert, 2023, Duration-based valuation of corporate bonds, Working Paper.

Blume, Marshall E., and Robert Stambaugh, 1983, Biases in computed returns: An application to the size effect, *Journal of Financial Economics* 12, 387–404.

Bredendiek, Maximilian, Giorgio Ottonello, and Rossen I Valkanov, 2023, Implementable corporate bond portfolios, *Available at SSRN 2886825* .

Cao, Jie, Amit Goyal, Xiao Xiao, and Xintong Zhan, 2023, Implied volatility changes and corporate bond returns, *Management Science* 69, 1375–1397.

Ceballos, Luis, 2023, Inflation volatility risk and the cross-section of corporate bond returns, *Working paper* .

Chen, Andrew Y., and Mihail Velikov, 2023, Zeroing in on the expected returns of anomalies, *Journal of Financial and Quantitative Analysis* 58, 968–1004.

Chen, Andrew Y., and Tom Zimmermann, 2022, Open source cross-sectional asset pricing, *Critical Finance Review* 11, 207–264.

Chen, Long, David A. Lesmond, and Jason Wei, 2007, Corporate yield spreads and bond liquidity, *Journal of Finance* 62, 119–149.

Chen, Xi, Junbo Wang, and Chunchi Wu, 2022, Jump and volatility risk in the cross-section of corporate bond returns, *Journal of Financial Markets* 60, 100733.

Choi, Jaewon, Martijn Cremers, and Timothy B Riley, 2023a, Why have actively managed bond funds remained popular?, *Working Paper* .

Choi, Jaewon, Yesol Huh, and Sean Seunghun Shin, 2023b, Customer liquidity provision: Implications for corporate bond transaction costs, *Management Science* forthcoming.

Choi, Jaewon, and Yongjun Kim, 2018, Anomalies and market (dis)integration, *Journal of Monetary Economics* 100, 16–34.

Chordia, Tarun, Amit Goyal, Yoshio Nozawa, Avanidhar Subrahmanyam, and Qing Tong, 2017, Are capital market anomalies common to equity and corporate bond markets? An empirical investigation, *Journal of Financial and Quantitative Analysis* 52, 1301–1342.

Chung, Kee H., Junbo Wang, and Chunchi Wu, 2019, Volatility and the cross-section of corporate bond returns, *Journal of Financial Economics* 133, 397 – 417.

Dang, Thuy Duong, Fabian Hollstein, and Marcel Prokopczuk, 2023, Which factors for corporate bond returns?, *The Review of Asset Pricing Studies* raad005.

Detzel, Andrew, Robert Novy-Marx, and Mihail Velikov, 2023, Model comparison with transaction costs, *Journal of Finance* .

Dick-Nielsen, Jens, Peter Feldhütter, Lasse Heje Pedersen, and Christian Stolborg, 2023, Corporate bond factors: Replication failures and a new framework, Copenhargen Business School Working Paper.

Dick-Nielsen, Jens, and Marco Rossi, 2018, The Cost of Immediacy for Corporate Bonds, *Review of Financial Studies* 32, 1–41.

Dickerson, Alexander, Philippe Mueller, and Cesare Robotti, 2023, Priced risk in corporate bonds, *Journal of Financial Economics* 150, article 103707.

Diebold, Francis, and Roberto Mariano, 1995, Comparing predictive accuracy, *Journal of Business & Economic Statistics* 13.

Duan, Tinghua, Frank Weikai Li, and Quan Wen, 2021, Is carbon risk priced in the cross-section of corporate bond returns?, *Journal of Financial and Quantitative Analysis* 1–52.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen, 2005, Over-the-counter markets, *Econometrica* 73, 1815–1847.

Edwards, Amy, Lawrence E. Harris, and Michael S. Piwowar, 2007, Corporate bond market transaction costs and transparency, *Journal of Finance* 62, 1421–1451.

Eisfeldt, Andrea L., Bernard Herskovic, and Shuo Liu, 2023, Interdealer price dispersion, UCLA Working Paper.

Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.

Fama, Eugene F, and Kenneth R French, 2010, Luck versus skill in the cross-section of mutual fund returns, *Journal of Finance* 65, 1915–1947.

Feldhütter, Peter, 2010, The same bond at different prices: Identifying search frictions and selling pressures, *Review of Financial Studies* 25, 1155–1206.

Friewald, Nils, and Florian Nagler, 2016, Dealer inventory and the cross-section of corporate bond returns, *Working paper* .

Friewald, Nils, and Florian Nagler, 2019, Over-the-counter market frictions and yield spread changes, *Journal of Finance* 74, 3217–3257.

Gebhardt, William R., Soeren Hvidkjaer, and Bhaskaran Swaminathan, 2005a, The cross-section of expected corporate bond returns: Betas or characteristics?, *Journal of Financial Economics* 75, 85–114.

Gebhardt, William R., Soeren Hvidkjaer, and Bhaskaran Swaminathan, 2005b, Stock and bond market interaction: Does momentum spill over?, *Journal of Financial Economics* 75, 651–690.

Goldberg, Jonathan, and Yoshio Nozawa, 2021, Liquidity supply in the corporate bond market, *Journal of Finance* 76, 755–796.

Goldstein, Michael A., and Edith S. Hotchkiss, 2020, Providing liquidity in an illiquid market: Dealer behavior in us corporate bonds, *Journal of Financial Economics* 135, 16–40.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, 2020, Empirical asset pricing via machine learning, *Review of Financial Studies* 33, 2223–2273.

Haesen, Daniel, Patrick Houweling, and Jeroen van Zundert, 2017, Momentum spillover from stocks to corporate bonds, *Journal of Banking & Finance* 79, 28–41.

Harvey, Campbell R, Yan Liu, and Heqing Zhu, 2016, ... and the cross-section of expected returns, *Review of Financial Studies* 29, 5–68.

He, Xin, Guanhao Feng, Junbo Wang, and Chunchi Wu, 2024, Predicting individual corporate bond returns, *Working Paper* .

He, Zhiguo, Paymon Khorrami, and Zhaogang Song, 2019, Commonality in credit spread changes: Dealer inventory and intermediary distress, Working Paper, University of Chicago.

Houweling, Patrick, and Jeroen van Zundert, 2017, Factor investing in the corporate bond market, *Financial Analysts Journal* 73, 100–115.

Huang, Jing-Zhi, Nan Qin, and Ying Wang, 2013, Breadth of ownership and the cross-section of corporate bond returns, *Management Science , forthcoming* .

Israel, Ronen, Diogo Palhares, and Scott A Richardson, 2017, Common factors in corporate bond returns, *Journal of Investment Management* forthcoming.

Ivashchenko, Alexey, 2023, Corporate bond price reversals, Working Paper.

Ivashchenko, Alexey, and Robert Kosowski, 2023, Transaction costs and capacity of systematic corporate bond strategies, Working Paper.

Jacobsen, Stacey, and Kumar Venkataraman, 2023, Asymmetric information and receiving investor outcomes in the block market for corporate bonds, Working Paper.

Jostova, Gergana, Stanislava Nikolova, Alexander Philipov, and Christof W. Stahel, 2013, Momentum in corporate bond returns, *Review of Financial Studies* 26, 1649–1693.

Kargar, Mahyar, Benjamin Lester, Sébastien Plante, and Pierre-Olibier Weill, 2023, Sequential search for corporate bonds, Working Paper.

Kelly, Bryan T., Diogo Palhares, and Seth Pruitt, 2021, Modeling corporate bond returns, *Journal of Finance* forthcoming.

Kyle, Albert S, and Anna A Obizhaeva, 2016, Market microstructure invariance: Empirical hypotheses, *Econometrica* 84, 1345–1404.

Li, Sophia Zhengzi, Peixuan Yuan, and Guofu Zhou, 2023, Risk-based momentum of corporate bonds, *Working Paper* .

Lin, Hai, Junbo Wang, and Chunchi Wu, 2011, Liquidity risk and expected corporate bond returns, *Journal of Financial Economics* 99, 628–650.

Lin, Hai, Chunchi Wu, and Guofu Zhou, 2017, Trend momentum in corporate bonds, *Working paper* .

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Novy-Marx, Robert, and Mihail Velikov, 2015, A Taxonomy of Anomalies and Their Trading Costs, *Review of Financial Studies* 29, 104–147.

Novy-Marx, Robert, and Mihail Velikov, 2019, Comparing cost-mitigation techniques, *Financial Analysts Journal* 75, 85–102.

Nozawa, Yoshio, Yancheng Qiu, and Yan Xiong, 2023, Disagreement and price drifts in the corporate bond market, Working Paper.

O'Hara, Maureen, and Xing Alex Zhou, 2021, The electronic evolution of corporate bond dealers, *Journal of Financial Economics* 140, 368–390.

O'Hara, Maureen, Yihui Wang, and Xing (Alex) Zhou, 2018, The execution quality of corporate bonds, *Journal of Financial Economics* 130, 308–326.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, 2011, Scikit-learn: Machine learning in Python, *Journal of Machine Learning Research* 12, 2825–2830.

Pinter, Gabor, Chaojun Wang, and Junyuan Zou, 2021, Size discount and size penalty trading costs in bond markets, Discussion Papers 2114, Centre for Macroeconomics (CFM).

Rapach, David E, Jack K Strauss, and Guofu Zhou, 2010, Out-of-sample equity premium prediction: Combination forecasts and links to the real economy, *Review of Financial Studies* 23, 821–862.

Reichenbacher, Michael, and Philipp Schuster, 2022, Size-adapted bond liquidity measures and their asset pricing implications, *Journal of Financial Economics* 146, 425–443.

Sandulescu, Mirela, 2022, How integrated are corporate bond and stock markets?, Working Paper.

Schestag, Raphael, Philipp Schuster, and Marliese Uhrig-Homburg, 2016, Measuring Liquidity in Bond Markets, *Review of Financial Studies* 29, 1170–1219.

Sharpe, William F, 1991, The arithmetic of active management, *Financial Analysts Journal* 47, 7–9.

Tao, Xinyuan, Bo Wang, Junbo Wang, and Chunchi Wu, 2022, Economic policy uncertainty and the cross-section of corporate bond returns, *Journal of Fixed Income* 32, 6–44.

Wu, Botao, 2022, Post-crisis regulations, trading delays, and increasing corporate bond liquidity premium, CUHK Working Paper.

Figure 2: Delays and Inventory Dynamics

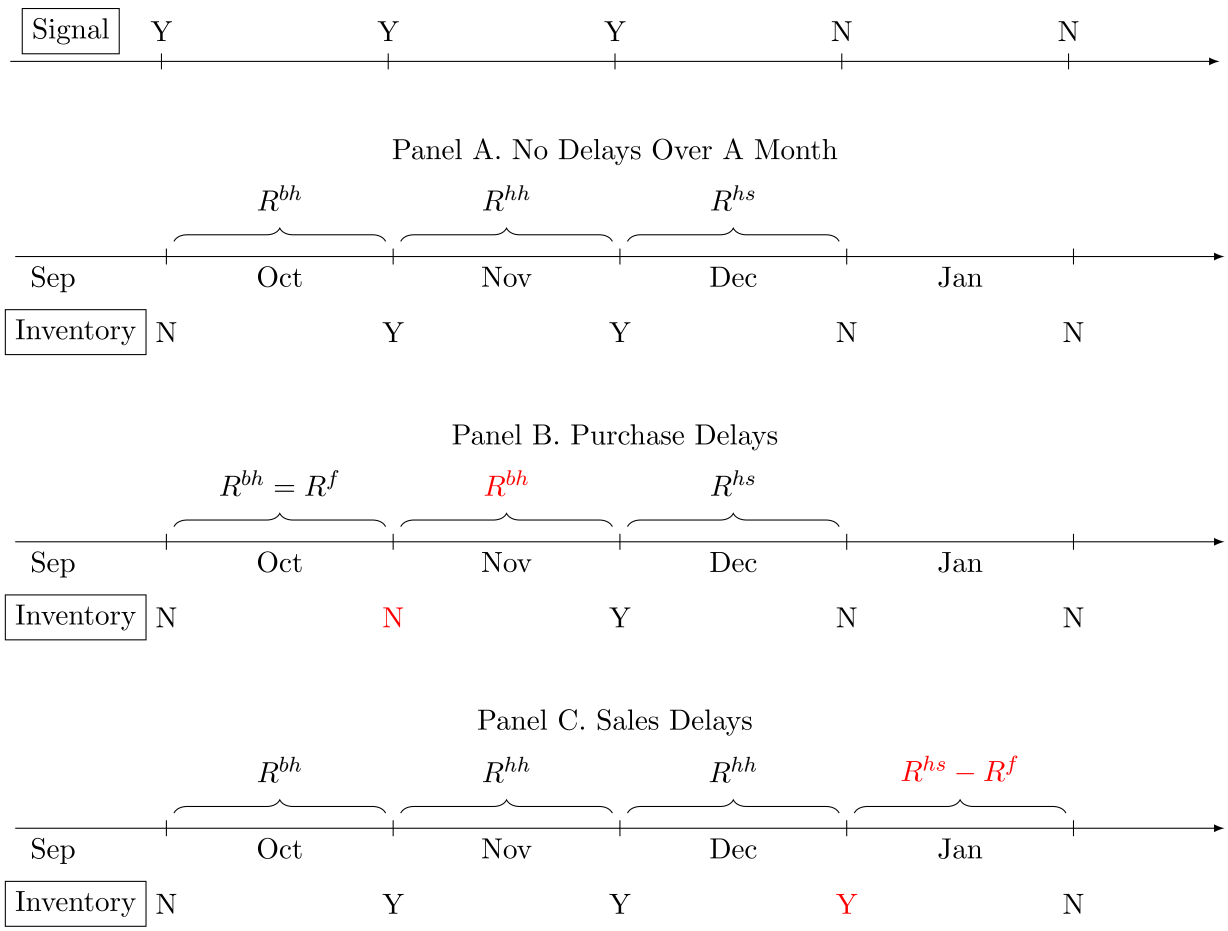

This figure illustrates a sequence of trades based on a trading signal. Panel A shows the types of returns used for a long position when there is no trading beyond one month. The two superscripts of a return R denote the action taken at the beginning and end of a month, respectively. The superscripts b indicate that the investor buys, s that the investor sells, and h that the investor holds the existing position. For inventory, Y indicates that the investor has a bond in inventory and N indicates that he does not. Panel B shows the case where there is a delay in purchasing a bond. Panel C shows the case where there is a delay in selling a bond.

Figure 3: Summary Statistics For Different Trade Size

Panel A. Mean and Median Returns By Size

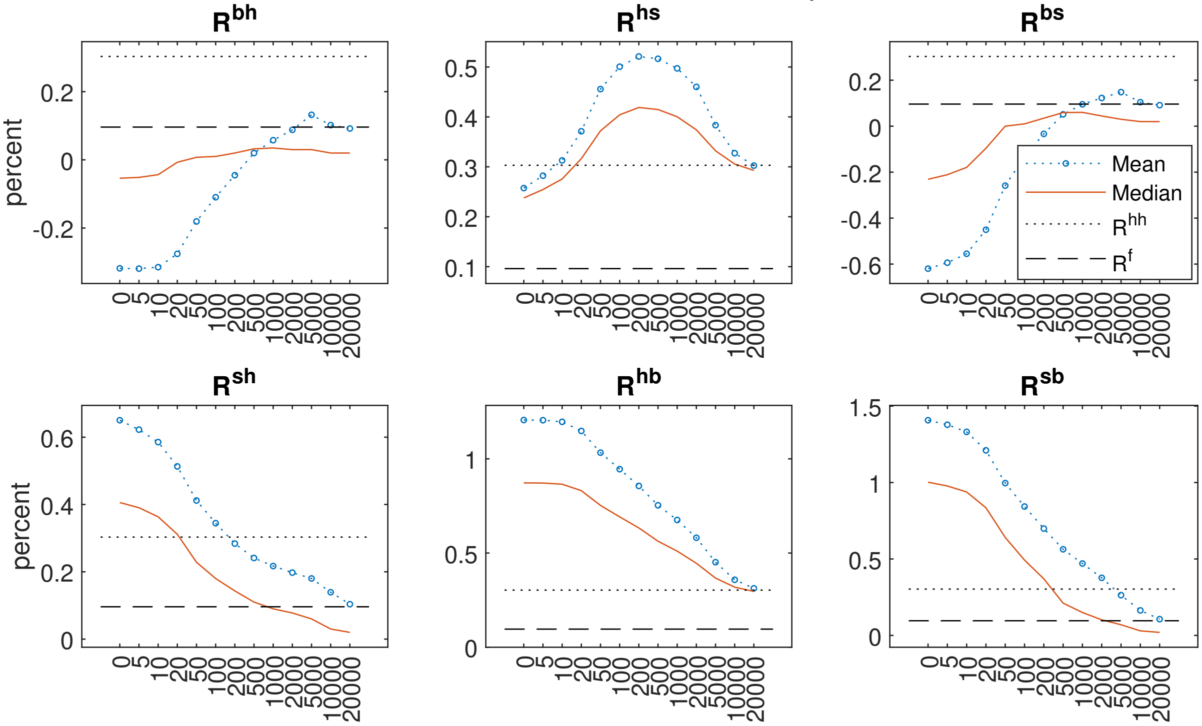

Panel B. Percentage of Observations with No Trade

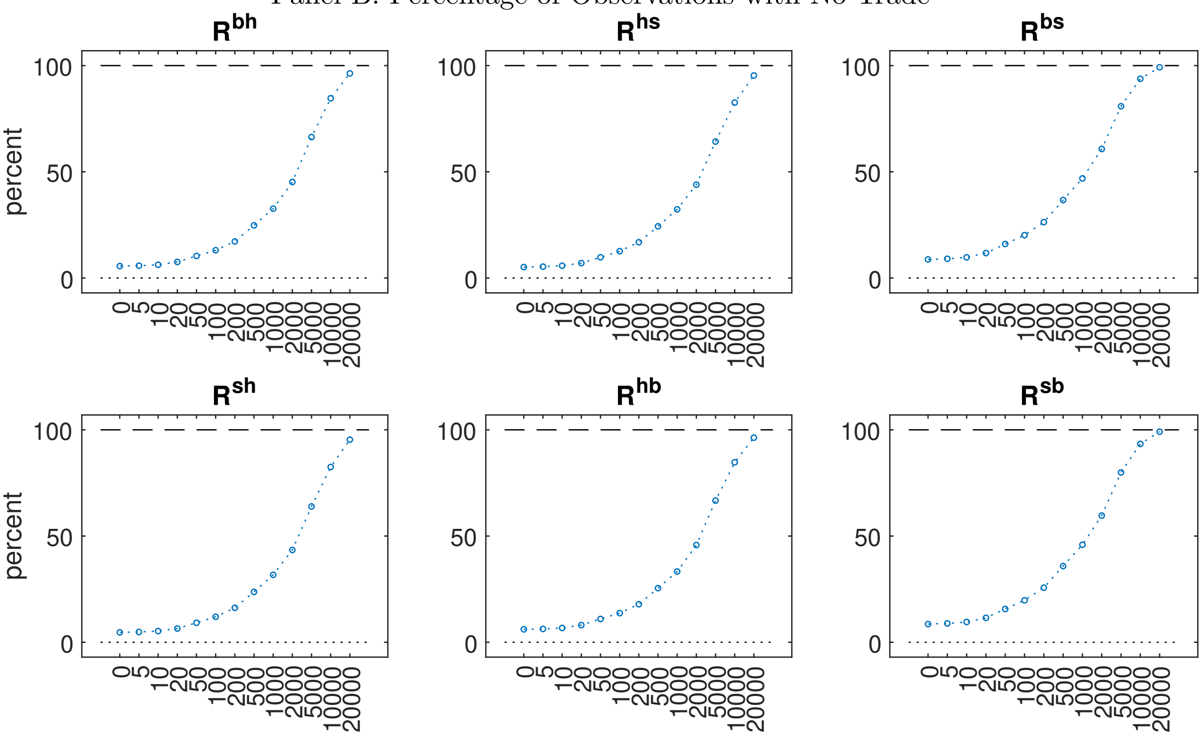

Panel A plots the mean and median net returns for different trade sizes. Panel B plots the percentage of observations where there is no trade to calculate a return in the month. Values on the x-axis are in thousand dollars.

Figure 4: Distribution of Trade Size: July 2002-December 2022

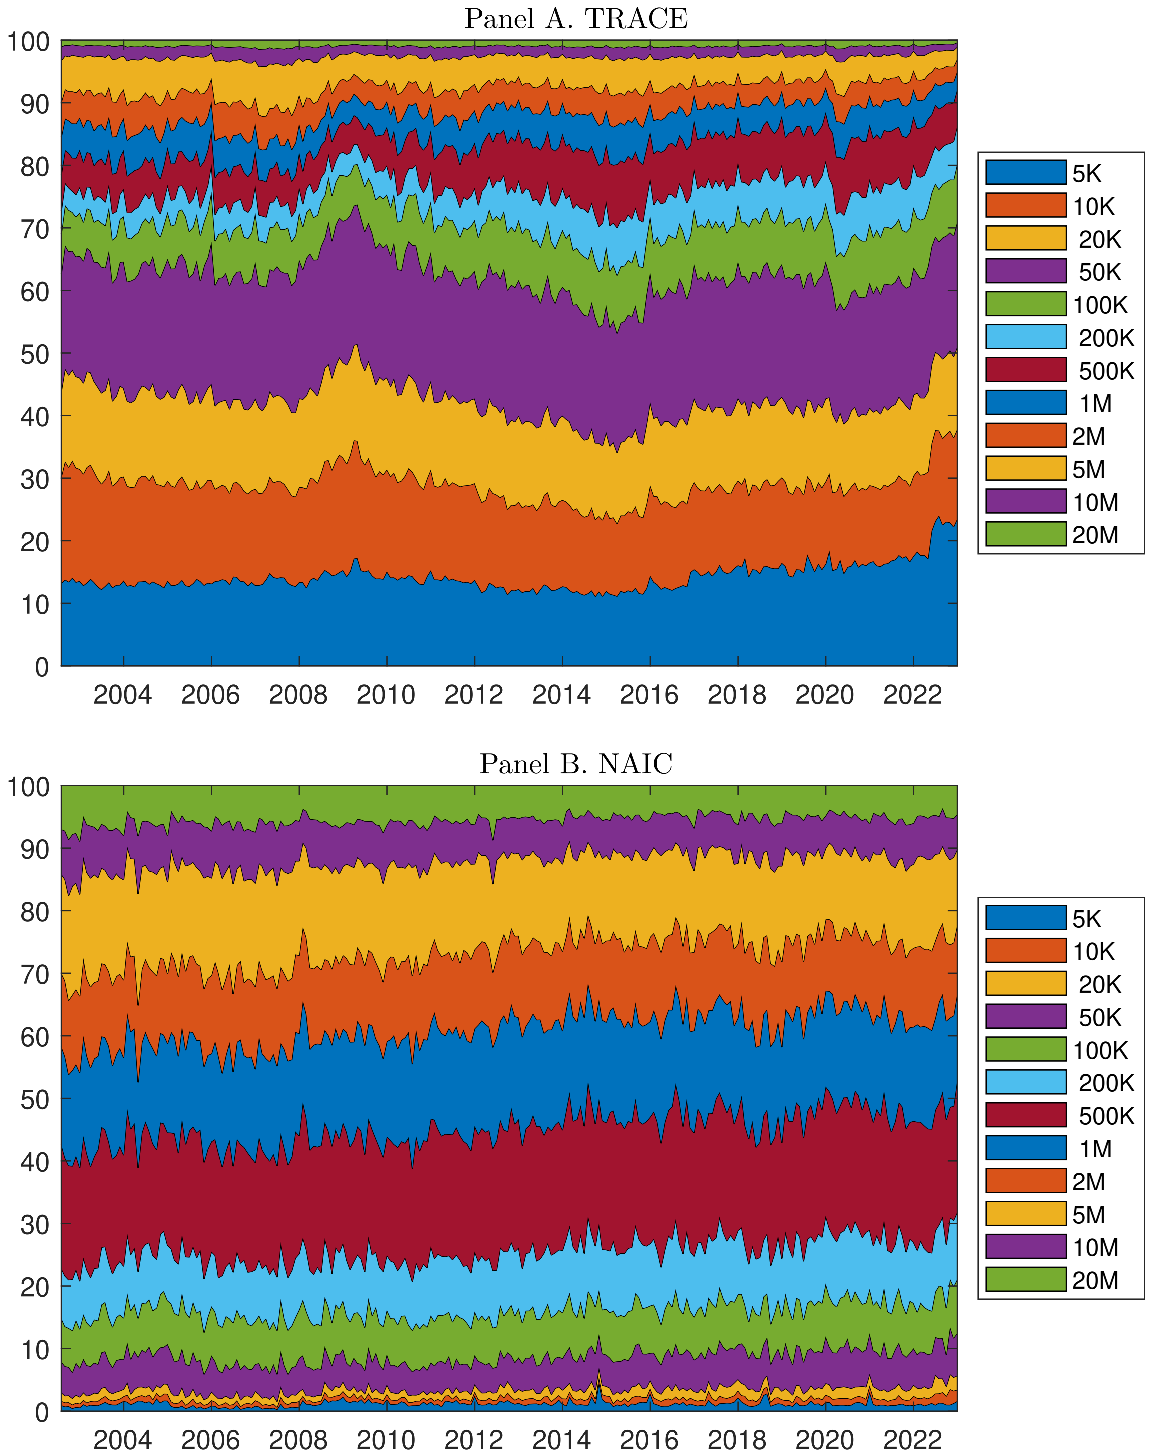

This figure plots the cumulative frequency of trade size observed in the corporate bond market (Panel A) and for insurance firms' trade size only (Panel B). For example, the area below 10K represents the number of transactions with a size below $10,000. The sample is from July 2002 to December 2022 and includes only dealer-customer trades.

Figure 5: Average Excess Returns and CAPMB Alphas of the ML Strategies: Optimal Volume

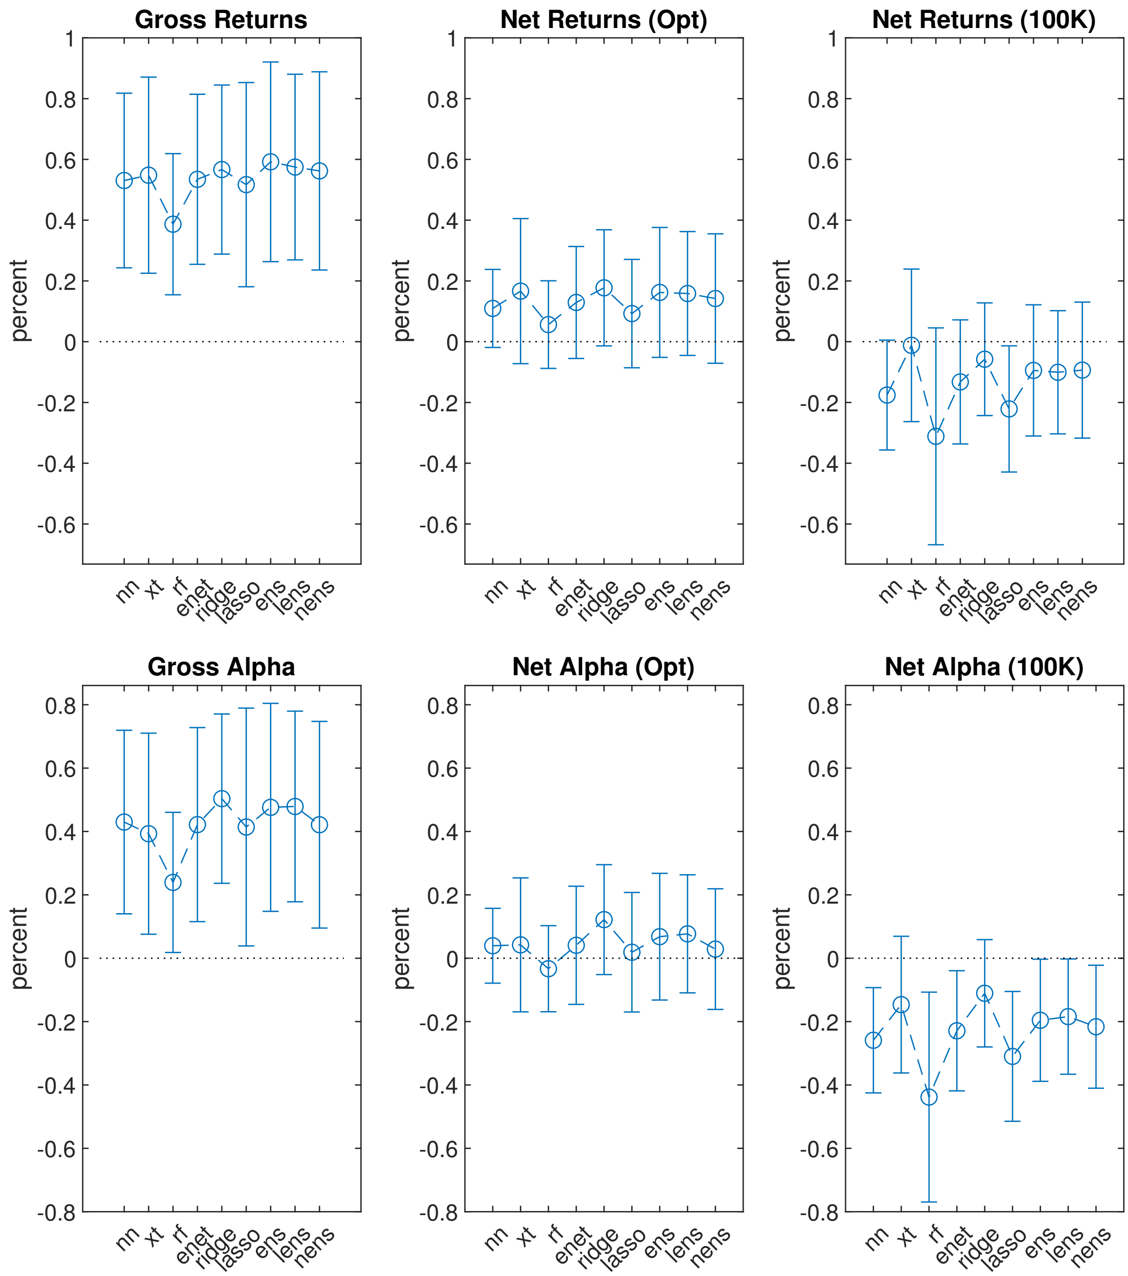

This figure shows the point estimates and associated two standard error bars for the long-short portfolios based on the expected returns generated by the machine learning algorithms. Gross returns and alphas are before transaction costs, and net returns and alphas are after costs. In the middle panels, transaction costs are calculated using the optimal threshold that maximizes the net alpha. In the right panels, transaction costs are calculated using trades with a size above $100,000.

Figure 6: t-Statistics of Individual Signals

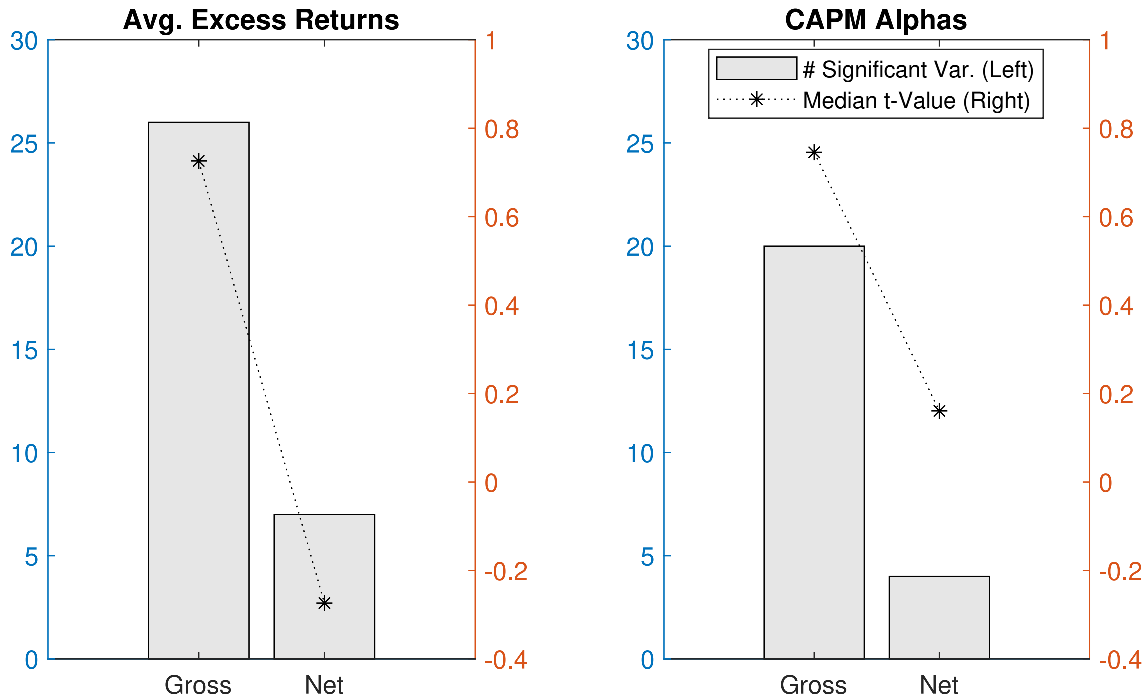

This figure characterizes the distribution of the t-statistics of the performance of 200 individual signals. In the left panel, we compute t-statistics of the average excess returns of the long-short portfolios based on each signal. In the right panel, we compute t-statistics of the CAPM alphas of the long-short portfolios. The bar chart shows the number of characteristics with a t-value greater than 1.96. The point shows the median t-value. Gross is the value before accounting for transaction costs and net is the value after costs at the optimal trade size. Standard errors are adjusted for serial correlation using Newey and West (1987) 12 lags.

Figure 7: Effect of Transaction Costs: ML Strategies

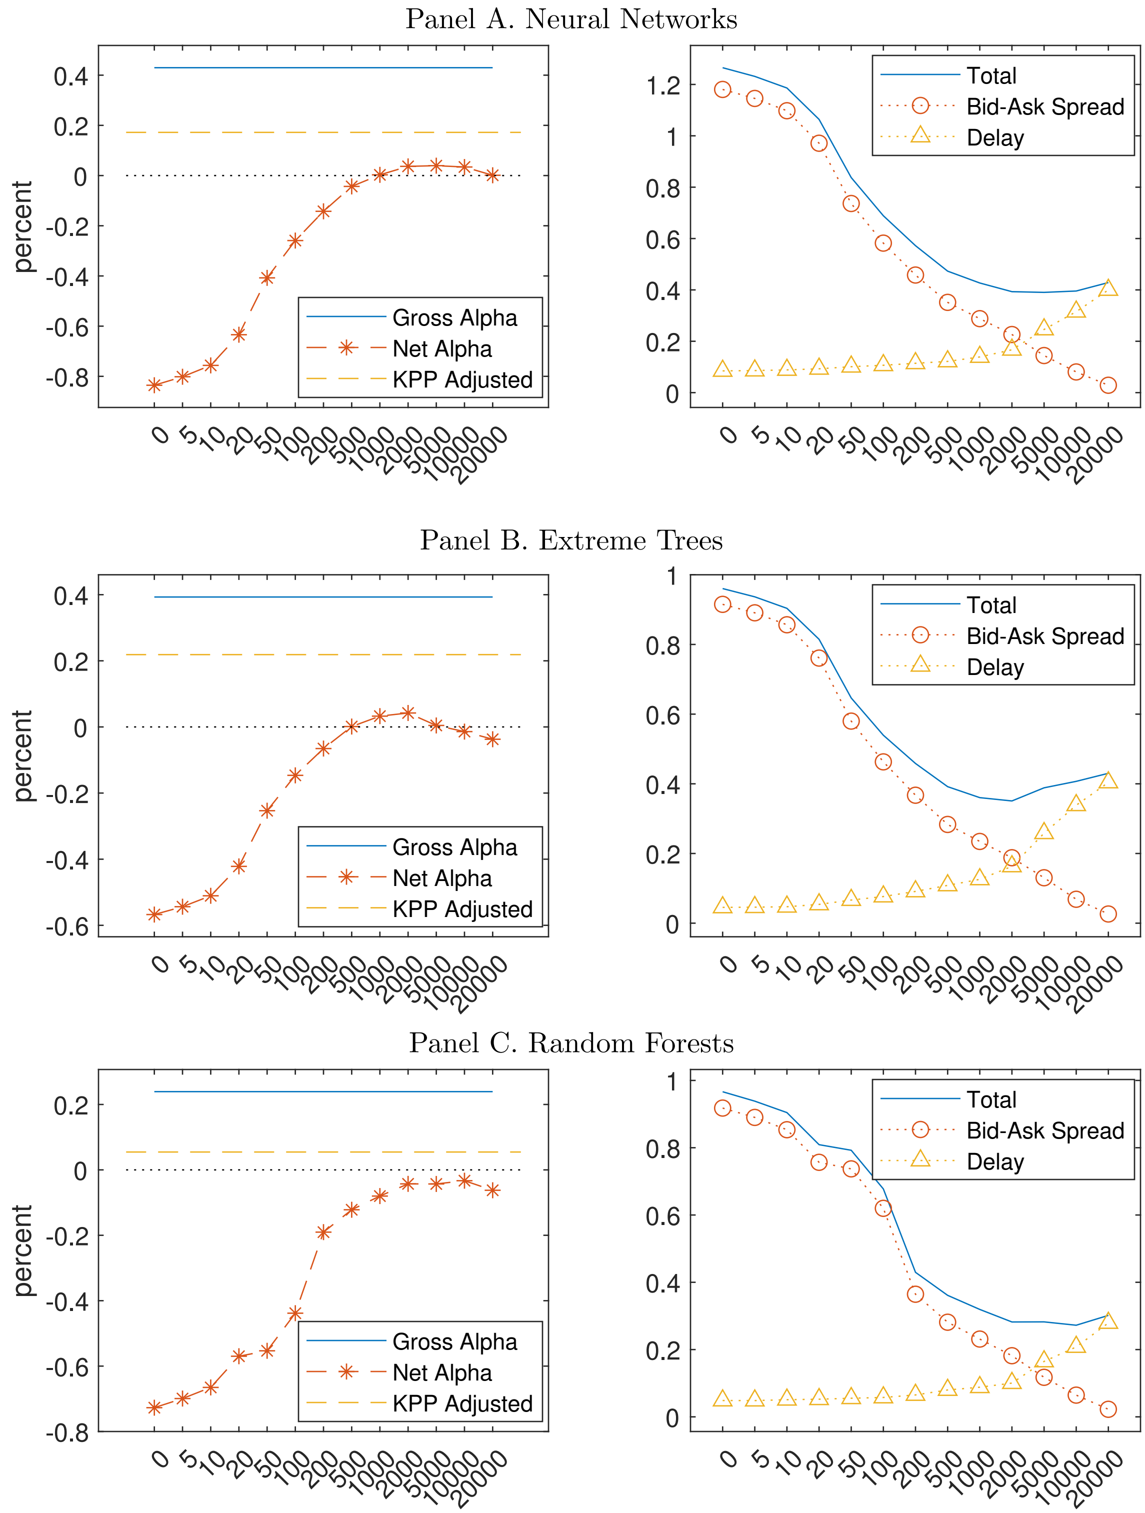

This figure plots the bond CAPM alphas on the long-short strategies before and after accounting for transaction costs (left panels). The transaction costs are decomposed into bid-ask spreads and delays (right panels). Values on the x-axis are in thousand dollars.

Figure 7, Continued

Panel D. Elastic Net

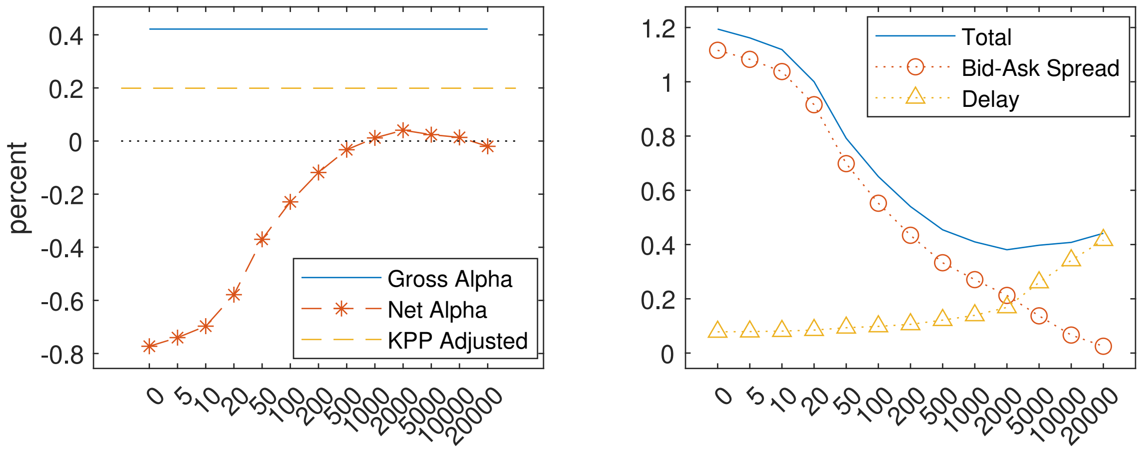

Panel E. Ridge

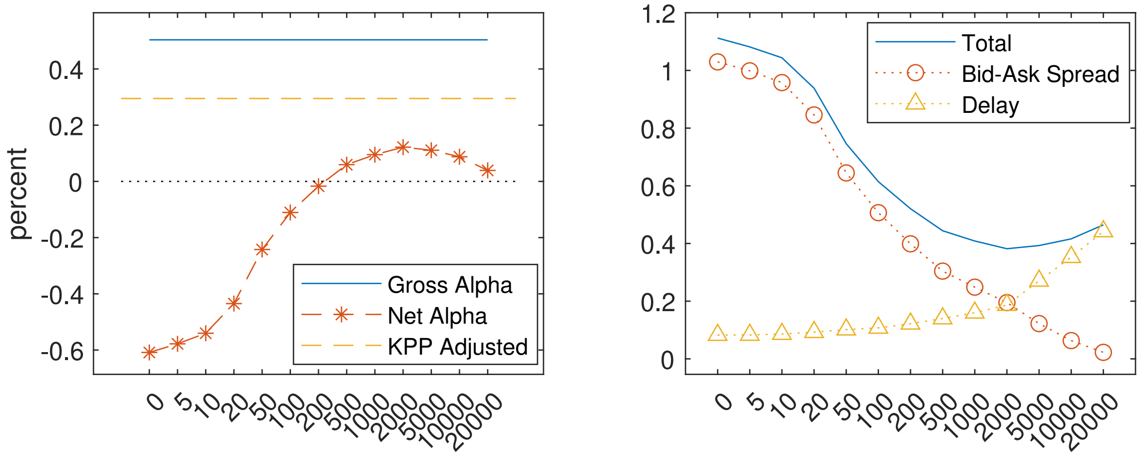

Panel F. Lasso

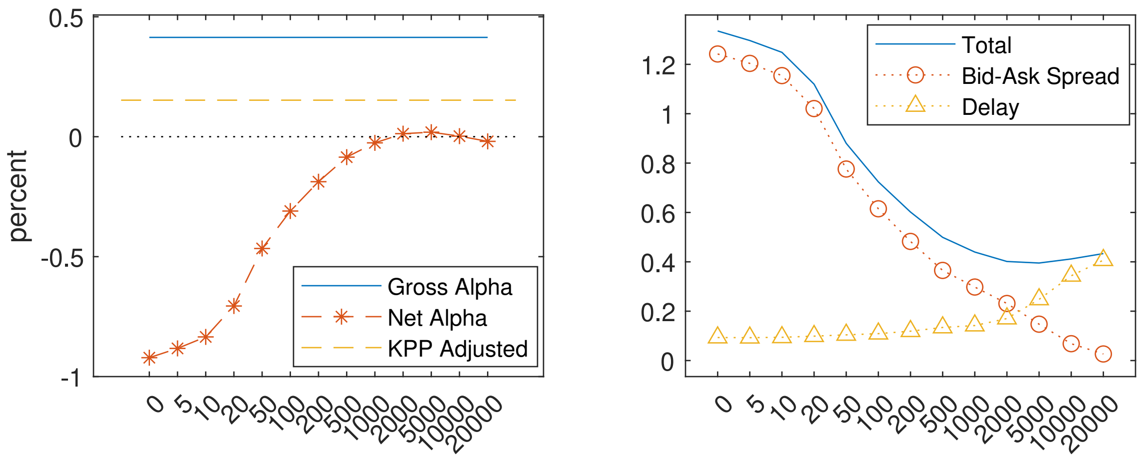

Figure 7, Continued

Panel G. Ensemble Across All Models

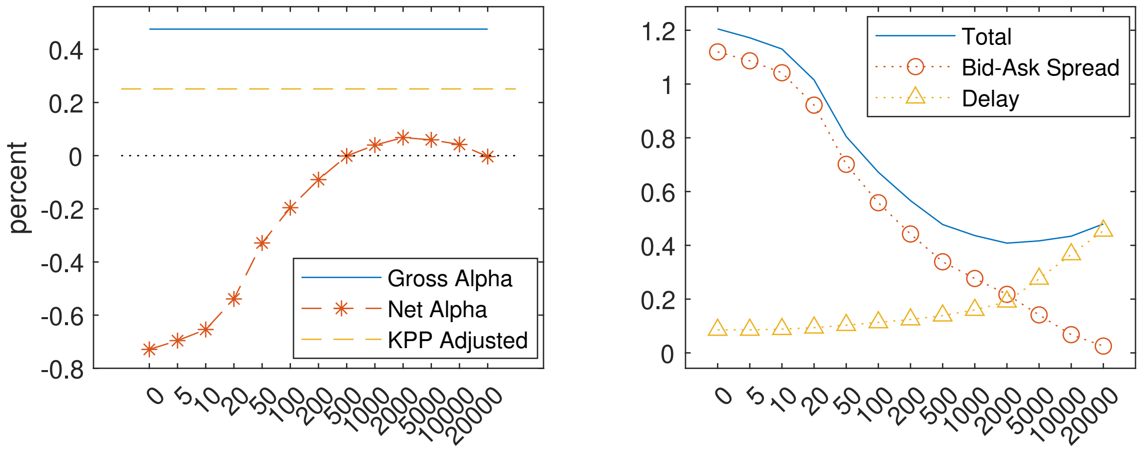

Panel H. Ensemble Across Linear Models

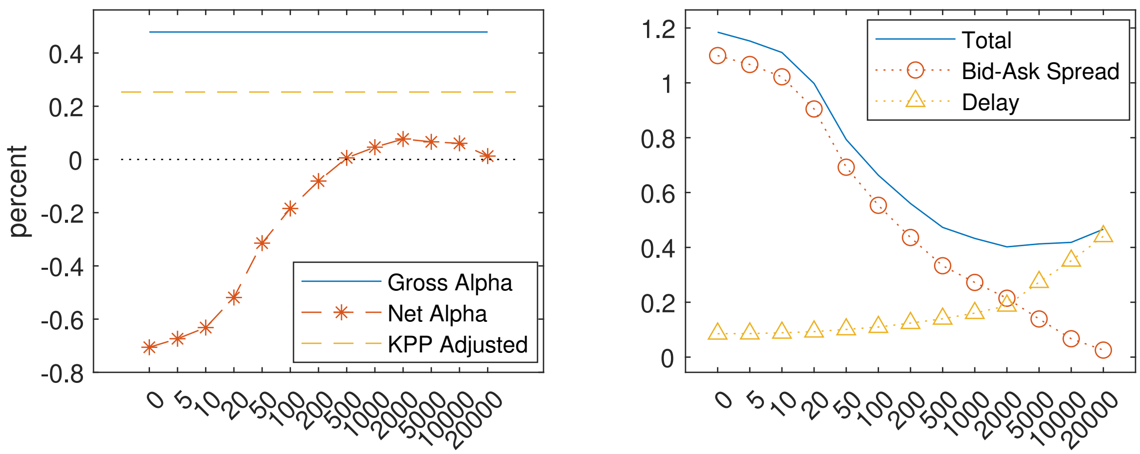

Panel I. Ensemble Across Nonlinear Models

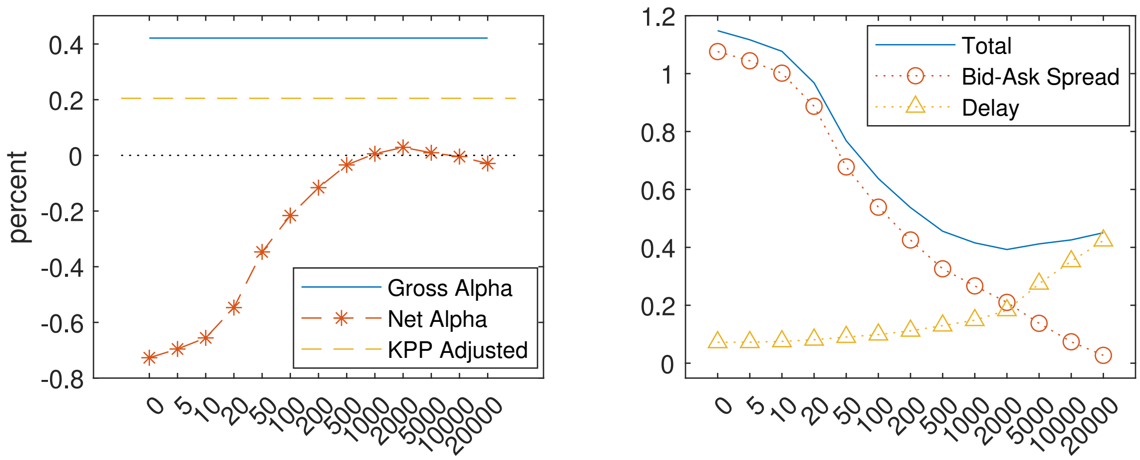

Figure 8: Subperiod Analysis for ENS

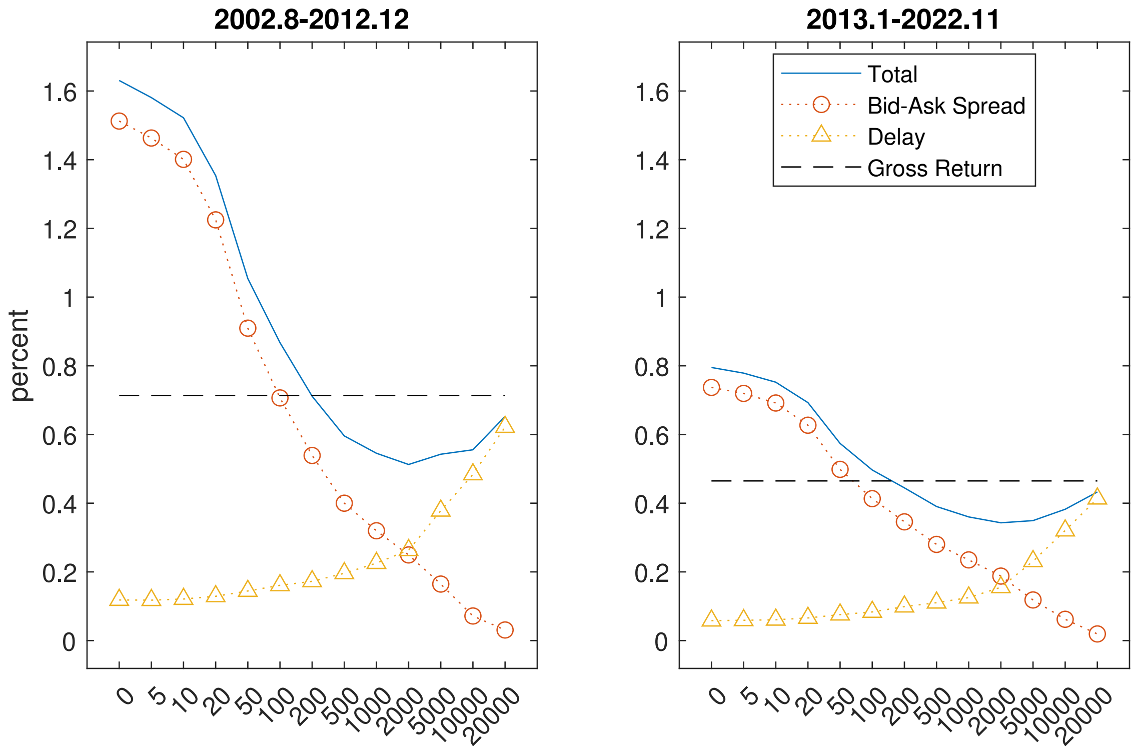

This figure shows the decomposition of transaction costs for the ENS strategy using the two subperiods. In this figure, total costs are the difference between gross and net average returns. The bid-ask spread cost is the difference between the gross average return and an alternative net average return in which transaction prices are replaced by quotes on the day of the transaction. Values on the x-axis are in thousand dollars.

Figure 9: Time-Variation in Transaction Costs for ENS

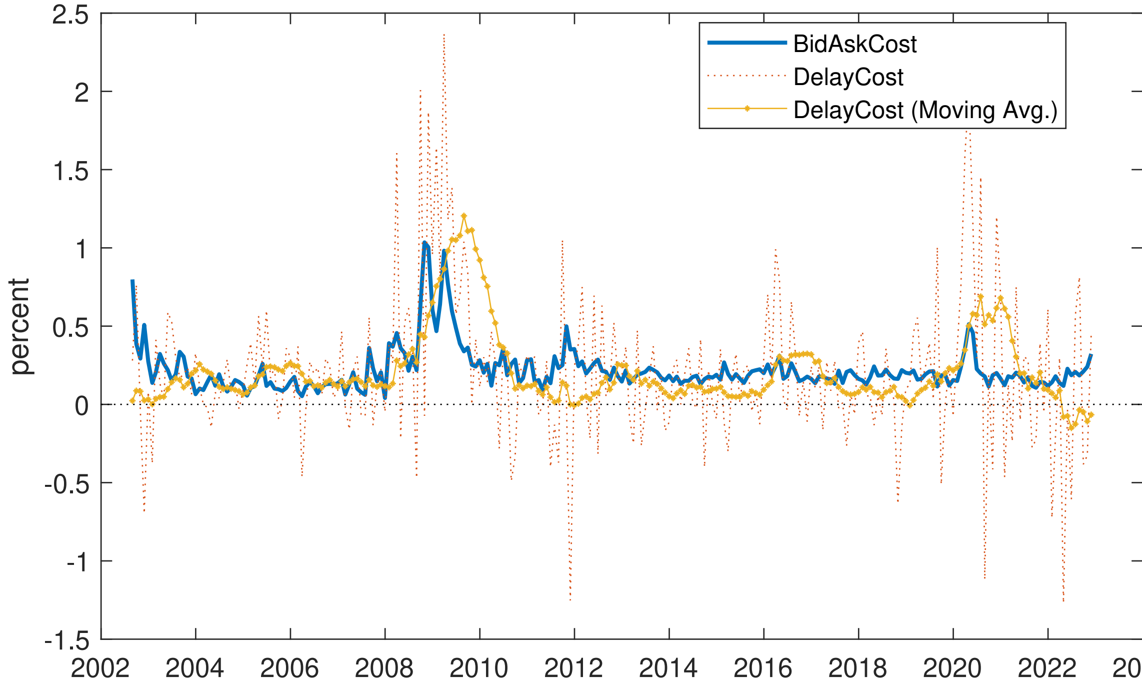

The figure plots the transaction costs for the ENS strategy with a trade size of $2 million, decomposed into the cost due to bid-ask spreads and portfolio turnover and the cost due to delays. Delay Cost (Moving Avg.) plots the 12-month moving average of the delay costs.

Figure 10: Turnover Rate, Gross and Net CAPM α

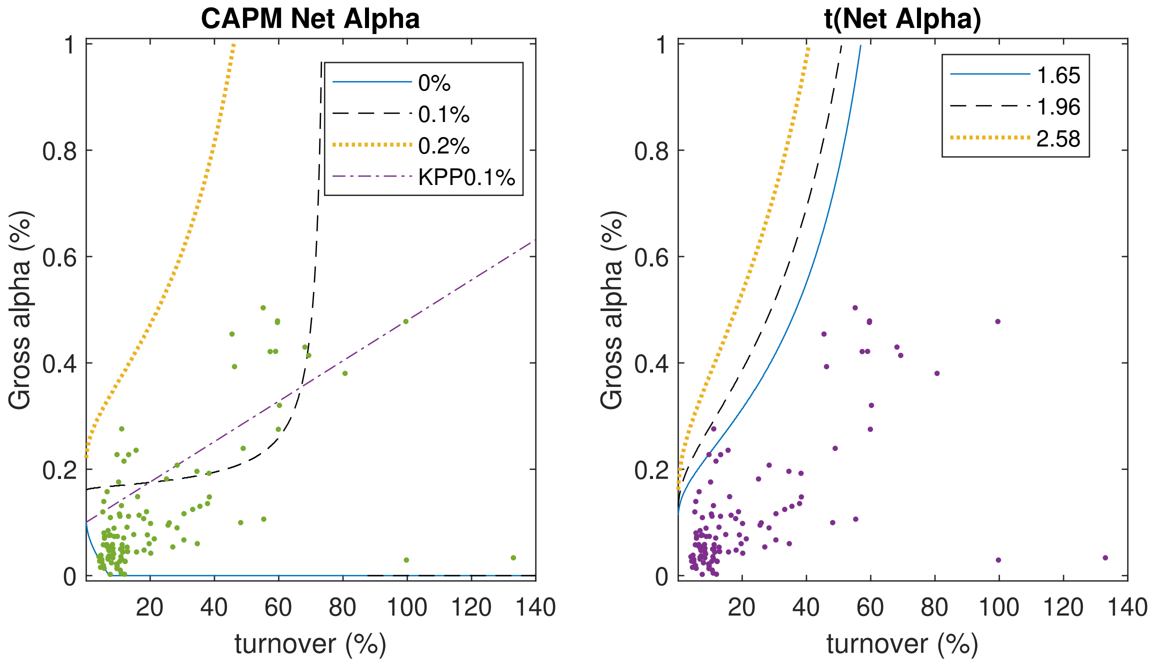

The figures plot the combination of gross CAPMB α and portfolio turnover rate that matches the target values of net CAPMB α and the associated t-statistics. The boundaries are estimated by regressing the net CAPMB α's and t-statistics on gross α, log portfolio turnover rate, and the product of the two. The regression uses 114 strategies for which we can find an optimal trade size.

Figure 11: Mutual Fund Alphas

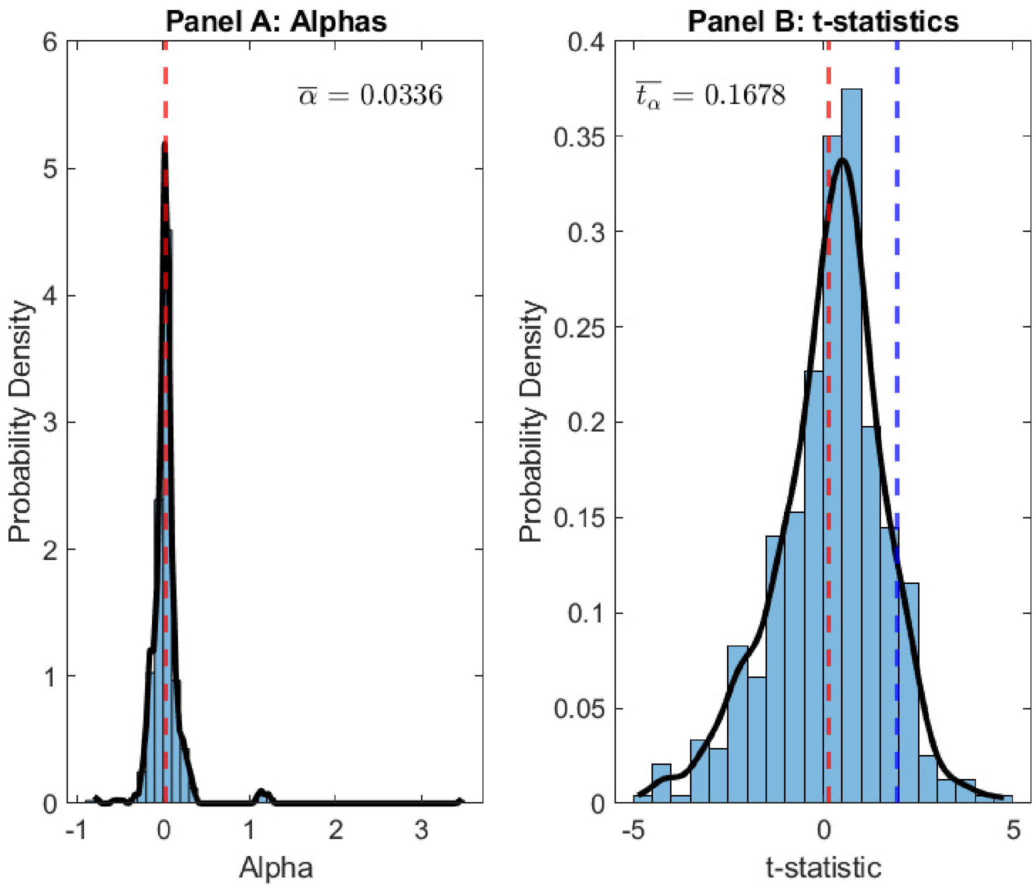

This figure plots the cross-sectional distribution of the corporate bond net of costs single-factor $MKTB_{Net}$ alphas and associated t-statistics. The dashed red lines indicate the mean values for the alphas (t-statistics). The dashed blue line represents the cut-off value for the 95% level of significance ($t = 1.96$). The sample includes 540 corporate bond mutual funds over the sample period 2002:07–2022:11.

Table 1: **List of Papers on the Cross-Section of Corporate Bonds**

Article	Cost Estimates
Panel A. Papers Without Transaction Costs	
Bai, Bali, and Wen (2019)	
Bai, Bali, and Wen (2021)	
Bali, Subrahmanyam, and Wen (2021a)	
Bali, Subrahmanyam, and Wen (2021b)	
Ceballos (2023)	
Chen, Wang, and Wu (2022)	
Chung, Wang, and Wu (2019)	
Dang, Hollstein, and Prokopczuk (2023)	
Dick-Nielsen, Feldhütter, Pedersen, and Stolborg (2023)	
Duan, Li, and Wen (2021)	
Friewald and Nagler (2016)	
Gebhardt, Hvidkjaer, and Swaminathan (2005a)	
Gebhardt, Hvidkjaer, and Swaminathan (2005b)	
Haesen, Houweling, and Zundert (2017)	
Huang, Qin, and Wang (2013)	
Li, Yuan, and Zhou (2023)	
Lin, Wang, and Wu (2011)	
Tao, Wang, Wang, and Wu (2022)	
Panel B. Papers Incorporating Transaction Costs	
Bali et al. (2020)	Roll measure of Bao et al. (2011)
Bali, Beckmeyer, and Goyal (2023)	Fixed at 35bps
Bartram, Grinblatt, and Nozawa (2023)	Portfolio-level bid-ask spreads
Bredendiek, Ottonello, and Valkanov (2023)	Round-trip transaction costs
Cao et al. (2023)	Estimates following Edwards et al. (2007)
Choi and Kim (2018)	Considers transaction costs as characteristics
Chordia et al. (2017)	Portfolio-level bid-ask spreads
He, Feng, Wang, and Wu (2024)	Fixed at 20 to 80bps
Houweling and Zundert (2017)	Maturity-rating, following Chen et al. (2007)
Israel, Palhares, and Richardson (2017)	Maturity-rating, following Chen et al. (2007)
Ivashchenko (2023)	Average 12m moving average of bond bid-ask spreads
Ivashchenko and Kosowski (2023)	Estimates following Kyle and Obizhaeva (2016)
Jostova et al. (2013)	Estimates following Edwards et al. (2007)
Kelly, Palhares, and Pruitt (2021)	Fixed at 19bps
Lin, Wu, and Zhou (2017)	Break-even transaction costs
Nozawa, Qiu, and Xiong (2023)	Bond-level bid-ask spreads

Table 2: Average and Median Half Spreads (Percent) of Insurance Firms' Bond Trade

		Investor Size Category				
Trade Size Category		Small	2	3	4	Large
(0,100K]	mean	0.486	0.354	0.405	0.268	0.237
(100K,500K]	mean	0.461	0.170	0.228	0.157	0.131
(500K,1M]	mean	0.312	0.113	0.116	0.083	0.086
(1M,∞]	mean	0.223	0.095	0.084	0.072	0.088
(0,100K]	median	0.199	0.130	0.168	0.087	0.090
(100K,500K]	median	0.233	0.060	0.078	0.048	0.041
(500K,1M]	median	0.151	0.047	0.044	0.032	0.031
(1M,∞]	median	0.119	0.040	0.021	0.017	0.025
(0,100K]	n	26,173	13,224	17,119	19,474	17,489
(100K,500K]	n	68,185	37,586	36,437	40,139	36,044
(500K,1M]	n	31,427	22,981	20,327	21,536	17,743
(1M,∞]	n	55,636	98,332	96,935	95,836	100,222

This table reports the summary statistics of half spreads for insurance firms' transaction of corporate bonds in TRACE. For each trade, a half spread is computed as the difference in price between the transaction and the latest interdealer transaction. Then, the spread is classified into bins based on their trade size and the insurance firm's total bond holding in the preceding quarter. The sample is from July 2002 to December 2022.

Table 3: Trading Costs and Trade Size in the Corporate Bond Market.

	(1)	(2)	(3)	(4)	(5)	(6)
Log(Volume)	−0.06***	−0.04***	−0.05***	−0.05***	−0.03***	−0.05***
	(−14.45)	(−13.14)	(−14.14)	(−14.59)	(−12.09)	(−14.26)
Day FE	Yes	Yes	Yes	Yes	Yes	Yes
Client FE	No	Yes	No	No	Yes	No
Broker FE	No	No	Yes	No	No	Yes
Controls	No	No	No	Yes	Yes	Yes
Obs.	881,504	881,474	881,369	731,080	731,042	730,960
R^2	0.039	0.088	0.059	0.055	0.090	0.067

This table regresses trading costs on trade size, various fixed effects and bond-level control variables. The bond level control variables include bond rating, coupon, maturity and the log of the issue size (amount outstanding). t-statistics in parentheses are based on two-way clustered standard errors at the day and client level. Asterisks denote significance levels (* p <0.1, ** p <0.05, *** p <0.01).

Table 4: Summary Statistics for Returns: Volume of $10K, $100K, $1M, and $10M

Variable	N	Mean	Std.	p1	p10	p50	p90	p99	NoTrade(%)
R^{hh}	854,250	0.30	3.71	-9.70	-2.36	0.29	2.89	9.49	
Panel A. $10K									
R^{bh}	854,250	-0.31	3.45	-10.08	-2.91	-0.04	1.97	7.62	6.24
R^{hs}	854,250	0.31	4.19	-10.81	-2.82	0.28	3.47	10.74	5.80
R^{bs}	854,250	-0.56	3.76	-11.28	-3.44	-0.18	1.88	7.94	9.76
R^{sh}	854,250	0.59	3.69	-8.43	-1.90	0.36	3.25	10.12	5.27
R^{hb}	854,250	1.20	4.28	-9.14	-1.86	0.87	4.57	13.03	6.73
R^{sb}	854,250	1.33	4.12	-7.77	-1.38	0.94	4.48	12.75	9.54
Panel B. $100K									
R^{bh}	854,250	-0.11	3.36	-8.81	-2.35	0.01	1.92	7.32	13.08
R^{hs}	854,250	0.50	4.27	-10.67	-2.58	0.40	3.67	11.41	12.63
R^{bs}	854,250	-0.13	3.74	-9.81	-2.66	0.01	2.13	8.34	20.19
R^{sh}	854,250	0.34	9.12	-7.97	-1.80	0.18	2.57	8.65	12.00
R^{hb}	854,250	0.95	4.32	-9.68	-2.09	0.69	4.21	12.56	13.70
R^{sb}	854,250	0.84	9.33	-7.94	-1.54	0.49	3.51	10.97	19.75
Panel C. $1M									
R^{bh}	854,250	0.06	2.97	-7.05	-1.48	0.03	1.57	6.69	32.65
R^{hs}	854,250	0.50	4.33	-10.89	-2.56	0.40	3.63	11.59	32.33
R^{bs}	854,250	0.10	3.44	-8.19	-1.74	0.06	1.89	8.08	46.90
R^{sh}	854,250	0.22	2.88	-6.76	-1.25	0.09	1.86	7.17	31.73
R^{hb}	854,250	0.68	4.32	-10.34	-2.35	0.51	3.85	12.06	33.30
R^{sb}	854,250	0.47	3.40	-7.40	-1.20	0.15	2.57	9.13	45.89
Panel D. $10M									
R^{bh}	854,250	0.10	1.74	-2.80	0.00	0.02	0.36	3.18	84.62
R^{hs}	854,250	0.33	4.12	-10.39	-2.46	0.31	3.09	10.26	82.59
R^{bs}	854,250	0.10	2.03	-3.22	0.00	0.02	0.37	3.60	93.78
R^{sh}	854,250	0.14	1.79	-2.69	0.00	0.03	0.40	3.60	82.51
R^{hb}	854,250	0.36	4.11	-10.23	-2.42	0.32	3.12	10.33	84.74
R^{sb}	854,250	0.16	2.14	-2.98	0.00	0.03	0.41	4.20	93.37

This table shows the summary statistics of the panel data used in the study. The sample spans from August 2002 to November 2022. Panels A, B, C, and D correspond to the statistics with volume thresholds of $10K, $100K, $1M, and $10M, respectively. NoTrade(%) is the percentage of monthly observations in which there is no trade in TRACE above the volume threshold.

Table 5: Summary Statistics for Returns: Volume of $10K$,100K, $1M, and $10M: Duration-Adjusted Returns

Variable	N	Mean	Std.	p1	p10	p50	p90	p99	NoTrade(%)
R^{hh}	854,216	0.20	3.51	-9.66	-1.47	0.13	2.05	8.95	
Panel A. $10K									
R^{bh}	854,216	-0.48	3.51	-10.36	-3.09	-0.18	1.72	7.89	6.24
R^{hs}	854,216	0.09	4.36	-13.15	-2.89	0.12	3.22	10.44	5.80
R^{bs}	854,216	-0.75	4.08	-13.27	-3.96	-0.32	2.01	8.54	9.76
R^{sh}	854,216	0.41	3.73	-8.49	-2.02	0.21	2.96	10.40	5.27
R^{hb}	854,216	0.97	4.38	-10.79	-1.86	0.67	4.24	12.80	6.73
R^{sb}	854,216	0.97	4.34	-9.45	-1.96	0.58	4.28	13.16	9.54
Panel B. $100K									
R^{bh}	854,216	-0.29	3.37	-8.90	-2.44	-0.05	1.53	7.48	13.08
R^{hs}	854,216	0.33	4.34	-12.47	-2.42	0.25	3.34	10.95	12.63
R^{bs}	854,216	-0.31	3.91	-11.29	-2.92	0.00	2.03	8.66	20.19
R^{sh}	854,216	0.15	9.12	-8.15	-1.87	0.08	2.16	8.80	12.00
R^{hb}	854,216	0.77	4.36	-11.03	-1.96	0.53	3.86	12.13	13.70
R^{sb}	854,216	0.52	9.39	-9.37	-1.96	0.25	3.19	11.05	19.75
Panel C. $1M									
R^{bh}	854,216	-0.10	2.95	-7.16	-1.52	0.01	1.02	6.59	32.65
R^{hs}	854,216	0.38	4.31	-11.60	-2.32	0.28	3.27	11.17	32.33
R^{bs}	854,216	-0.05	3.48	-8.71	-1.80	0.01	1.47	8.11	46.90
R^{sh}	854,216	0.06	2.87	-6.96	-1.25	0.03	1.29	7.14	31.73
R^{hb}	854,216	0.57	4.30	-11.03	-2.11	0.40	3.50	11.62	33.30
R^{sb}	854,216	0.24	3.44	-8.18	-1.34	0.09	2.00	9.06	45.89
Panel D. $10M									
R^{bh}	854,216	0.06	1.72	-2.78	0.00	0.02	0.32	2.61	84.62
R^{hs}	854,216	0.31	4.11	-10.42	-2.37	0.28	2.98	10.16	82.59
R^{bs}	854,216	0.07	2.02	-3.09	0.00	0.02	0.32	3.06	93.78
R^{sh}	854,216	0.09	1.76	-2.65	0.00	0.02	0.34	3.03	82.51
R^{hb}	854,216	0.34	4.10	-10.27	-2.34	0.30	3.01	10.21	84.74
R^{sb}	854,216	0.11	2.11	-2.90	0.00	0.02	0.36	3.54	93.37

This table reports the summary statistics of the panel data used for the study. The sample is from August 2002 to November 2022. Panels A, B, C, and D corresponds to the statistics with the volume threshold of $10K, $100K, $1M, and $10M, respectively. NoTrade(%) is the percentage of monthly observations where there is no trade in TRACE that is above the volume threshold.

Table 6: Performance of Top 15 Signals

	CAPMB α		Info. Ratio			CAPM α		Info. Ratio	
Signal	Gross	Net	Gross	Net	Signal	Gross	Net	Gross	Net
Panel A. Sort By Gross IR					Panel B. Sort By Net IR				
strev	0.478	0.068	1.375	0.340	bondkurtosis	0.276	0.178	0.783	0.696
	(3.18)	(1.08)				(2.69)	(2.36)		
pead	0.192	0.063	1.232	0.429	coskewacx	0.182	0.108	0.800	0.644
	(5.50)	(1.90)				(3.54)	(2.78)		
mom6mspread¶	0.454	0.178	0.884	0.554	mom6mspread¶	0.454	0.178	0.884	0.554
	(2.93)	(1.67)				(2.93)	(1.67)		
trendfactor	0.275	0.345	0.879	0.292	skew¶	0.215	0.121	0.623	0.472
	(2.53)	(1.30)				(1.72)	(1.38)		
dvolcall¶	0.243	-0.127	0.852	-0.574	duration¶	0.158	0.163	0.368	0.454
	(4.44)	(-3.38)				(1.69)	(2.04)		
dvolput	0.217	-0.148	0.823	-0.743	mrreversal	0.060	0.082	0.221	0.451
	(4.18)	(-4.07)				(0.87)	(1.62)		
coskewacx	0.182	0.108	0.800	0.644	pead	0.192	0.063	1.232	0.429
	(3.54)	(2.78)				(5.50)	(1.90)		
bondkurtosis	0.276	0.178	0.783	0.696	tmt¶	0.139	0.140	0.361	0.418
	(2.69)	(2.36)				(1.71)	(1.91)		
returnskew¶	0.121	-0.165	0.731	-1.176	exchswitch	0.130	0.222	0.271	0.408
	(2.94)	(-5.66)				(1.14)	(1.70)		
chtax	0.117	0.028	0.711	0.198	lrreversal	0.077	0.091	0.269	0.395
	(2.48)	(0.70)				(0.98)	(1.53)		
mom3mspread¶	0.320	0.069	0.668	0.336	dnoa	0.058	0.046	0.436	0.375
	(2.47)	(1.04)				(1.31)	(1.22)		
skew¶	0.215	0.121	0.623	0.472	cheq	0.051	0.054	0.282	0.375
	(1.72)	(1.38)				(1.34)	(1.78)		
mom12off	0.208	0.055	0.590	0.197	volsd¶	0.080	0.107	0.214	0.373
	(2.61)	(0.92)				(1.46)	(2.22)		
deltarecomd	0.079	-0.214	0.548	-1.921	divomit	0.380	0.510	0.283	0.369
	(1.90)	(-8.41)				(1.45)	(1.79)		
mom6yrseas	0.116	-0.135	0.517	-0.893	strev	0.478	0.068	1.375	0.340
	(1.74)	(-3.33)				(3.18)	(1.08)		

This table reports the CAPMB alphas of the long-short portfolios built on bond characteristics. Bond characteristics are defined in Table A3. ¶ indicates that we multiply the characteristic by minus one. Each month, we select the top and bottom 20% of bonds in terms of bond characteristics and form a long-short strategy. Gross alphas are before transaction costs and net alphas are after costs. The costs are calculated using the optimal trade size which maximizes the net alpha. If we fail to find the optimum, then we use the size of $2 million. Values in parentheses are t-statistics adjusted for Newey and West (1987) 12 lags. Annualized IR is the information ratio defined by the ratio of alpha to the standard deviation of the CAPM regression residuals times square root of 12. The sample period is August 2002 through November 2022.

Table 7: Performance of ML Strategies

Signal	Excess Returns Gross	Excess Returns Net Optimal	CAPMB α Gross	CAPMB α Net Optimal	Information Ratio Gross	Information Ratio Net Optimal	Optimal Volume	Turnover (%)
NN	0.531	0.110	0.430	0.039	1.192	0.177	5000	49.10
	(3.69)	(1.71)	(2.97)	(0.67)				
XT	0.548	0.166	0.393	0.042	0.901	0.118	2000	39.71
	(3.40)	(1.39)	(2.48)	(0.40)				
RF	0.387	0.056	0.239	-0.033	0.595	-0.125	10000	32.58
	(3.33)	(0.78)	(2.16)	(-0.49)				
ENET	0.535	0.129	0.422	0.041	0.980	0.121	2000	48.67
	(3.82)	(1.40)	(2.75)	(0.44)				
RIDGE	0.567	0.177	0.504	0.122	1.371	0.425	2000	46.22
	(4.07)	(1.86)	(3.77)	(1.41)				
LASSO	0.517	0.093	0.414	0.019	0.981	0.071	5000	49.92
	(3.08)	(1.04)	(2.21)	(0.20)				
ENS	0.592	0.162	0.476	0.068	1.152	0.215	2000	49.32
	(3.60)	(1.52)	(2.90)	(0.68)				
LENS	0.575	0.159	0.479	0.077	1.226	0.258	2000	49.46
	(3.76)	(1.56)	(3.18)	(0.83)				
NENS	0.562	0.142	0.421	0.029	1.038	0.092	2000	47.67
	(3.45)	(1.33)	(2.58)	(0.30)				

This table reports the average excess returns and CAPMB alphas of the long-short portfolios built on the expected returns generated by the machine learning algorithms. Each month, we select the top and bottom 20% of bonds in terms of expected returns and form a long-short strategy. Gross returns and alphas are before transaction costs and net returns and alphas are after costs. Net costs are calculated using the optimal trade size reported in the "Optimal Volume" column. Information ratio is the ratio of the CAPM alpha to the standard deviation of the residual of the regression times the square root of 12. Optimal Volume is the transaction size is in thousand dollars. Turnover is the monthly turnover rate averaged over the two legs of the strategy. Values in parentheses are t-statistics adjusted for Newey and West (1987) 12 lags. The sample period is August 2002 to November 2022.

Table 8: Performance of ML Strategies: Duration-Adjusted Returns

	Dur-Adj. Returns		CAPMB α		Information Ratio			
Signal	Gross	Net Optimal	Gross	Net Optimal	Gross	Net Optimal	Optimal Volume	Turnover (%)
NN	0.495	0.102	0.449	0.063	1.241	0.229	2000	56.00
	(3.76)	(1.34)	(3.38)	(0.83)				
XT	0.542	0.144	0.469	0.077	1.073	0.214	2000	39.74
	(3.72)	(1.37)	(3.27)	(0.76)				
RF	0.394	0.071	0.330	0.021	0.825	0.082	10000	32.61
	(4.36)	(1.33)	(3.60)	(0.36)				
ENET	0.490	0.102	0.445	0.060	1.025	0.171	2000	48.71
	(3.96)	(1.26)	(3.27)	(0.67)				
RIDGE	0.514	0.162	0.497	0.139	1.379	0.471	2000	46.27
	(3.79)	(1.70)	(3.88)	(1.59)				
LASSO	0.493	0.101	0.454	0.065	1.143	0.220	2000	57.06
	(3.30)	(1.08)	(2.84)	(0.67)				
ENS	0.552	0.139	0.507	0.097	1.242	0.296	2000	49.36
	(3.64)	(1.38)	(3.38)	(0.97)				
LENS	0.529	0.145	0.492	0.105	1.268	0.337	2000	49.51
	(3.63)	(1.45)	(3.49)	(1.11)				
NENS	0.553	0.133	0.490	0.074	1.221	0.234	2000	47.71
	(3.84)	(1.43)	(3.43)	(0.82)				

This table reports the average duration-adjusted returns and CAPMB alphas of the long-short portfolios built on the expected returns generated by the machine learning algorithms. Each month, we select the top and bottom 20% of bonds in terms of expected returns and form a long-short strategy. Gross returns and alphas are before transaction costs and net returns and alphas are after costs. Net costs are calculated using the optimal trade size reported in the "Optimal Volume" column. Information ratio is the ratio of the CAPM alpha to the standard deviation of the residual of the regression times the square root of 12. Optimal Volume is the transaction size is in thousand dollars. Turnover is the monthly turnover rate averaged over the two legs of the strategy. Values in parentheses are t-statistics adjusted for Newey and West (1987) 12 lags. The sample period is August 2002 to November 2022.

Table 9: Decomposition of Transaction Costs

	$100K			Optimal Volume			KPP
Signal	Total	BidAsk	Delay	Total	BidAsk	Delay	Cost
NN	0.689	0.582	0.107	0.390	0.144	0.246	0.258
XT	0.539	0.463	0.076	0.351	0.188	0.163	0.174
RF	0.678	0.620	0.058	0.272	0.065	0.208	0.185
ENET	0.651	0.552	0.098	0.381	0.213	0.168	0.223
RIDGE	0.614	0.507	0.108	0.382	0.195	0.187	0.209
LASSO	0.724	0.615	0.109	0.395	0.147	0.248	0.262
ENS	0.672	0.559	0.114	0.408	0.218	0.191	0.225
LENS	0.663	0.554	0.110	0.402	0.215	0.187	0.226
NENS	0.638	0.539	0.099	0.393	0.210	0.183	0.216

This table reports the two components of transaction costs: BidAsk costs and Delay costs. The total costs are the difference between gross α and net α. To measure delay costs, we compute an alternative version of net returns using quote prices on TRACE transaction dates to compute all returns. Delay costs are the difference between gross returns and the alternative net returns. Bid-ask costs are the difference between the alternative net returns and the (original) net returns. KPP costs are computed following Kelly et al. (2021) assuming the half spread of 19 bps.

Table 10: Optimal Volume and Transaction Costs

		Gross α			Gross α			Avg.
		Low	Middle	High	Low	Middle	High	Turnover
		Panel A. Optimal Vol. ($ mil.)			Panel B. Total Cost (%)			
Turnover	Low	3.93	6.75		0.014	0.003		6.67
	Middle	6.13	4.78	2.60	0.012	0.060	0.106	12.44
	High	7.43	5.33	4.80	-0.001	0.058	0.284	46.83
		Panel C. Half-Spread Cost (%)			Panel D. Delay Cost (%)			
Turnover	Low	0.021	0.016		-0.007	-0.013		6.67
	Middle	0.020	0.044	0.058	-0.008	0.016	0.049	12.44
	High	0.008	0.052	0.129	-0.009	0.006	0.155	46.83
Average α		0.049	0.136	0.354	0.049	0.136	0.354	

This table reports the average of optimal transaction volume (Panel A), the total cost at the optimum (Panel B), the half-spread cost (Panel C), and the delay cost (Panel D) using the 114 strategies for which we can find an optimal trade size. The signals/strategies are classified into three groups based independently on their turnover rate (calculated using the minimum trade size) and gross CAPMB α. The cutoff value of gross α is $[0,0.1)$, $[0.1,0.2)$, $[0.2,\infty)$ for low, middle, and high, respectively.

Table 11: Robustness: Long-Only Strategies

Signal	Excess Returns		CAPM α		Information Ratio		Optimal Volume	Turnover (%)
	Gross	Net Optimal	Gross	Net Optimal	Gross	Net Optimal		
NN	0.558	0.342	0.234	0.087	1.128	0.575	2000	55.96
	(2.91)	(2.29)	(2.65)	(1.61)				
XT	0.593	0.405	0.227	0.085	0.891	0.390	1000	42.03
	(2.86)	(2.20)	(2.79)	(1.29)				
RF	0.501	0.338	0.145	0.044	0.609	0.224	2000	41.23
	(2.61)	(2.09)	(2.32)	(0.80)				
ENET	0.564	0.360	0.236	0.083	1.012	0.473	1000	52.21
	(3.02)	(2.31)	(2.91)	(1.44)				
RIDGE	0.561	0.369	0.246	0.113	1.170	0.708	2000	46.22
	(2.96)	(2.40)	(3.07)	(2.00)				
LASSO	0.547	0.325	0.220	0.065	0.921	0.352	2000	57.02
	(2.83)	(2.11)	(2.14)	(0.98)				
ENS	0.594	0.371	0.251	0.093	1.066	0.550	2000	49.32
	(2.93)	(2.29)	(2.73)	(1.51)				
LENS	0.572	0.356	0.245	0.091	1.100	0.574	2000	49.46
	(2.95)	(2.29)	(2.87)	(1.59)				
NENS	0.598	0.391	0.241	0.085	1.047	0.464	1000	51.04
	(2.90)	(2.24)	(2.66)	(1.34)				

This table reports the average excess returns (in excess of T-bill rates) and CAPMB alphas of the long-only portfolios built on the expected returns generated by the machine learning algorithms. Each month, we select the top 20% of bonds in terms of expected returns and form a long-only strategy. Gross returns and alphas are before transaction costs and net returns and alphas are after costs. Net costs are calculated using the optimal trade size reported in the "Optimal Volume" column. Information ratio is the ratio of the CAPM alpha to the standard deviation of the residual of the regression times the square root of 12. Optimal Volume is the transaction size is in thousand dollars. Turnover is the monthly turnover rate averaged over the two legs of the strategy. Values in parentheses are t-statistics adjusted for Newey and West (1987) 12 lags. The sample period is August 2002 to November 2022.

Table 12: Robustness: Smaller Number of Bonds

Signal	Excess Returns Gross	Excess Returns Net Optimal	CAPM α Gross	CAPM α Net Optimal	Information Ratio Gross	Information Ratio Net Optimal	Optimal Volume	Turnover (%)
NN	1.104	0.393	0.962	0.296	1.541	0.619	2000	68.59
	(4.75)	(2.73)	(4.16)	(2.09)				
XT	1.290	0.505	1.154	0.389	1.002	0.407	1000	73.11
	(3.85)	(1.94)	(3.92)	(1.53)				
RF	1.171	0.446	1.149	0.406	1.284	0.553	1000	63.63
	(4.47)	(2.36)	(3.98)	(2.17)				
ENET	0.923	0.282	0.772	0.154	1.007	0.278	2000	68.02
	(4.36)	(1.69)	(3.92)	(1.05)				
RIDGE	1.206	0.510	1.072	0.408	1.204	0.628	2000	62.95
	(3.66)	(2.06)	(3.62)	(1.91)				
LASSO	1.145	0.303	1.043	0.186	1.072	0.272	1000	74.30
	(3.24)	(1.24)	(2.57)	(0.77)				
ENS	1.365	0.574	1.214	0.450	1.256	0.612	1000	69.06
	(3.94)	(2.21)	(3.75)	(1.96)				
LENS	1.208	0.497	1.046	0.374	1.215	0.596	2000	65.76
	(3.81)	(2.09)	(3.57)	(1.80)				
NENS	1.388	0.592	1.259	0.486	1.310	0.644	2000	66.74
	(4.19)	(2.61)	(4.18)	(2.32)				

This table reports the average excess returns and CAPMB alphas of the long-short portfolios built on the expected returns generated by the machine learning algorithms. Each month, we select the top and bottom 2% of bonds in terms of expected returns and form a long-short strategy. Gross returns and alphas are before transaction costs and net returns and alphas are after costs. Net costs are calculated using the optimal trade size reported in the "Optimal Volume" column. Information ratio is the ratio of the CAPM alpha to the standard deviation of the residual of the regression times the square root of 12. Optimal Volume is the transaction size is in thousand dollars. Turnover is the monthly turnover rate averaged over the two legs of the strategy. Values in parentheses are t-statistics adjusted for Newey and West (1987) 12 lags. The sample period is August 2002 to November 2022.

Table 13: Robustness: Subsample of Liquid Bonds

	Excess Returns		CAPM α		Information Ratio			
Signal	Gross	Net Optimal	Gross	Net Optimal	Gross	Net Optimal	Optimal Volume	Turnover (%)
Panel A. Above-Median Trade Counts								
NN	0.543	0.172	0.434	0.073	1.038	0.206	2000	66.87
	(4.04)	(1.77)	(3.49)	(0.88)				
XT	0.559	0.219	0.413	0.081	0.853	0.191	2000	50.36
	(3.49)	(1.66)	(2.69)	(0.70)				
RF	0.409	0.093	0.293	0.003	0.668	0.008	10000	40.38
	(3.51)	(1.10)	(2.38)	(0.03)				
ENET	0.572	0.210	0.450	0.101	0.943	0.233	2000	58.41
	(4.19)	(1.91)	(3.26)	(0.95)				
RIDGE	0.561	0.213	0.508	0.162	1.110	0.396	2000	55.25
	(4.06)	(2.01)	(3.59)	(1.61)				
LASSO	0.561	0.162	0.453	0.060	0.928	0.147	2000	66.93
	(3.40)	(1.32)	(2.64)	(0.51)				
ENS	0.603	0.204	0.503	0.106	1.065	0.263	2000	59.40
	(4.00)	(1.82)	(3.26)	(1.03)				
LENS	0.586	0.199	0.502	0.118	1.081	0.292	2000	58.99
	(4.03)	(1.83)	(3.35)	(1.15)				
NENS	0.596	0.213	0.462	0.091	1.008	0.228	2000	57.59
	(3.81)	(1.75)	(3.05)	(0.84)				
Panel B. Above-Median Number of Trade Days								
NN	0.518	0.156	0.404	0.052	0.984	0.150	2000	67.84
	(3.87)	(1.56)	(3.33)	(0.63)				
XT	0.530	0.203	0.381	0.062	0.795	0.147	2000	51.31
	(3.34)	(1.54)	(2.50)	(0.54)				
RF	0.392	0.098	0.274	0.012	0.632	0.039	10000	41.61
	(3.50)	(1.21)	(2.34)	(0.16)				
ENET	0.552	0.198	0.430	0.086	0.917	0.208	2000	59.34
	(4.08)	(1.83)	(3.11)	(0.84)				
RIDGE	0.534	0.197	0.477	0.140	1.059	0.347	2000	56.39
	(3.80)	(1.79)	(3.34)	(1.36)				
LASSO	0.541	0.155	0.432	0.052	0.891	0.128	2000	67.82
	(3.32)	(1.28)	(2.56)	(0.45)				
ENS	0.579	0.192	0.476	0.090	1.019	0.227	2000	60.36
	(3.81)	(1.67)	(3.09)	(0.87)				
LENS	0.559	0.162	0.472	0.098	1.036	0.324	10000	47.55
	(3.87)	(1.80)	(3.20)	(1.12)				
NENS	0.573	0.201	0.434	0.071	0.960	0.183	2000	58.71
	(3.66)	(1.63)	(2.91)	(0.68)				

This table reports the average excess returns and CAPMB alphas of the long-short portfolios built on the expected returns generated by the machine learning algorithms. Each month, we first divide bonds into two categories based on the trade counts and the number of trade days in the previous month. Within the bonds in the high category, we select the top and bottom 20% of bonds in terms of expected returns and form a long-short strategy. The definitions of each column can be found in the notes to Table 7. Values in parentheses are t-statistics adjusted for Newey and West (1987) 12 lags. The sample period is August 2002 to November 2022.

Table 14: Summary Statistics for Corporate Bond Mutual Funds

	N	Mean	Std.	p1	p10	p50	p90	p99
Panel A: Fund Characteristics								
Fund TNA ($ millions)	53,213	574.6	1574	11.15	19.17	125.1	1283	6525
Fund NAV ($ millions)	53,213	11.17	5.483	3.396	6.658	10.46	14.57	32.80
(Annual) expense ratio (%)	44,684	0.903	0.434	0.200	0.455	0.807	1.590	1.903
(Annual) turnover (%)	44,695	119.0	116.6	10.73	21.92	79.96	255.4	530.3
Panel B: Cross-Section of Fund Performance								
(Monthly) Excess gross return (%)	485	0.33	0.26	-0.17	0.06	0.31	0.57	1.18
(Monthly) Excess net return (%)	485	0.26	0.26	-0.24	0.01	0.24	0.50	1.06
(Monthly) Gross alpha (%)	485	0.05	0.25	-0.31	-0.11	0.04	0.16	1.23
(Monthly) Net alpha (%)	485	0.03	0.24	-0.33	-0.14	0.02	0.15	1.14
$MKTB_{Net}$ beta	485	0.67	0.30	-0.23	0.39	0.64	0.97	1.47
$MKTB_{Net}$ R^2	485	0.76	0.19	0.00	0.51	0.80	0.94	0.98

This table reports time-series averages of cross-sectional summary statistics for various fund characteristics in Panel A. Panel B reports average fund performance statistics for the cross-section of corporate bond mutual funds. The monthly gross (net) alpha is computed from time-series regressions of each funds excess gross (net) return on the gross and net of fees bond market factor, MKT_{Gross} (MKT_{Net}). The sample period is August 2002 through to December 2022 (245 Months) consisting of 485 bond mutual funds.

Table 15: Corporate Bond Value Added ($\widehat{S}_i$)

Panel A: Cross-Sectional Weighted Value-Add

	Equal weights	Time weights	Expense weights
Value-add ($\widehat{S}_i$)	-0.396	-0.300	-0.341
Standard error	0.104	0.113	0.077
t-statistic	(-3.83)	(-2.66)	(-4.44)

Panel B: Cross-Sectional Percentiles

	p1	p10	p50	p90	p99	% $\widehat{S}_i < 0$
Value-add ($\widehat{S}_i$)	-9.122	-1.125	-0.060	0.141	2.634	75.05

This table reports the average monthly value-add, $\widehat{S}_i$, defined as the total lagged inflation adjusted assets of each fund multiplied by the difference between the funds gross return and the gross return of the passive benchmark. The average cross-sectional mean of the value-add is computed with equal weights (Column 1), time weights (Column 2) and expense ratio weights (Column 3). We report standard errors and the associated t-statistic below the mean. Panel B reports the percentiles of the cross-sectional distribution of $\widehat{S}_i$ and the percentage of funds that generate a negative value-add. Numbers are reported in US$ millions per month. The sample period is August 2002 through to December 2022 (245 Months) consisting of 485 bond mutual funds.

Table 16: Impact of Luck on Performance

Panel A. Proportion of Unskilled and Skilled Funds

	Zero alpha ($\widehat{\pi}_0$)	Non-zero alpha	Unskilled ($\widehat{\pi}_A^-$)	Skilled ($\widehat{\pi}_A^+$)
Proportion (%)	76.45 [3.49]	23.55	8.07 [2.49]	15.48 [2.24]
Number	371	114	39	75

Panel B. Impact of Luck in the Left and Right Tails

	Left tail				Right tail				
Signif. level (γ)	0.05	0.10	0.15	0.20	0.20	0.15	0.10	0.05	Signif. level (γ)
Signif. $\widehat{S}_\gamma^-$(%)	7.01	7.84	9.07	10.93	14.02	12.37	9.90	6.39	Signif. $\widehat{S}_\gamma^+$(%)
	[1.16]	[1.22]	[1.30]	[1.42]	[1.58]	[1.50]	[1.36]	[1.11]	
Unlucky $\widehat{F}_\gamma^-$(%)	1.91	3.82	5.73	7.64	7.64	5.73	3.82	1.91	Lucky $\widehat{F}_\gamma^+$(%)
	[0.09]	[0.17]	[0.26]	[0.35]	[0.35]	[0.26]	[0.17]	[0.09]	
Unskilled $\widehat{T}_\gamma^-$(%)	5.10	4.01	3.34	3.28	6.38	6.64	6.07	4.48	Skilled $\widehat{T}_\gamma^+$(%)
	[1.19]	[1.28]	[1.41]	[1.57]	[1.61]	[1.57]	[1.46]	[1.25]	
Alpha (% year)	-1.99	-1.94	-1.96	-1.97	2.03	2.06	2.17	2.30	Alpha (% year)

This table reports the estimated proportions of zero-alpha, unskilled and skilled funds ($\widehat{\pi}_0$,$\widehat{\pi}_A^-$,$\widehat{\pi}_A^+$) for the population of our 'Corporate Bond' specific mutual funds ($N = 485$) from August 2002 through to December 2022 (245 Months). The fund alphas are computed for each fund using net of fees excess returns and the single-factor $MKTB_{Net}$ bond market factor. Panel B counts the proportions of significant funds in the left and right tails of the cross-sectional distribution of fund alphas ($\widehat{S}_\gamma^-$,$\widehat{S}_\gamma^+$) at four pre-defined significance levels (γ = 0.05, 0.10, 0.15, 0.20). The columns on the left decompose the proportion of significantly negative fund alphas into unlucky and unskilled funds ($\widehat{F}_\gamma^-$,$\widehat{T}_\gamma^-$). The columns on the right decompose the proportion of significantly positive fund alphas into lucky and skilled funds ($\widehat{F}_\gamma^+$,$\widehat{T}_\gamma^+$). The final row of the table present the average alpha in the left and right tail of the cross-sectional distribution of fund alphas. Standard errors are presented in square brackets.

A Inventory and Dynamic Portfolio Choice

Because the choice of a return depends on the investor's past bond holdings, we must keep track of her inventory. To do this, introducing some notation is useful. A month $t+1$ return on a bond is characterized by the investor's actions at the end of months t and $t+1$. Let I_t be the inventory (or existing short position) at the end of month t, x_t be the signal which is either Y (i.e., take a position on the bond) or N (i.e., do not take a position), $f(I_t, x_t) = \{b, h, s\}$ be an action function at the end of month t to start the trade, and $g(x_{t+1}) = \{b, h, s\}$ be the function in $t+1$ to close it. Thus, the selected returns based on these actions are expressed as $R_{t+1}^{f \cdot g}$.

The trading process is shown in Figure 12. It can be summarized as follows:

1. At the end of month t, the investor receives the signal x_t and receives the inventory I_t. She then decides whether to take a position on a bond (Y) or not (N) using the function $f(x_t, I_t)$.

2. Her order is sent to the dealers and executed if possible.

3. At the end of the month $t+1$, she receives the signal x_{t+1}. After observing it, she decides whether or not to keep the existing position, as encoded by the function $g(x_{t+1})$.

4. Her order is sent to the dealers and executed if possible. The result determines her return for month $t+1$, $R_{t+1}^{f \cdot g}$.

5. The result of the previous two order executions determines her inventory level I_{t+1}. Given x_{t+1} and I_{t+1}, we return to step 1 to compute a return in month $t+2$.

This procedure explicitly accounts for delays in order execution. The action $f(x_t, I_t)$ is executed either at the end of month t or sometime in month $t+1$. If the trade does not occur, the return and inventory are adjusted accordingly at the end of $t+1$. Similarly, the action $g(x_{t+1})$ is executed either at the end of month $t+1$ or sometime in month $t+2$. As long as the execution occurs during this period, the bond is recorded in the inventory record I_{t+1} as if the transaction were executed at the end of month $t+1$. We adjust for any excess holding costs by charging the risk-free rate until the action in month $t+2$ is taken. If the trade is not executed in month $t+2$, it is added to I_{t+1}

as unintended inventory. Therefore, the result of the order execution in both months t and $t+1$ together determines the inventory level in $t+1$. This in turn influences the next month's action $f(x_{t+1}, I_{t+1})$.

Figure 12: Flow Chart To Compute Net Returns

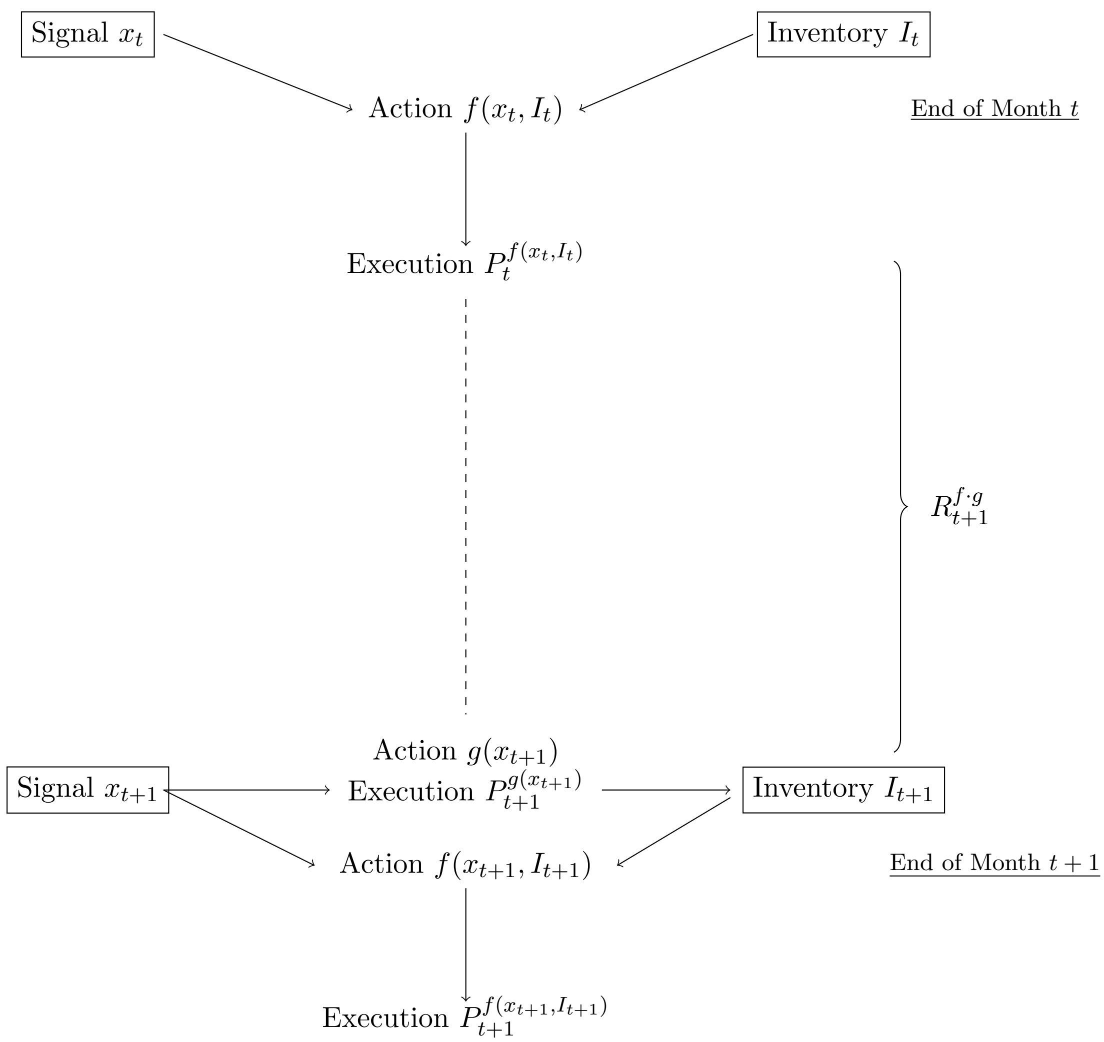

To concretely describe the set of actions in each month, we consider seven bonds as shown in Table A1. Panel A describes the action function $f(I_t, x_t)$ and $g(x_{t+1})$ for a long position. In this case, a possible action in month t is either to maintain the previous long position (h) or to buy a bond (b). The action depends not only on the month-t signal, but also on the inventory of bonds held from the previous months. If the signal is 'Y' and the inventory is also 'Y' (bonds A and B), the action is to maintain the existing position (h). On the other hand, if the signal and the inventory pair is (Y, N), as for bonds D and E, the action is to buy the bond (b). There are cases (bonds F and G) where the signal is 'N' but the inventory is 'Y' because the sales were not executed in month t. In this case, the investor's initial action is to hold the long position (h). It is important to realize that month-t action depends only on the signal at that time and inventory, not how the investor ended up with the inventory (intentional or unintentional). The distinction between intentional and unintentional inventory only affects the return calculation at the end, because unintended inventory must be financed individually by risk-free lending and borrowings.

Table A1: Return Computation

Panel A: Long Position

Time (end of month)		A	B	C	D	E	F	G
Inventory I_t		Y	Y	N	N	N	Y	Y
Unintended inventry?							Yes	Yes
t	Buy signal x_t	Y	Y	N	Y	Y	N	N
	Action $f(I_t, x_t)$	h	h	-	b	b	h	h
$t+1$	Buy signal x_{t+1}	Y	N		Y	N	N	Y
	Action $g(x_{t+1})$	h	s		h	s	s	h
Return in $t+1$		R^{hh}	R^{hs}	0	R^{bh}	R^{bs}	$R^{hs} - R^f$	$R^{hh} - R^f$

Panel B: Short Position

Time (end of month)		A	B	C	D	E	F	G
Existing position I_t		Y	Y	N	N	N	Y	Y
Unintended position?							Yes	Yes
t	Sell signal x_t	Y	Y	N	Y	Y	N	N
	Action $f(I_t, x_t)$	h	h	-	s	s	h	h
$t+1$	Sell signal x_{t+1}	Y	N		Y	N	N	Y
	Action $g(x_{t+1})$	h	b		h	b	b	h
Return in $t+1$		$-R^{hh}$	$-R^{hb}$	0	$-R^{sh}$	$-R^{sb}$	$-R^{hb} + R^f$	$-R^{hh} + R^f$

At the end of month $t+1$, the investor tries to either close the long position (s) or hold it (h). This intended action function $g(x_{t+1})$ is simple in that it depends only on the signal of month $t+1$.

For a long position, signal Y corresponds to no action (h), while signal N corresponds to intended sales (s). If the intended purchases and sales did not materialize due to excessive delays, they are reflected in the month $t+1$ inventory, I_{t+1}, and influence the next month's action.

Panel B of Table A1 describes the same descriptions for short positions. These actions can be obtained by simply replacing b in Panel A with s.

B Net of Fees Corporate Bond Market Factor

We risk-adjust our net of cost strategies with a realistic corporate bond market factor that combines tradable passively managed investment grade and high yield exchange traded funds (ETFs). We source the BlackRock iShares iBoxx Investment Grade (ticker: LQD) and High Yield (ticker: HYG) ETF net returns from the CRSP Mutual Funds database as provided by WRDS. The LQD ETF has an inception date of 2002:06 which spans the full length of our out-of-sample period. The HYG inception date is 2007:03. To address the shorter sample period for HYG, we source high yield gross return data from the Bloomberg-Barclays (BB) High-Yield bond index. Thereafter, we estimate a simple OLS regression of the HYG net returns on the BB gross returns such that we can extrapolate values for HYG before 2007:03,

$$\begin{aligned} R_{HYG,t} &= \beta_0 + \beta_{BB} \cdot R_{BB,t} + \varepsilon_t, \\ \widehat{R_{HYG,t}} &= \underset{(-2.010)}{-0.095} + \underset{(60.13)}{0.883} \cdot R_{BB,t}, \end{aligned}$$

where $R_{HYG,t}$ and $R_{BB,t}$ are the net of cost and gross returns of the HYG ETF and BB High-Yield bond index over the sample period 2007:03–2023:06 ($T = 251$). The intercept, β_0 is estimated at -9.5 basis points (statistically significant from zero at the 5% nominal level), which captures the fact that HYG is adversely impacted by trading costs and ETF fees. From the OLS estimation above, we set the net return value of the HYG index to $\widehat{R_{HYG}}$ before 2007:03 and to the actual net return of the HYG index thereafter. We denote this return R_{HYG}.

To generate the $MKTB_{Net}$ factor, we require appropriate weights for the representative investor to apportion their funds between HYG and LQD. To do this, we source *all* bonds that are included in the Bank of America Merrill Lynch Investment Grade (C0A0) and High Yield (H0A0) corporate indices and compute their respective market capitalizations (Clean Price × Units Outstanding). The weight for each index for each month is simply the sum of the respective index market capitalization at month t divided by the total market capitalization. On average, over the sample period, the investor apportions 19.90% to the high yield index and 80.10% to the investment grade index. Finally, the $MKTB_{Net}$ factor is computed as,

$$R^{Net}_{MKTB,t+1} = (R_{HYG,t+1} \cdot \omega_{HYG,t} + R_{LQD,t+1} \cdot \omega_{LQD,t}) - R_{f,t+1},$$

where $\omega_{HYG,t}$ is the weight in the HYG ETF, $\omega_{LQD,t}$ is the weight in the LQD ETF and $R_{f,t+1}$ is the one-month risk-free rate of return from Kenneth French's webpage.

We report summary statistics for the $MKTB_{Net}$, $MKTB_{Gross}$ (computed using the same weights as above with the Bloomberg-Barclays Investment Grade and High Yield index gross returns) and $MKTB$ available from `openbondassetpricing.com`.

Table A2: Summary statistics for the corporate bond market factor.

Panel A: Corporate bond market factor statistics			
	$MKTB_{Net}$	$MKTB_{Gross}$	$MKTB$
Mean	0.316	0.367	0.364
	(2.14)	(2.36)	(2.32)
SD	2.06	1.95	1.91
SR	0.53	0.65	0.66

Panel B: Pairwise correlations			
	$MKTB_{Net}$	$MKTB_{Gross}$	$MKTB$
$MKTB_{Net}$	1		
$MKTB_{Gross}$	0.982	1	
$MKTB$	0.973	0.992	1

Panel A reports the monthly factor means (Mean), the monthly factor standard deviations (SD), and the annualized Sharpe ratios. The $MKTB_{Net}$ factor is constructed as the weighted-average of the BlackRock iShares iBoxx Investment Grade (ticker: LQD) and High Yield (ticker: HYG) ETF net returns from the CRSP Mutual Funds database. The $MKTB_{Gross}$ factor is constructed as the weighted-average of the Bloomberg-Barclays Investment Grade and High Yield index gross returns. The $MKTB$ factor is the value-weighted bond market factor publicly available from `openbondassetpricing.com`. Panels A and B are based on the sample period 2002:08 to 2022:12 (245 months). t-statistics are in round brackets computed with the Newey-West adjustment with 12-lags.

C Variable Definitions

Table A3: List of the Corporate Bond and Stock Characteristics.

Num.	ID	Characteristic Name and Description	Reference	Source
		Panel A: Bond Characteristics Computed by the Authors		
1	ave12mspread	Rolling 12-month moving average of bond option adjusted credit spreads skipping the prior month	Elkamhi et al. (2021)	BAML/ICE
2	bondage	Bond age. The number of years the bond has been in issuance	Israel et al. (2018)	BAML/ICE
3	bondkurtosis	Bond kurtosis. Rolling bond excess kurtosis computed with a minimum amount of rolling observations equalling 12 which then expands up to 60-months	–	BAML/ICE
4	bondsize	Bond market capitalization. Computed as bond units amount outstanding multiplied by the clean price of the bond	Houweling and Van Zundert (2017)	BAML/ICE
5	coupon	Bond coupon. The annualised bond coupon payment in percent (%)	Chung et al. (2019)	BAML/ICE
6	dspread	First difference in bond option adjusted credit spread	–	BAML/ICE
7	dts	Duration-times-spread. Annualized bond duration multiplied by the bond option adjusted credit spread	Dor et al. (2007)	BAML/ICE
8	duration	Bond duration. The derivative of the bond value to the credit spread divided by the bond value, and is calculated by ICE	Israel et al. (2018)	BAML/ICE
9	faceval	Face value. The bond amount outstanding in units	Israel et al. (2018)	BAML/ICE
10	idiospread	Idiosyncratic component of bond credit spread. First, we run cross-sectional regressions of the log of bond option adjusted credit spreads onto the 3-month change in spreads, maturity and credit ratings. Thereafter we compute the idiosyncratic spread as the difference between the exponential of the fitted spread and the actual spread	Houweling and Van Zundert (2017)	BAML/ICE
11	impliedspread	Systematic component of bond credit spread. The fitted value from the cross-sectional regression described above	Houweling and Van Zundert (2017)	BAML/ICE
12	mom3mspread	Mom. 3m log(Spread). The log of the spread 3 months earlier minus current log spread	–	BAML/ICE
13	mom6	Corporate bond momentum. The sum of the last 6-months of bond returns minus the prior month	Gebhardt et al. (2005)	BAML/ICE
14	mom6ind	Corporate bond portfolio industry momentum. The sum of the last 6-months of bond portfolio returns minus the prior month. Portfolios are formed based on the Fama-French Industry 17 classification	Kelly et al. (2021)	BAML/ICE
15	mom6mspread	Mom. 6m log(Spread). The log of the spread 6 months earlier minus current log spread	–	BAML/ICE
16	mom6xrtg	Corporate bond momentum multiplied by bond rating. The sum of the last 6-months of bond returns minus the prior month multiplied by the bond's numerical rating AAA = 1, ... , D = 22	Kelly et al. (2021)	BAML/ICE
17	rating	Bond S&P rating. Bond numerical rating. AAA = 1, ... , D = 22	Kelly et al. (2021)	BAML/ICE
18	ratingxspread	Bond rating multiplied by credit spread	–	BAML/ICE
19	skew	Bond skewness. The rolling 60-month skewness of bond total returns. We require a minimum of 12 observations, once this threshold is hit, the rolling window expands upward to 60-months	Kelly et al. (2021)	BAML/ICE
20	spread	Bond option adjusted credit spread. The option adjusted spread of the bond provided by ICE	Kelly et al. (2021)	BAML/ICE
21	spreadvol	Volatility of the first difference of the bond option adjusted credit spread. Rolling period of 24-months with a minimum required observations of 12	–	BAML/ICE
22	strevb	Short-term bond reversal. Defined as the previous months bond return	–	BAML/ICE
23	tmt	Bond time to maturity	–	BAML/ICE
24	value	Bond value. Defined as the percentage difference between the actual credit spread and the fitted ("fair") credit spread for each bond	Houweling and Van Zundert (2017)	BAML/ICE

25	var	Historical 95% value-at-risk. Rolling 36-month bond total 95% value-at-risk. We require a minimum of 12 observations, once this threshold is hit, the rolling window expands upward to 36-months	Bai et al. (2019)	BAML/ICE
26	vixbeta	VIX beta. Rolling 60-month regression of bond returns on the Fama French 3-factors (Mkt-RF,SMB,HML, the default risk factor DEF, and the interest rate risk factor, $TERM$ and the first difference in the CBOE VIX and lagged VIX. The VIX beta in month t is the sum of the coefficient on VIX and lagged VIX. We require a minimum of 12 observations, once this threshold is hit, the rolling window expands upward to 60-months	Chung et al. (2019)	BAML/ICE
27	volatility	Bond return volatility. Rolling 36-month bond total return volatility. We require a minimum of 12 observations, once this threshold is hit, the rolling window expands upward to 36-months	Kelly et al. (2021)	BAML/ICE
		Panel B: Equity Characteristics Computed by the Authors		
28	booklev	Book leverage. Shareholder's equity and long-/short-term debt (`DLTTQ` + `DLCQ`) and minority interest (`MIBTQ`) minus cash and inventories (`CHEQ`), divided by shareholder's equity minus preferred stock	Kelly et al. (2021)	COMP
29	bookprc	Book-to-price. Firm Book-to-price is the sum of shareholder's equity and preferred stock divided by equity market capitalization for the issuing firm	Kelly et al. (2021)	CRSP/COMP
30	chggpat	Profitability change. The 5-year change in gross profitability	Asness et al. (2019)	COMP
31	d2d	Distance-to-default. Computed as in Bharath and Shumway (2008)	Bharath and Shumway (2008)	CRSP/COMP
32	debtebitda	Debt-to-EBITDA. Total debt (`DLTTQ` + `DLCQ`) divided by EBITDA (`SALEQ` − `COGSQ` − `XSGAQ`)	Kelly et al. (2021)	CRSP/COMP
33	eqtyvol	Equity volatility defined as the month-end value from a 180-day rolling-period	Campbell and Taksler (2003)	CRSP
34	gpat	Profitability. Sales (`REVTQ`) minus cost-of-goods-sold (`COGSQ`), divided by assets (`ATQ`)	Choi and Kim (2018)	COMP
35	marketcap	Equity market capitalization	Choi and Kim (2018)	CRSP
36	mktlev	Market leverage. Market capitalization and long-/short-term debt (`DLTTQ` + `DLCQ`) and minority interest (`MIBTQ`) and preferred stock minus cash and inventories (`CHEQ`), divided by market capitalization	Kelly et al. (2021)	CRSP/COMP
37	nime	Earnings-to-price. Net income (`NIQ`) divided by market equity	Correia et al. (2012)	CRSP/COMP
38	operlvg	Operating leverage. Sales (`SALEQ`) minus EBITDA (`SALEQ` − `COGSQ` − `XSGAQ`), divided by EBITDA	Gamba and Saretto (2013)	COMP
39	ret61	Six month stock momentum. Computed as the sum of the last 6 months of stock returns minus the prior month	Kelly et al. (2021)	CRSP/COMP
40	sprtod2d	Spread-to-Distance-to-Default. Spread-to-D2D is the option-adjusted spread, divided by one minus the CDF of the distance-to-default	Kelly et al. (2021)	CRSP/COMP
41	imsprtod2d	Implied (Systematic) Spread-to-Distance-to-Default. Spread-to-D2D is the fitted (systematic) option-adjusted spread, divided by one minus the CDF of the distance-to-default	Kelly et al. (2021)	CRSP/COMP
42	totaldebt	Total firm debt (`DLTTQ` + `DLCQ`)	Kelly et al. (2021)	COMP
43	turnvol	Turnover volatility. Turnover volatility is the quarterly standard deviation of sales (`SALEQ`) divided by assets (`ATQ`). The volatility is computed over 80 quarters, with a minimum required period of 10 quarters. Thereafter, the volatility is averaged (smoothed) over the preceding 4-quarters in a rolling fashion	Kelly et al. (2021)	CRSP/COMP
44	strev	Stock short-term reversal. Prior months stock return	–	CRSP
		Panel C: Open Source Asset Pricing Equity Characteristics		
45	abmaccruals	Abnormal Accruals	Xie (2001)	OSAP
46	accruals	Accruals	Sloan (1996)	OSAP
47	am	Total assets to market	Fama and French (1992)	OSAP
48	analystrevn	EPS forecast revision	Hawkins et al. (1984)	OSAP
49	analystvalue	Analyst Value	Frankel and Lee (1998)	OSAP

50	aop	Analyst Optimism	Frankel and Lee (1998)	OSAP
51	assetgrowth	Asset growth	Cooper et al. (2008)	OSAP
52	beta	CAPM beta	Fama and MacBeth (1973)	OSAP
53	betafp	Frazzini-Pedersen Beta	Frazzini and Pedersen (2014)	OSAP
54	betaliqityps	Pastor-Stambaugh liquidity beta	Pastor and Stambaugh (2003)	OSAP
55	betatailrisk	Tail risk beta	Kelly and Jiang (2014)	OSAP
56	bidaskspread	Bid-ask spread	Amihud and Mendelsohn (1986)	OSAP
57	bm	Book to market, original (Stattman 1980)	Stattman (1980)	OSAP
58	bmdec	Book to market using December ME	Fama and French (1992)	OSAP
59	bookleverage	Book leverage (annual)	Fama and French (1992)	OSAP
60	bpebm	Leverage component of BM	Penman et al. (2007)	OSAP
61	cash	Cash to assets	Palazzo (2012)	OSAP
62	cashprod	Cash Productivity	Chandrashekar and Rao (2009)	OSAP
63	cboperprof	Cash-based operating profitability	Ball et al. (2016)	OSAP
64	cf	Cash flow to market	Lakonishok et al. (1994)	OSAP
65	cfp	Operating Cash flows to price	Desai et al. (2004)	OSAP
66	chassetto	Change in Asset Turnover	Soliman (2008)	OSAP
67	cheq	Growth in book equity	Lockwood and Prombutr (2010)	OSAP
68	chinv	Inventory Growth	Thomas and Zhang (2002)	OSAP
69	chinvia	Change in capital inv (ind adj)	Abarbanell and Bushee (1998)	OSAP
70	chnncoa	Change in Net Noncurrent Op Assets	Soliman (2008)	OSAP
71	chnwc	Change in Net Working Capital	Soliman (2008)	OSAP
72	chtax	Change in Taxes	Thomas and Zhang (2011)	OSAP
73	compdebtiss	Composite debt issuance	Lyandres et al. (2008)	OSAP
74	compequiss	Composite equity issuance	Daniel and Titman (2006)	OSAP
75	convdebt	Convertible debt indicator	Valta (2016)	OSAP
76	coskewacx	Coskewness using daily returns	Ang et al. (2006)	OSAP
77	coskewness	Coskewness	Harvey and Siddique (2000)	OSAP
78	cpvolspread	Call minus Put Vol	Bali and Hovakimian (2009)	OSAP
79	dcpvolspread	Change in put vol minus change in call vol	An et al. (2014)	OSAP
80	debtissuance	Debt Issuance	Spiess and Affleck-Graves (1999)	OSAP
81	delbreadth	Breadth of ownership	Chen, Hong and Stein (2002)	OSAP
82	delcoa	Change in current operating assets	Richardson et al. (2005)	OSAP
83	delcol	Change in current operating liabilities	Richardson et al. (2005)	OSAP
84	delequ	Change in equity to assets	Richardson et al. (2005)	OSAP
85	delfinl	Change in financial liabilities	Richardson et al. (2005)	OSAP
86	dellti	Change in long-term investment	Richardson et al. (2005)	OSAP
87	delnetfin	Change in net financial assets	Richardson et al. (2005)	OSAP
88	deltarecomd	Change in recommendation	Jegadeesh et al. (2004)	OSAP
89	divinit	Dividend Initiation	Michaely et al. (1995)	OSAP
90	divomit	Dividend Omission	Michaely et al. (1995)	OSAP
91	divseason	Dividend seasonality	Hartzmark and Salomon (2013)	OSAP
92	divyieldst	Predicted div yield next month	Litzenberger and Ramaswamy (1979)	OSAP
93	dnoa	Change in net operating assets	Hirshleifer et al. (2004)	OSAP
94	dolvol	Past trading volume	Brennan et al. (1998)	OSAP

95	downrecomm	Down forecast EPS	Barber et al. (2001)	OSAP
96	dvolcall	Change in call vol	An et al. (2014)	OSAP
97	dvolput	Change in put vol	An et al. (2014)	OSAP
98	earnforedisp	Long-vs-short EPS forecasts	Da and Warachka (2011)	OSAP
99	earnstreak	Earnings surprise streak	Loh and Warachka (2012)	OSAP
100	earnsurpise	Earnings Surprise	Foster et al. (1984)	OSAP
101	ebm	Enterprise component of BM	Penman et al. (2007)	OSAP
102	entmult	Enterprise Multiple	Loughran and Wellman (2011)	OSAP
103	ep	Earnings-to-Price Ratio	Basu (1977)	OSAP
104	eqtydur	Equity Duration	Dechow et al. (2004)	OSAP
105	exchswitch	Exchange Switch	Dharan and Ikenberry (1995)	OSAP
106	exclexp	Excluded Expenses	Doyle et al. (2003)	OSAP
107	feps	Analyst earnings per share	Cen et al. (2006)	OSAP
108	fgr5yrlag	Long-term EPS forecast	La Porta (1996)	OSAP
109	firmage	Firm age based on CRSP	Barry and Brown (1984)	OSAP
110	foredisp	EPS Forecast Dispersion	Diether et al. (2002)	OSAP
111	fr	Pension Funding Status	Franzoni and Marin (2006)	OSAP
112	gp	gross profits / total assets	Novy-Marx (2013)	OSAP
113	grcapx	Change in capex (two years)	Anderson and Garcia-Feijoo (2006)	OSAP
114	grcapx3y	Change in capex (three years)	Anderson and Garcia-Feijoo (2006)	OSAP
115	grltnoa	Growth in long term operating assets	Fairfield et al. (2003)	OSAP
116	grsaletogrinv	Sales growth over inventory growth	Abarbanell and Bushee (1998)	OSAP
117	herf	Industry concentration (sales)	Hou and Robinson (2006)	OSAP
118	herfasset	Industry concentration (assets)	Hou and Robinson (2006)	OSAP
119	herfbe	Industry concentration (equity)	Hou and Robinson (2006)	OSAP
120	high52	52 week high	George and Hwang (2004)	OSAP
121	hire	Employment growth	Bazdresch et al. (2014)	OSAP
122	idiovol3f	Idiosyncratic risk (3 factor)	Ang et al. (2006)	OSAP
123	idiovolaht	Idiosyncratic risk (AHT)	Ali et al. (2003)	OSAP
124	illiquidity	Amihud's illiquidity	Amihud (2002)	OSAP
125	indipo	Initial Public Offerings	Ritter (1991)	OSAP
126	indmom	Industry Momentum	Grinblatt and Moskowitz (1999)	OSAP
127	intanbm	Intangible return using BM	Daniel and Titman (2006)	OSAP
128	intancfp	Intangible return using CFtoP	Daniel and Titman (2006)	OSAP
129	intanep	Intangible return using EP	Daniel and Titman (2006)	OSAP
130	intansp	Intangible return using Sale2P	Daniel and Titman (2006)	OSAP
131	intmom	Intermediate Momentum	Novy-Marx (2012)	OSAP
132	investment	Investment to revenue	Titman et al. (2004)	OSAP
133	investppeinv	Change in ppe and inv/assets	Lyandres et al. (2008)	OSAP
134	iomomcust	Customers momentum	Menzly and Ozbas (2010)	OSAP
135	iomomsupp	Suppliers momentum	Menzly and Ozbas (2010)	OSAP
136	leverage	Market leverage	Bhandari (1988)	OSAP
137	lrreversal	Long-run reversal	De Bondt and Thaler (1985)	OSAP
138	maxret	Maximum return over month	Bali et al. (2011)	OSAP
139	mom11yroff	Off season reversal years 11 to 15	Heston and Sadka (2008)	OSAP
140	mom11yrseas	Return seasonality years 11 to 15	Heston and Sadka (2008)	OSAP

141	mom12m	Momentum (12 month)	Jegadeesh and Titman (1993)	OSAP
142	mom12off	Momentum without the seasonal part	Heston and Sadka (2008)	OSAP
143	mom16yroff	Off season reversal years 16 to 20	Heston and Sadka (2008)	OSAP
144	mom16yrseas	Return seasonality years 16 to 20	Heston and Sadka (2008)	OSAP
145	mom6yroff	Off season reversal years 6 to 10	Heston and Sadka (2008)	OSAP
146	mom6yrseas	Return seasonality years 6 to 10	Heston and Sadka (2008)	OSAP
147	momoffseason	Off season long-term reversal	Heston and Sadka (2008)	OSAP
148	momseason	Return seasonality years 2 to 5	Heston and Sadka (2008)	OSAP
149	momseasshort	Return seasonality last year	Heston and Sadka (2008)	OSAP
150	momvol	Momentum in high volume stocks	Lee and Swaminathan (2000)	OSAP
151	mrreversal	Medium-run reversal	De Bondt and Thaler (1985)	OSAP
152	netdebtfin	Net debt financing	Bradshaw et al. (2006)	OSAP
153	neteqtyfin	Net equity financing	Bradshaw et al. (2006)	OSAP
154	netpoyld	Net Payout Yield	Boudoukh et al. (2007)	OSAP
155	noa	Net Operating Assets	Hirshleifer et al. (2004)	OSAP
156	numearnincr	Earnings streak length	Loh and Warachka (2012)	OSAP
157	operprof	Operating profits / book equity	Fama and French (2006)	OSAP
158	operprofrd	Operating profitability R&D adjusted	Ball et al. (2016)	OSAP
159	opleverage	Operating leverage	Novy-Marx (2011)	OSAP
160	optionvolume1	Option to stock volume	Johnson and So (2012)	OSAP
161	optionvolume2	Option volume to average	Johnson and So (2012)	OSAP
162	payoutyield	Payout Yield	Boudoukh et al. (2007)	OSAP
163	pctacc	Percent Operating Accruals	Hafzalla et al. (2011)	OSAP
164	pcttotacc	Percent Total Accruals	Hafzalla et al. (2011)	OSAP
165	pead	Earnings announcement return	Chan et al. (1996)	OSAP
166	predictedfe	Predicted Analyst forecast error	Frankel and Lee (1998)	OSAP
167	rankrevgr	Revenue Growth Rank	Lakonishok et al. (1994)	OSAP
168	rdipo	IPO and no R&D spending	Gou et al. (2006)	OSAP
169	rds	Real dirty surplus	Landsman et al. (2011)	OSAP
170	realizedvol	Realized (Total) Volatility	Ang et al. (2006)	OSAP
171	residmom	Momentum based on FF3 residuals	Blitz et al. (2011)	OSAP
172	retconglom	Conglomerate return	Cohen and Lou (2012)	OSAP
173	returnskew	Return skewness	Bali et al. (2015)	OSAP
174	returnskew3f	Idiosyncratic skewness (3F model)	Bali et al. (2015)	OSAP
175	rev6	Earnings forecast revisions	Chan et al. (1996)	OSAP
176	revsurp	Revenue Surprise	Jegadeesh and Livnat (2006)	OSAP
177	rivolspread	Realized minus Implied Vol	Bali and Hovakimian (2009)	OSAP
178	roaq	Return on assets (qtrly)	Balakrishnan et al. (2010)	OSAP
179	roe	Net income / book equity	Haugen and Baker (1996)	OSAP
180	saleoverhead	Sales growth over overhead growth	Abarbanell and Bushee (1998)	OSAP
181	shareiss1y	Share issuance (1 year)	Pontiff and Woodgate (2008)	OSAP
182	shareiss5y	Share issuance (5 year)	Daniel and Titman (2006)	OSAP
183	sharerepo	Share repurchases	Ikenberry et al. (1995)	OSAP
184	shortinterest	Short Interest	Dechow et al. (2001)	OSAP
185	skew1	Volatility smirk near the money	Xing et al. (2010)	OSAP
186	smileslope	Put volatility minus call volatility	Yan (2011)	OSAP

187	sp	Sales-to-price	Barbee et al. (1996)	OSAP
188	spinoff	Spinoffs	Cusatis et al. (1993)	OSAP
189	tax	Taxable income to income	Lev and Nissim (2004)	OSAP
190	totalaccruals	Total accruals	Richardson et al. (2005)	OSAP
191	trendfactor	Trend Factor	Han et al. (2016)	OSAP
192	uprecomm	Up Forecast	Barber et al. (2001)	OSAP
193	varcf	Cash-flow to price variance	Haugen and Baker (1996)	OSAP
194	volmkt	Volume to market equity	Haugen and Baker (1996)	OSAP
195	volsd	Volume Variance	Chordia et al. (2001)	OSAP
196	volumetrend	Volume Trend	Haugen and Baker (1996)	OSAP
197	xfin	Net external financing	Bradshaw et al. (2006)	OSAP
198	zerotrade	Days with zero trades	Liu (2006)	OSAP
199	ztradealt1	Days with zero trades	Liu (2006)	OSAP
200	ztradealt2	Days with zero trades	Liu (2006)	OSAP

This table presents information on the 200 characteristics we use to form our predictions from the various machine learning (ML) models we employ. Panel A reports the 27 characteristics which relate to *bond*-only characteristics which are constructed using the BAML/ICE corporate bond dataset. Panel B reports the 17 characteristics which relate to *equity-and-bond* characteristics which are constructed using CRSP and COMPUSTAT (COMP). Panel C reports the 156 characteristics which we download from `openassetpricing.com` (OSAP). We provide more detailed descriptions for the characteristics we construct from source in Panels A and B respectively. Additional resources and description notes for the `openassetpricing.com`-based data can be downloaded here.

Table A4: Net CAPM α

Turnover	Gross α (%)									
Rate (%)	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
5	0.055	0.124	0.193	0.262	0.331	0.400	0.469	0.538	0.607	0.676
10	0.065	0.116	0.167	0.219	0.270	0.322	0.373	0.424	0.476	0.527
15	0.071	0.112	0.153	0.194	0.235	0.276	0.317	0.358	0.399	0.440
20	0.075	0.109	0.142	0.176	0.210	0.244	0.277	0.311	0.345	0.379
25	0.078	0.106	0.134	0.162	0.190	0.218	0.247	0.275	0.303	0.331
30	0.081	0.104	0.128	0.151	0.174	0.198	0.221	0.245	0.268	0.292
35	0.083	0.102	0.122	0.142	0.161	0.181	0.200	0.220	0.239	0.259
40	0.085	0.101	0.117	0.133	0.149	0.166	0.182	0.198	0.214	0.230
45	0.087	0.100	0.113	0.126	0.139	0.152	0.165	0.179	0.192	0.205
50	0.088	0.099	0.109	0.120	0.130	0.140	0.151	0.161	0.172	0.182
55	0.090	0.098	0.106	0.114	0.122	0.130	0.138	0.146	0.154	0.162
60	0.091	0.097	0.102	0.108	0.114	0.120	0.126	0.132	0.137	0.143
65	0.092	0.096	0.100	0.103	0.107	0.111	0.115	0.118	0.122	0.126
70	0.093	0.095	0.097	0.099	0.101	0.103	0.104	0.106	0.108	0.110
75	0.094	0.094	0.094	0.094	0.095	0.095	0.095	0.095	0.095	0.095
80	0.095	0.093	0.092	0.090	0.089	0.087	0.086	0.085	0.083	0.082
85	0.096	0.093	0.090	0.087	0.084	0.081	0.078	0.075	0.072	0.069
90	0.097	0.092	0.088	0.083	0.079	0.074	0.070	0.065	0.061	0.056
95	0.097	0.092	0.086	0.080	0.074	0.068	0.062	0.056	0.051	0.045
100	0.098	0.091	0.084	0.077	0.070	0.062	0.055	0.048	0.041	0.034

The table reports the net CAPMB α as a function of gross α and portfolio turnover rate. The values are estimated by regressing the net CAPMB α's on gross α, portfolio turnover rate, and the product of the two. The regression uses 114 strategies for which we can find the optimal trade size. All variables are in percentage per month.

Internet Appendix to

The Low Frequency Trading Arms Race: Machines Versus Delays

(not for publication)

Abstract

This Internet Appendix presents supplementary material and results not included in the main body of the paper.

A Data and Variable Construction

The following sections describe the various databases that we use in the paper. Across all databases, we filter out bonds which have a time-to-maturity of less than 1-year. Furthermore, for consistency, across all databases, we define bond ratings as those provided by Standard & Poors (S&P). We include the full spectrum of ratings (AAA to D), but exclude bonds which are unrated. For each database that we consider, we (the authors) *do not* winsorize or trim bond returns in any way.

A.1 Corporate Bond Databases

Mergent Fixed Income Securities Database (FISD) database

Mergent Fixed Income Securities Database (FISD) for academia is a comprehensive database of publicly offered U.S. bonds. Research market trends, deal structures, issuer capital structures, and other areas of fixed income debt research.

We apply the standard filters to the FISD data as they relate to empirical asset pricing in corporate bonds,

1. Only keep bonds that are issued by firms domiciled in the United States of America, `COUNTRY_DOMICILE == 'USA'`.
2. Remove bonds that are private placements, `PRIVATE_PLACEMENT == 'N'`.
3. Only keep bonds that are traded in U.S. Dollars, `FOREIGN_CURRENCY == 'N'`.
4. Bonds that trade under the 144A Rule are discarded, `RULE_144A == 'N'`.
5. Remove all asset-backed bonds, `ASSET_BACKED == 'N'`.
6. Remove convertible bonds, `CONVERTIBLE == 'N'`.
7. Only keep bonds with a fixed or zero coupon payment structure, i.e., remove bonds with a floating (variable) coupon, `COUPON_TYPE != 'V'`.

8. Remove bonds that are equity linked, agency-backed, U.S. Government, and mortgage-backed, based on their `BOND_TYPE`.

9. Remove bonds that have a "non-standard" interest payment structure or bonds not caught by the variable coupon filter (`COUPON_TYPE`). We remove bonds that have an `INTEREST_FREQUENCY` equal to −1 (N/A), 13 (Variable Coupon), 14 (Bi-Monthly), and 15 and 16 (undocumented by FISD). Additional information on `INTEREST_FREQUENCY` is available on Page 60 of 67 of the FISD Data Dictionary 2012 document.

Bank of America Merrill Lynch (BAML) database

The BAML data is provided by the Intercontinental Exchange (ICE) and provides daily bond price quotes, accrued interest, and a host of pre-computed corporate bond characteristics such as the bond option-adjusted credit spread (OAS), the asset swap spread, duration, convexity, and bond returns in excess of a portfolio of duration-matched Treasuries. The ICE sample spans the time period 1997:01 to 2022:12 and includes constituent bonds from the ICE Bank of America High Yield (H0A0) and Investment Grade (C0A0) Corporate Bond Indices.

ICE Bond Filters. We follow Binsbergen, Nozawa, and Schwert (2023) and take the last quote of each month to form the bond-month panel. We then merge the ICE data to the filtered Mergent FISD database.

The following ICE-specific filters are then applied:

1. Only include corporate bonds, `Ind_Lvl_1 == 'corporate'`

2. Only include bonds issued by U.S. firms, `Country == 'US'`

3. Only include corporate bonds denominated in U.S. Dollars, `Currency == 'USD'`

BAML/ICE Bond Returns. Total bond returns are computed in a standard manner in ICE, and no assumptions about the timing of the last trading day of the month are made because the data is quote based, i.e., there is always a valid quote at month-end to compute a bond return. This means that each bond return is computed using a price quote at exactly the end of the month, each and every month. This introduces homogeneity into the bond returns because prices are sampled at exactly the same time each month. ICE only provides bid-side pricing, meaning bid-ask bias is inherently not present in the monthly sampled prices, returns and credit spreads. The monthly ICE return variable is (as denoted in the original database), is `trr_mtd_loc`, which is the month-to-date return on the last business day of month t. We use this return specification (in excess of the one-month risk free rate of return) and the bond returns in excess of a portfolio duration matched U.S. Treasury bond returns denoted as `ex_rtn_mtd` in the ICE dataset as the dependent variables to train the machine learning models.

Enhanced TRACE Database

TRACE provides data on corporate bond transactions. Since we measure the profitability of factor investing from an end-user perspective, we use only dealer-customer transactions (`cntra_mp_id = 'C'`). We remove trades that are i) when-issued (`wis_fl != 'Y'`), ii) locked-in (`lckd_in_ind != 'Y'`), iii) with special conditions (`sale_cndtn_cd = '@' or sale_cndtn_cd = ''`).

In addition, we restrict our sample to those with standard settlement days (`days_to_sttl_ct = '' or days_to_sttl_ct = '000' or days_to_sttl_ct = '001' or days_to_sttl_ct = '002'`).

However, some transaction records contain prices that appear to reflect clerical/recording errors. We avoid simply removing outliers in terms of prices and returns because such procedures bias the standard deviation of returns downward and inflate Sharpe ratios. Furthermore, if we simply removed very low returns, we would eliminate the bond that defaulted, leading to a spurious profitability of a strategy. To avoid these problems, we apply the reversal filter of Bessembinder et al. (2008) with a wider band. That is, we examine the log price changes of a bond using two consecutive transactions. If a product of two adjacent log price changes is less than -0.25 (i.e., a 50% decline followed by a 50% increase), then we consider the price record in the middle to be an error and remove it.

After applying these filters, we compute the average price of a bond on a day, separately for dealer buys and dealer sells. These daily averages are used to calculate net returns in the main results of the paper.

A.2 CRSP Mutual Fund Data

Data Filters We identify corporate bond mutual funds by CRSP's fund classification. In particular, we choose the subcategory 'Corporate' among 'Fixed Income' funds. We remove funds with less than 36 monthly observations and observations with total net assets (`TNA`) less than $10 million. We are careful to remove index funds (i.e., those which track a market index). Where a funds expense ratio is missing, we set it value to the cross-sectional mean of the expense ratio in month t.

Handling Data Errors After filtering, 8 observations (all from different funds) but for the same date (`2022-09-30`) have monthly gross returns (`mret`) greater than +100%. To identify if these returns are real we impute the funds return using `TNA` and document that the return is an order of magnitude smaller. We remove these observations, since they are obviously data errors.

B Machine Learning Model Estimation and Cross-Validation

For all of our machine learning models, we cross-validate the model hyperparameters every five-years and re-train the model every 12-months with an expanding window. Within each window we perform the cross-validation with a 70:30 training-validation split. For example, if we have window of 1,000 temporally ordered observations, 1-700 are used to train the model and the remaining 300 are used for validation. We graphically depict the sample splitting strategy for the training and cross-validation in Figure A.1. For all models except for the feed forward neural network we utilize the `sklearn` Python package (Pedregosa et al., 2011). We use the `tensorflow` Python package to estimate the neural network.

We report the respective sets of hyperparameters which we cross-validate over in Table A.3.

Linear Models with Penalties Panel A reports the hyperparameters for the linear models with penalties for the Lasso (LASSO), Ridge (RIDGE) and the Elastic Net (ENET). For the LASSO-style penalty, we cross-validate over 100 possible ℓ_1 penalties which change dynamically with the sample. The 100 potential ℓ_1 penalties are set by default with `sklearn` with a logarithmic scale. The maximum penalty is set to be the smallest value such that the coefficients are all set to zero. The minimum penalty is set to be 0.001 scaled by the maximum penalty. The ℓ_2 (RIDGE) penalties are defined as 100 values between 0.0001 and 1 with a logarithmic scale. The elastic net model hyperparameters are tuned with the 100 possible ℓ_1 penalties which change dynamically with the sample and a set of ℓ_1 vs. ℓ_2 ratios.

Nonlinear Tree-Based Ensembles Panel B reports the hyperparameters for the tree-based nonlinear ensemble models which includes the Random Forest (RF) and Extremely Randomized Trees (XT). For both ensemble models, we use 100 estimators (trees). We also follow Gu et al. (2020) and set the maximum tree depth to be $\in [2, 4, 6]$. Thereafter, we allow the trees to consider a maximum of 5,10,15 or 30 features (characteristics) at each split point. Finally, at each end node of the tree (final leaf), we impose a minimum of 1, 10 or 50 samples (i.e., bond returns) in each leaf.

Feed Forward Neural Network Ensemble Panel C reports the hyperparameters for the feed forward neural network (NN). We estimate a shallow network with a single layer and 32 neurons. Since our sample starts off with a relatively smaller sample size than that of Gu et al. (2020) and other work which utilizes equity data only, we set the batch size to 1024 (with batch normalization) and the number of epochs to 100. We cross-validate over the learning rate which is $\in [0.001, 0.01]$ and an ℓ_1 penalty $\in [0.001, 0.01]$. We also implement early stopping with the 'patience' parameter set to 5. The prediction variance of each individually estimated neural network is high. In order to reduce prediction variance across estimated neural network models, at each training date we estimate 10 models with different randomly assigned initial weights. In doing so, we select the best performing 5 models based on the smallest mean squared error estimated in the validation sample at that training date. This means, that at each date $t + 1$, we produce five predictions from the five best performing models estimated at the training date. The overall $t+1$ prediction is the average over these five best performing models. At each training date, we then repeat this process ten times, yielding ten ensembled predictions. The final NN prediction for each month $t+1$ is the average over these ten ensembled predictions, i.e., an ensemble over the ensemble.

C Machine Learning Model Explanatory Power

In this section we report the machine learning model out-of-sample R-square values, R^2_{OS} (Panel A) and Diebold-Mariano t-statistics (Panel B) in Table A.4.

Out-of-sample R-squared In Panel A, the R^2_{OS}'s are computed as,

$$R^2_{OS} = 1 - \frac{\sum_{i,t}(r_{i,t+1} - \sum_l \widehat{r}_{i,t+1})^2}{\sum_{i,t}(r_{i,t+1})^2},$$

where $r_{i,t+1}$ is the excess bond return over $t{:}t+1$ and $\widehat{r}_{i,t+1}$ is the machine learning model forecast available to the investor at month t for the period $t{:}t+1$.

Overall, the nonlinear ML models outperform the penalized linear regression models. The R^2_{OS} ranges from -1.71% (RIDGE) to 0.091% for the elastic net (ENET). The linear ensemble (LENS) delivers a negative R^2_{OS} of -0.14%.

In contract, the nonlinear tree ensembles comprising the extremely randomized trees (XT) and the random forests (RF) deliver R^2_{OS}s of 0.48% and 0.35% respectively. The shallow feed forward neural network (NN) marginally outperforms the XT method with an R^2_{OS} of 0.46%. Given the above, the nonlinear ensemble (NENS) performs the best amongst the considered models with an R^2_{OS} of 0.69%.

Diebold and Mariano (1995) Tests To compare the out-of-sample predictive power between two ML models we use a modified version of the Diebold and Mariano (1995) which follows Gu et al. (2020). We first compute the time-series of average forecast differences between model x and

y as the cross-sectional average of the differences in each month $t+1$,

$$d_{x,y,t+1} = \frac{1}{T}\sum_{i=1}^{N}\left((\hat{e}^{x}_{i,t+1})^2 - (\hat{e}^{y}_{i,t+1})^2\right)$$

The modified Diebold-Mariano statistic is defined as,

$$DMx, y = \frac{\overline{d}_{x,y}}{\hat{\sigma}_{\overline{d}}},$$

where $\overline{d}_{x,y}$ is the time-series average of the cross-sectional average differences between the forecast errors of the predictions from the ML model x and y and $\hat{\sigma}_{\overline{d}}$ is the Newey-West adjusted standard error.

In Panel B of Table A.4, we report the t-statistics from the above procedure for the column vs. row model. Positive t-statistics imply that the column model outperforms the row model. Confirming the results related to the R^2_{OS}, the NENS nonlinear ensemble outperforms all of the other models except for the XT and NN. The RIDGE penalized linear model is outperformed by all other models and ensembles.

D eMAXX Net Quarterly Changes

TRACE does not provide the identity of end-users and thus it is a challenge to identify who is likely to enjoy lower transaction costs with large trades. As an alternative, we investigate eMAXX institutional holding data which provides the institutional ownership of corporate bonds at the quarter end from 1998Q2 to 2021Q2.

If we assume institutions trade each bond only once in a quarter, then the absolute value of quarterly changes in positions provides information about the transaction size. Clearly, this is a strong assumption as institutions can trade multiple times spreading trades within a quarter. With this caveat in mind, we examine quarterly absolute changes in the positions of financial institutions. In doing so, we discard observations with no changes and treat non-zero changes as transactions. We also discard any position changes in the quarter in which the bond is issued or matures because such changes do not incur transaction costs. For each investor and each quarter, we compute the average transaction sizes across bonds. Then, we calculate the mean and median across institutions to arrive at the trade-size statistics.

Panel A of Figure A.2 plots the average and median transaction sizes over time in eMAXX data. The average spikes in some quarters with no obvious events and are likely to reflect measurement errors. The median has a downward trend in the period from 1998 to 2004 and remains stable since then.

Panel B presents the median within institution types, including insurance firms, mutual funds, and others. After 2004, the median transaction size is nearly unchanged at around $500,000 dollars. In TRACE data (Figure 4), a transaction with size $500,000 is at roughly the 80 percentile of the size distribution. Thus, the median eMAXX investors are relatively large and likely to pay lower transaction costs than the average TRACE investors.

Panel C shows the breakdown by portfolio sizes. Every quarter, we classify investors based on the total size (in face value) of their corporate bond portfolios at the end of the previous quarter. We then compute the median transaction size within each size quintile and plot it in Panel C.

The figure shows that, naturally, investors with a large portfolio size tend to have large quarterly changes in positions. If these position changes are implemented in one trade, then the eMAXX investors in the top quintile (whose transaction size is around $1.5 million after 2004) enjoy lower transaction costs than smaller investors.

E CRSP Corporate Bond Mutual Fund Size

In this section, we study summary statistics of the corporate bond mutual funds. In Figure A.3, we plot the number of corporate bond mutual funds and the Herfindhal index for total net assets. Neither figure shows a clear trend.

F Order Splitting

In this section, we depart from the main findings of an investor placing an order with a fixed trade size. It can be argued, based on intuition in the stock market, that an investor can break down a large order into smaller pieces to minimize costs. When dealing with bonds, it may be feasible to decrease the cost of delays by rapidly executing a portion of the order through dividing a large order, despite higher bid-ask spreads.

In this section, we allow the investor to execute a portion of their trades by following the order of opportunities that arise within a given month. To achieve this, we sort TRACE transactions in ascending order according to their trade date and time, grouped by bond, and calculate the cumulative volume for each month. When the cumulative volume for a particular bond reaches the target size (for example, $2 million), we utilize all trades conducted up to that point to compute the net returns. We classify eligible trades into 12 size groups following the same procedure as in the main analysis. For instance, suppose the target is $2 million. In that case, trades eligible in a month may include two orders of $500,000 and five orders of $200,000. We then calculate the volume-weighted average of the net return corresponding to each trade size in that month. For this example, we take 50% weight on the net return for $500,000 trades and 50% weight on the net return for $200,000 trades. This method allows us to determine the net return with order splitting. If the total monthly volume fails to reach the targeted trade size, the observation is considered a trade failure. While this creates a look-ahead bias, it is a necessary assumption to avoid the complexity of tracking partial inventory.

Figure A.5 compares the total cost before (x-axis) and after (y-axis) order splitting. We find that splitting a large order slightly increases transaction costs. As an investor deviates from the optimal trade size, she pays higher bid-ask spreads and this outweighs the benefit of executing a part of her order quickly. In summary, our main results are not impacted by order splitting.

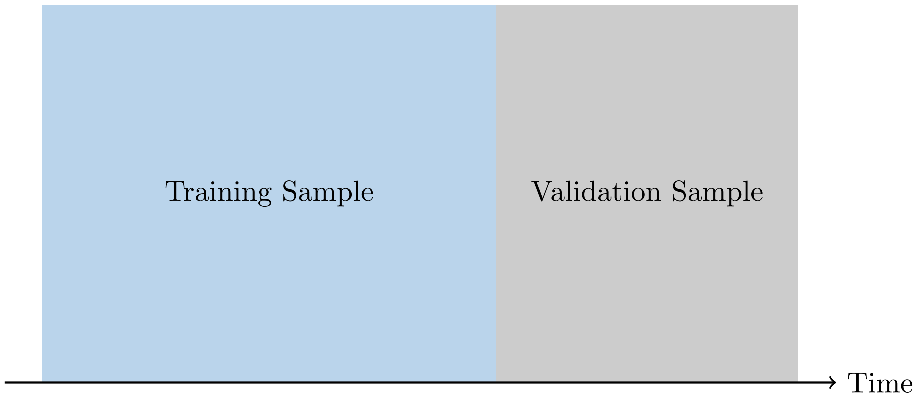

Figure A.1: Sample Splitting for Cross-Validation of Model Hyperparameters.

This figure shows the sample splitting scheme used for cross-validation of the machine learning model hyperparameters for the various machine learning models we consider. The forecasting exercise involves an expanding window that starts in January 1998. The initial window spans 1998:01–2002:07 ($T = 55$), and then expands forward each and every month until the sample end on 2022:12. The first (last) out-of-sample forecast is made in 2002:07 (2022:11) for the following month 2002:08 (2022:12). Hence, the out-of-sample ML portfolio returns commence in 2002:08 and end in 2022:12, $T = 245$. For each window, the blue area represents the training sample and the grey area represents the validation sample. The former consists of the first 70% of the observations while the latter consists of the final 30% of observations. The training and the validation samples are contiguous in time and not randomly selected in order to preserve the time series dependence of the data.

Figure A.2: Mean and Median Quarterly Changes in Positions

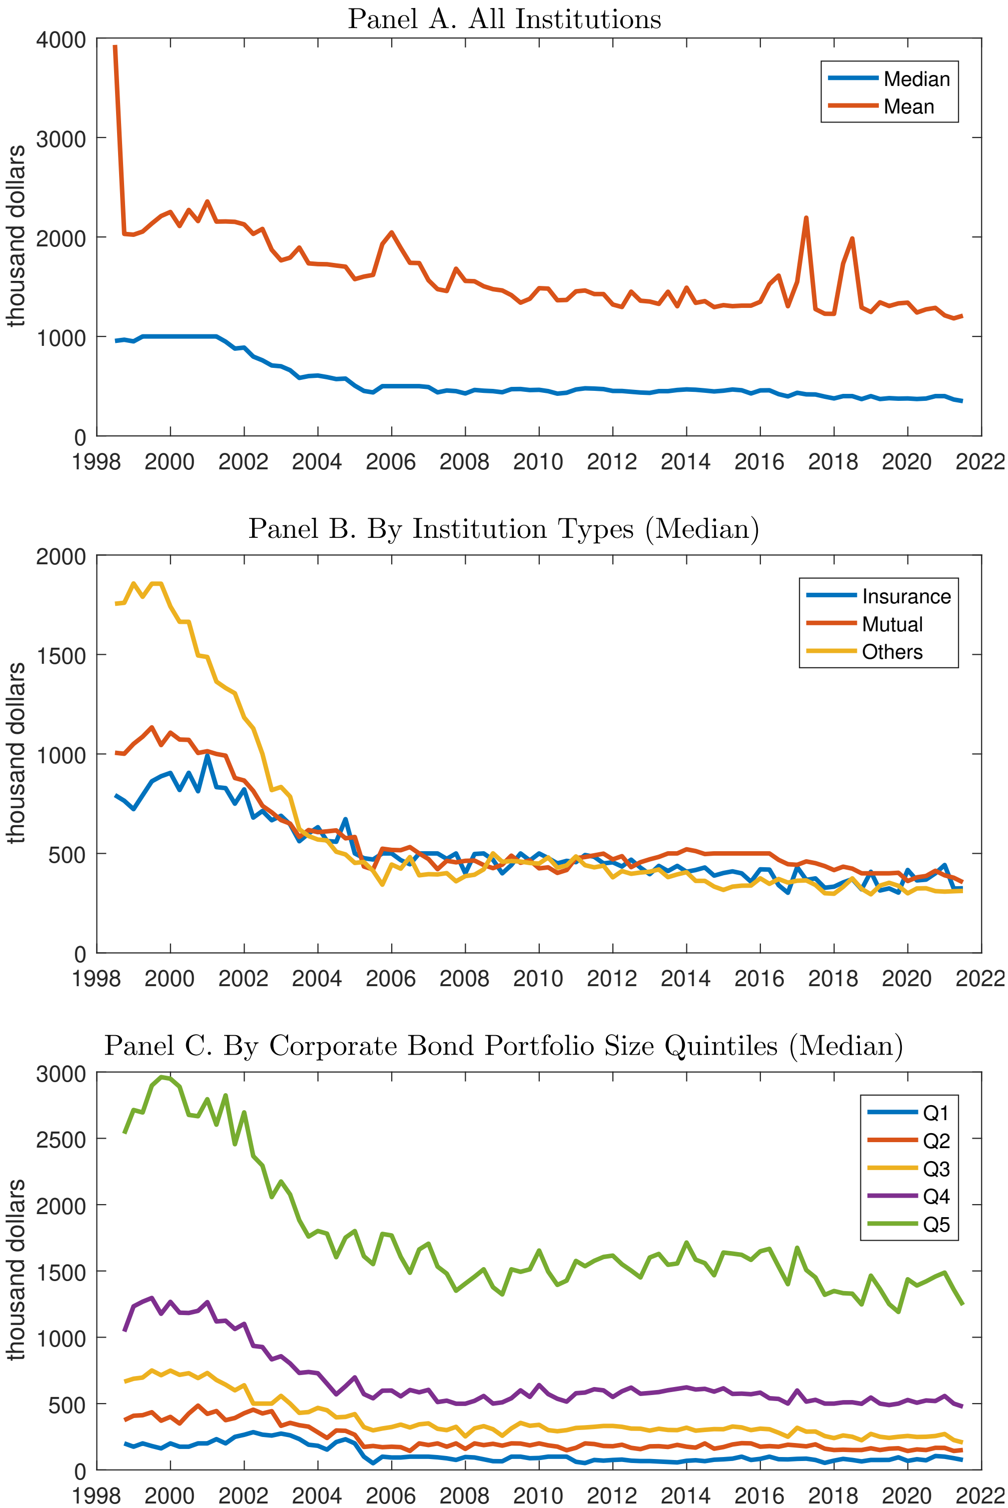

Figure A.3: Number of Corporate Bond Mutual Funds

This figure plots the number of corporate bond mutual funds (top panel) and the Herfindahl index of the total net assets (bottom panel).

Figure A.4: Unable to Find the Optimal Trade Size: Example of Bond Age-Sorted Portfolio

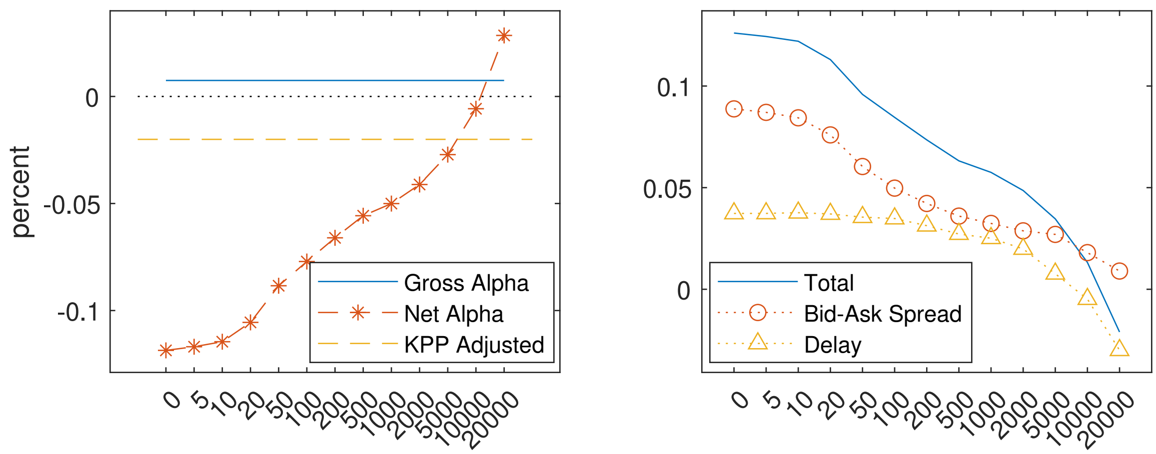

This figure plots the bond CAPM alphas of the long-short strategies based on corporate bonds' age before and after accounting for transaction costs (left panel). The transaction costs are decomposed into the bid-ask spread costs and delay costs (right panel). Values on the x-axis are the trade size in thousand dollars.

Figure A.5: Comparison of Total Transaction Costs: Order Splitting

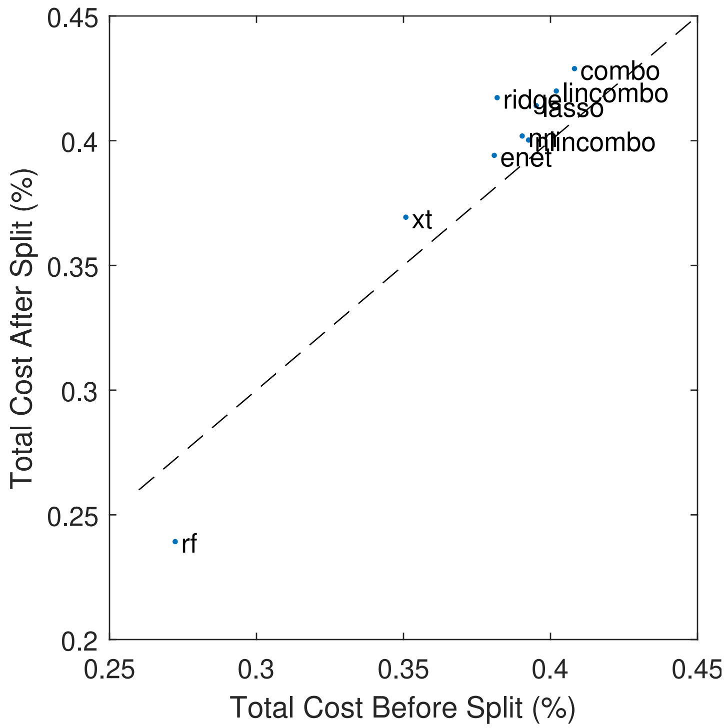

This chart displays the total cost, which is the disparity between the gross alpha and the net alpha with the most efficient trade size. The x-axis represents the fixed trade size's total cost, while the y-axis portrays the total cost when a trade is broken into pieces for quicker execution.

Table A.1: Performance of Individual Characteristics

	CAPM α		Info. Ratio			CAPM α		Info. Ratio	
Signal	Gross	Net	Gross	Net	Signal	Gross	Net	Gross	Net
bondkurtosis	0.276	0.178	0.783	0.696	feps¶	0.075	0.100	0.189	0.284
	(2.69)	(2.36)				(0.69)	(1.08)		
coskewacx	0.182	0.108	0.800	0.644	spreadvol	0.228	0.159	0.367	0.277
	(3.54)	(2.78)				(1.53)	(1.08)		
mom6mspread¶	0.454	0.178	0.884	0.554	beta	0.109	0.104	0.247	0.269
	(2.93)	(1.67)				(0.86)	(0.92)		
skew¶	0.215	0.121	0.623	0.472	nime¶	0.112	0.058	0.427	0.264
	(1.72)	(1.38)				(1.73)	(0.91)		
duration¶	0.158	0.163	0.368	0.454	betatailrisk	0.085	0.084	0.238	0.261
	(1.69)	(2.04)				(0.93)	(1.00)		
mrreversal	0.060	0.082	0.221	0.451	saleoverhead¶	0.074	0.049	0.325	0.257
	(0.87)	(1.62)				(1.55)	(1.24)		
pead	0.192	0.063	1.232	0.429	operprofrd¶	0.074	0.065	0.284	0.257
	(5.50)	(1.90)				(1.05)	(0.93)		
tmt¶	0.139	0.140	0.361	0.418	idiovolaht¶	0.116	0.113	0.222	0.255
	(1.71)	(1.91)				(0.76)	(0.90)		
exchswitch	0.130	0.222	0.271	0.408	rev6	0.095	0.050	0.316	0.253
	(1.14)	(1.70)				(1.69)	(1.36)		
lrreversal	0.077	0.091	0.269	0.395	tax	0.051	0.045	0.276	0.252
	(0.98)	(1.53)				(1.21)	(1.15)		
dnoa	0.058	0.046	0.436	0.375	mom6ind	0.125	0.047	0.395	0.249
	(1.31)	(1.22)				(1.91)	(1.14)		
cheq	0.051	0.054	0.282	0.375	eqtyvol	0.132	0.117	0.248	0.243
	(1.34)	(1.78)				(0.92)	(0.91)		
volsd¶	0.080	0.107	0.214	0.373	betafp	0.089	0.104	0.186	0.238
	(1.46)	(2.22)				(0.75)	(0.98)		
divomit	0.380	0.510	0.283	0.369	ebm	0.048	0.323	0.244	0.235
	(1.45)	(1.79)				(0.87)	(1.15)		
strev	0.478	0.068	1.375	0.340	operprof¶	0.055	0.047	0.256	0.226
	(3.18)	(1.08)				(1.00)	(0.90)		
mom3mspread¶	0.320	0.069	0.668	0.336	spread	0.228	0.119	0.378	0.216
	(2.47)	(1.04)				(1.74)	(0.97)		
idiospread	0.120	0.082	0.355	0.318	value	0.107	0.051	0.347	0.215
	(1.51)	(1.26)				(1.64)	(1.02)		
bmdec	0.071	0.076	0.278	0.306	momvol	0.148	0.249	0.434	0.209
	(1.02)	(1.10)				(2.24)	(0.97)		
roe¶	0.059	0.080	0.203	0.300	cash	0.077	0.045	0.273	0.209
	(0.78)	(1.20)				(1.49)	(1.10)		
trendfactor	0.275	0.345	0.879	0.292	delfinl	0.070	0.030	0.506	0.209
	(2.53)	(1.30)				(1.96)	(0.76)		

Table A.1, Continued.

	CAPM α		Info. Ratio			CAPM α		Info. Ratio	
Signal	Gross	Net	Gross	Net	Signal	Gross	Net	Gross	Net
divinit	0.067	0.401	0.060	0.203	sharerepo	0.037	0.032	0.172	0.146
	(0.33)	(1.00)				(0.80)	(0.66)		
betaliqityps¶	0.046	0.037	0.221	0.201	neteqtyfin¶	0.047	0.045	0.136	0.142
	(0.96)	(0.74)				(0.51)	(0.53)		
chassetto¶	0.050	0.031	0.304	0.200	totalaccruals	0.044	0.028	0.218	0.136
	(1.39)	(0.94)				(0.91)	(0.54)		
coupon¶	0.026	0.054	0.080	0.200	rds	0.054	0.026	0.262	0.130
	(0.35)	(0.86)				(1.22)	(0.71)		
chtax	0.117	0.028	0.711	0.198	investment¶	0.027	0.018	0.176	0.130
	(2.48)	(0.70)				(0.66)	(0.54)		
mom12off	0.208	0.055	0.590	0.197	divyieldst	-0.006	0.116	-0.034	0.125
	(2.61)	(0.92)				(-0.11)	(0.59)		
rating	0.120	0.109	0.206	0.191	strevb	0.029	0.175	0.054	0.124
	(0.78)	(0.72)				(0.38)	(0.57)		
chinv	0.039	0.026	0.272	0.187	revsurp	0.098	0.025	0.415	0.121
	(1.15)	(0.77)				(1.49)	(0.53)		
ratingxspread	0.176	0.103	0.301	0.185	pcttotacc¶	0.027	0.021	0.109	0.117
	(1.32)	(0.79)				(0.44)	(0.53)		
bookprc	0.039	0.062	0.106	0.182	dts	0.112	0.040	0.287	0.114
	(0.45)	(0.75)				(1.29)	(0.51)		
idiovol3f¶	0.106	0.051	0.231	0.177	delcol¶	0.027	0.020	0.132	0.113
	(0.87)	(0.60)				(0.60)	(0.57)		
ret61	0.136	0.036	0.391	0.171	firmage	0.036	0.029	0.133	0.111
	(1.99)	(0.84)				(0.50)	(0.42)		
cashprod¶	0.045	0.029	0.216	0.170	volumetrend¶	0.041	0.032	0.123	0.107
	(0.77)	(0.61)				(0.59)	(0.51)		
sprtod2d	0.236	0.073	0.516	0.170	dcpvolspread¶	0.034	0.125	0.192	0.107
	(1.95)	(0.72)				(0.62)	(0.50)		
herf¶	0.031	0.026	0.176	0.166	illiquidity	0.058	0.047	0.124	0.103
	(0.84)	(0.84)				(0.46)	(0.39)		
realizedvol¶	0.100	0.059	0.203	0.163	impliedspread	0.149	0.047	0.281	0.103
	(0.75)	(0.59)				(1.25)	(0.47)		
assetgrowth	0.054	0.029	0.301	0.161	compdebtiss	0.044	0.017	0.245	0.100
	(1.11)	(0.56)				(1.06)	(0.43)		
delequ	0.036	0.023	0.221	0.158	volatility	0.074	0.032	0.199	0.099
	(0.89)	(0.61)				(0.94)	(0.46)		
booklev¶	0.054	0.029	0.269	0.155	bondsize¶	0.070	0.033	0.204	0.098
	(1.44)	(0.82)				(0.96)	(0.45)		
varcf¶	0.070	0.067	0.148	0.147	operlvg	0.051	0.022	0.207	0.097
	(0.55)	(0.54)				(1.26)	(0.54)		

Table A.1, Continued.

	CAPM α		Info. Ratio			CAPM α		Info. Ratio	
Signal	Gross	Net	Gross	Net	Signal	Gross	Net	Gross	Net
maxret¶	0.175	0.031	0.403	0.091	gp	0.029	0.006	0.141	0.030
	(1.54)	(0.32)				(0.55)	(0.11)		
compequiss	0.016	0.015	0.078	0.081	var	0.033	0.011	0.084	0.029
	(0.35)	(0.36)				(0.35)	(0.12)		
indipo¶	-0.021	0.027	-0.040	0.079	turnvol	0.027	0.009	0.084	0.029
	(-0.20)	(0.28)				(0.55)	(0.19)		
ave12mspread	0.058	0.042	0.104	0.078	cf¶	0.091	0.008	0.308	0.028
	(0.46)	(0.34)				(1.50)	(0.13)		
spinoff¶	0.061	0.032	0.134	0.074	hire	0.028	0.004	0.179	0.023
	(0.40)	(0.22)				(0.63)	(0.08)		
bookleverage¶	0.020	0.022	0.059	0.066	high52	0.196	0.010	0.413	0.023
	(0.30)	(0.36)				(1.46)	(0.09)		
zerotrade¶	0.018	0.020	0.045	0.057	mom11yroff	0.054	0.003	0.271	0.021
	(0.17)	(0.21)				(1.16)	(0.10)		
investppeinv	0.031	0.010	0.171	0.056	cfp¶	0.071	0.004	0.185	0.011
	(0.63)	(0.19)				(0.73)	(0.04)		
grltnoa	0.031	0.010	0.119	0.053	residmom	0.090	0.001	0.427	0.008
	(0.58)	(0.28)				(2.00)	(0.03)		
dellti	0.005	0.007	0.030	0.052	earnstreak	0.069	0.000	0.287	0.001
	(0.11)	(0.20)				(2.22)	(0.01)		
coskewness	0.006	0.008	0.031	0.051	analystvalue¶	0.023	-0.002	0.095	-0.008
	(0.13)	(0.23)				(0.40)	(-0.04)		
predictedfe¶	0.021	0.008	0.105	0.050	convdebt	0.002	-0.003	0.009	-0.010
	(0.46)	(0.22)				(0.03)	(-0.04)		
mktlev¶	0.076	0.016	0.225	0.048	delcoa¶	0.003	-0.003	0.014	-0.016
	(0.83)	(0.20)				(0.05)	(-0.06)		
fgr5yrlag	0.002	0.006	0.016	0.047	pctacc¶	0.022	-0.004	0.071	-0.016
	(0.07)	(0.19)				(0.34)	(-0.09)		
cboperprof¶	0.012	0.009	0.059	0.046	herfasset	0.003	-0.003	0.016	-0.018
	(0.25)	(0.19)				(0.07)	(-0.08)		
chinvia¶	0.045	0.006	0.265	0.043	accruals¶	0.004	-0.005	0.015	-0.020
	(1.30)	(0.24)				(0.07)	(-0.10)		
numearnincr	0.042	0.005	0.240	0.041	opleverage	0.016	-0.009	0.043	-0.025
	(1.21)	(0.19)				(0.28)	(-0.16)		
noa	0.020	0.014	0.054	0.040	imsprtod2d	0.114	-0.011	0.259	-0.025
	(0.30)	(0.22)				(1.08)	(-0.11)		
ztradealt2¶	0.019	0.014	0.049	0.037	netdebtfin	0.030	-0.005	0.154	-0.026
	(0.18)	(0.14)				(0.66)	(-0.10)		
totaldebt¶	0.038	0.015	0.086	0.035	earnsurpise	0.084	-0.005	0.472	-0.030
	(0.50)	(0.20)				(1.96)	(-0.16)		

Table A.1, Continued.

Signal	CAPM α Gross	CAPM α Net	Info. Ratio Gross	Info. Ratio Net	Signal	CAPM α Gross	CAPM α Net	Info. Ratio Gross	Info. Ratio Net
xfin	0.032	-0.009	0.120	-0.037	bpebm	0.060	-0.015	0.371	-0.091
	(0.59)	(-0.18)				(1.81)	(-0.30)		
am¶	0.035	-0.017	0.074	-0.038	debtebitda¶	-0.005	-0.036	-0.013	-0.096
	(0.29)	(-0.15)				(-0.07)	(-0.57)		
leverage¶	0.031	-0.018	0.067	-0.040	debtissuance	-0.001	-0.015	-0.004	-0.101
	(0.27)	(-0.16)				(-0.02)	(-0.52)		
gpat	0.019	-0.013	0.058	-0.041	rankrevgr	0.001	-0.022	0.003	-0.104
	(0.26)	(-0.19)				(0.02)	(-0.49)		
herfbe	-0.002	-0.007	-0.010	-0.042	delbreadth	0.110	-0.023	0.411	-0.104
	(-0.05)	(-0.20)				(1.90)	(-0.53)		
shareiss5y¶	0.012	-0.007	0.067	-0.042	earnforedisp	0.053	-0.020	0.241	-0.108
	(0.32)	(-0.19)				(1.10)	(-0.48)		
chggpat¶	0.005	-0.005	0.041	-0.043	delnetfin	0.015	-0.019	0.089	-0.109
	(0.13)	(-0.15)				(0.35)	(-0.41)		
shareiss1y¶	0.017	-0.010	0.069	-0.043	d2d	0.032	-0.055	0.060	-0.110
	(0.32)	(-0.19)				(0.22)	(-0.42)		
fr	0.018	-0.011	0.079	-0.052	vixbeta¶	0.069	-0.027	0.265	-0.119
	(0.32)	(-0.23)				(1.00)	(-0.46)		
marketcap	0.009	-0.023	0.020	-0.052	grsaletogrinv¶	0.010	-0.017	0.067	-0.133
	(0.07)	(-0.19)				(0.26)	(-0.71)		
eqtydur¶	0.024	-0.014	0.102	-0.057	chnncoa	0.029	-0.022	0.181	-0.134
	(0.47)	(-0.27)				(0.91)	(-0.76)		
mom12m	0.099	-0.016	0.259	-0.064	bidaskspread	0.081	-0.049	0.182	-0.134
	(1.07)	(-0.24)				(0.74)	(-0.51)		
bm¶	0.014	-0.016	0.054	-0.064	chnwc¶	0.009	-0.022	0.050	-0.136
	(0.18)	(-0.22)				(0.17)	(-0.47)		
payoutyield	0.012	-0.009	0.085	-0.065	grcapx¶	0.013	-0.020	0.086	-0.143
	(0.39)	(-0.28)				(0.32)	(-0.54)		
abmaccruals	0.016	-0.014	0.083	-0.071	volmkt	0.024	-0.061	0.054	-0.144
	(0.36)	(-0.28)				(0.20)	(-0.56)		
aop	0.004	-0.013	0.023	-0.074	netpoyld¶	0.015	-0.025	0.090	-0.151
	(0.10)	(-0.29)				(0.42)	(-0.68)		
entmult¶	0.036	-0.021	0.121	-0.075	momoffseason	0.011	-0.040	0.041	-0.162
	(0.51)	(-0.33)				(0.17)	(-0.68)		
dolvol	0.028	-0.030	0.066	-0.076	intancfp¶	0.044	-0.060	0.121	-0.171
	(0.28)	(-0.32)				(0.57)	(-0.81)		
shortinterest	0.027	-0.035	0.062	-0.087	roaq	-0.002	-0.073	-0.004	-0.187
	(0.22)	(-0.31)				(-0.01)	(-0.68)		
faceval¶	-0.001	-0.029	-0.002	-0.087	sp¶	0.007	-0.068	0.017	-0.187
	(-0.01)	(-0.51)				(0.06)	(-0.71)		

Table A.1, Continued.

Signal	CAPM α Gross	CAPM α Net	Info. Ratio Gross	Info. Ratio Net	Signal	CAPM α Gross	CAPM α Net	Info. Ratio Gross	Info. Ratio Net
grcapx3y¶	-0.001	-0.028	-0.003	-0.192	momseason¶	0.113	-0.128	0.455	-0.631
	(-0.01)	(-0.89)				(2.22)	(-3.10)		
indmom	0.082	-0.062	0.235	-0.205	iomomcust	0.077	-0.147	0.299	-0.662
	(1.16)	(-0.99)				(1.84)	(-4.13)		
intansp¶	0.005	-0.079	0.012	-0.206	dspread	0.085	-0.223	0.199	-0.663
	(0.05)	(-0.86)				(1.06)	(-3.26)		
ep¶	0.008	-0.035	0.043	-0.215	retconglom	0.108	-0.127	0.511	-0.734
	(0.18)	(-0.90)				(1.91)	(-3.55)		
intanbm¶	0.005	-0.086	0.013	-0.226	dvolput	0.217	-0.148	0.823	-0.743
	(0.06)	(-0.98)				(4.18)	(-4.07)		
mom6	0.016	-0.126	0.024	-0.241	optionvolume2¶	0.058	-0.112	0.322	-0.745
	(0.11)	(-1.13)				(1.09)	(-2.58)		
mom16yroff¶	0.020	-0.051	0.090	-0.273	mom6yrseas	0.116	-0.135	0.517	-0.893
	(0.37)	(-1.02)				(1.74)	(-3.33)		
mom6xrtg	-0.004	-0.139	-0.007	-0.281	analystrevn	0.085	-0.138	0.503	-0.990
	(-0.03)	(-1.25)				(1.51)	(-4.64)		
rdipo¶	-0.168	-0.153	-0.295	-0.286	mom16yrseas	0.003	-0.162	0.015	-1.093
	(-1.06)	(-1.07)				(0.07)	(-4.63)		
ztradealt1	0.013	-0.101	0.035	-0.289	momseasshort¶	0.011	-0.215	0.041	-1.101
	(0.12)	(-1.02)				(0.20)	(-5.08)		
foredisp	0.034	-0.109	0.084	-0.299	rivolspread	0.010	-0.234	0.040	-1.111
	(0.28)	(-1.00)				(0.17)	(-4.46)		
intanep¶	0.003	-0.097	0.009	-0.304	cpvolspread¶	0.015	-0.180	0.079	-1.172
	(0.04)	(-1.31)				(0.25)	(-4.68)		
optionvolume1	0.005	-0.081	0.016	-0.308	returnskew¶	0.121	-0.165	0.731	-1.176
	(0.08)	(-1.37)				(2.94)	(-5.66)		
intmom	0.009	-0.099	0.026	-0.345	returnskew3f¶	0.092	-0.172	0.510	-1.234
	(0.10)	(-1.16)				(1.86)	(-5.28)		
bondage	0.008	-0.041	0.065	-0.352	skew1	0.034	-0.210	0.142	-1.244
	(0.23)	(-1.34)				(0.49)	(-5.20)		
mom6yroff	0.026	-0.055	0.162	-0.381	smileslope	0.017	-0.195	0.096	-1.253
	(0.83)	(-1.77)				(0.34)	(-5.04)		
iomomsupp	0.103	-0.108	0.395	-0.476	mom11yrseas	0.042	-0.202	0.226	-1.410
	(2.32)	(-2.69)				(1.14)	(-8.01)		
dvolcall¶	0.243	-0.127	0.852	-0.574	divseason	0.001	-0.202	0.007	-1.413
	(4.44)	(-3.38)				(0.03)	(-6.14)		
uprecomm	0.013	-0.104	0.119	-0.576	downrecomm	0.027	-0.283	0.279	-1.716
	(0.82)	(-3.90)				(1.36)	(-8.01)		
exclexp	0.017	-0.075	0.105	-0.583	deltarecomd	0.079	-0.214	0.548	-1.921
	(0.38)	(-2.11)				(1.90)	(-8.41)		

This table reports the gross and net CAPM α and information ratios of the long-short strategies based on the underlying characteristics. The results are sorted by net information ratio.

Table A.2: Average Slope Coefficients of Monthly Risk-Adjusted Returns on Bond Mutual Funds on Size Decile Dummies

Parameter		Estimate	(s.e.)
TNA Dummy	1	0.003	(0.02)
	2	-0.035	(0.02)
	3	-0.017	(0.02)
	4	-0.017	(0.02)
	5	-0.024	(0.02)
	6	-0.006	(0.03)
	7	-0.010	(0.02)
	8	-0.022	(0.02)
	9	-0.019	(0.02)
Intercept		0.068	(0.02)
Number of Funds		1537	
R-Squared		0.013	

This table reports the average slope coefficients of the regression of mutual fund returns adjusted for the market risk on ten dummy variables based on the fund TNA in the previous month. The sample is from August 2002 to November 2022. Values in parentheses are standard errors.

Table A.3: Hyperparameters Across the Machine Learning Models.

Panel A: Linear models with penalties: LASSO, RIDGE & ENET		
Parameter	`sklearn` mnemonic	Value
Intercept	`fit_intercept=True`	True
ℓ_1 penalty	`alphas`	Variable
ℓ_2 penalty	`alphas`	$\in [0.0001, \ldots, 1]$
Num. Penalties	`n_alphas`	100
ℓ_1 ratio	`l1_ratio`	$\in$ [0.001, 0.01,0.99,0.999]

Panel B: Tree-based ensembles: RF and XT		
Parameter	`sklearn` mnemonic	Value
Num. Trees	`n_estimators`	100
Max depth	`max_depth`	$\in$ [2,4,6]
Split features	`max_features`	$\in$ [5,10,20]
Min leaf samples	`min_samples_leaf`	$\in$ [1,10,50]

Panel C: Feed forward neural network: NN		
Parameter	`tensorflow` mnemonic	Value
Layers	`Dense`	1
Neurons	`Dense`	32
Activation	`activation='relu'`	ReLu
Epochs	`epochs`	100
Batch size	`batch_size`	1024
Batch normalization	`BatchNormalization`	True
Optimizer	`optimizers.Adam`	Adam
Patience	`patience`	5
Learning rate	`learning_rate`	$\in$ [0.001, 0.01]
ℓ_1 penalty	`regularizers.l1`	$\in$ [0.001, 0.01]
Ensemble	-	10
Grand Ensemble	-	10

This table reports the respective hyperparameters that are chosen via a cross-validation scheme with a 70:30 train-validate split that maintains the temporal ordering of the data. The cross-validation is conducted every 5-years commencing on 2002:07 using an expanding window. The set of hyperparameters are chosen which yield the smallest mean squared error (MSE) in the validation sample. Panel A reports the hyperparameters for the linear models which include Lasso (LASSO), Ridge (RIDGE) and Elastic Net (ENET) penalties respectively. Panel B reports the hyperparameters for the set of tree-based ensembling nonlinear models which includes the random forest (RF) and extremely randomized trees (XT). Panel C reports the hyperparameters for the feed forward neural network (NN). All models except for the NN are estimated with `sklearn`. The NN is estimated with `tensorflow`.

Table A.4: Out-of-Sample R-Square Values and Diebold-Mariano Tests.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	LASSO	RIDGE	ENET	RF	XT	NN	LENS	NENS	ENS
					Panel A: Out-of-sample R^2_{OS}				
R^2_{OS}	0.026	-1.706	0.091	0.348	0.483	0.455	-0.140	0.687	0.461
					Panel B: Diebold-Mariano t-statistics				
LASSO		−2.125	1.324	1.211	1.700	1.268	−0.629	2.758	2.171
RIDGE			2.295	2.477	3.490	2.671	2.572	3.211	3.176
ENET				0.805	1.592	1.111	−1.124	2.543	2.004
RF					0.841	0.596	−1.303	1.882	0.802
XT						0.122	−3.543	1.092	−0.177
NN							−1.715	0.620	−0.249
LENS								3.380	4.048
NENS									−2.213

This table reports the out-of-sample R-square values (R^2_{OS}) in Panel A and pairwise Diebold-Mariano t-statistics in Panel B. In Panel A, the out-of-sample R-square for each model is computed as: $R^2_{OS} = 1 - \frac{\sum_{i,t}(r_{i,t+1} - \sum_l \widehat{r}_{i,t+1})^2}{\sum_{i,t}(r_{i,t+1})^2}$, where $r_{i,t+1}$ is the excess bond return over $t{:}t+1$ and $\widehat{r}_{i,t+1}$ is the machine learning model forecast available to the investor at month t for the period $t{:}t+1$. The R^2_{OS} values are only computed with the out-of-sample data (not including the *training* data). In panel B, the Diebold-Mariano Newey-West adjusted t-statistics indicate whether a column model outperforms the row model. Positive t-values greater (smaller) than 1.96 (−1.96) indicates that the column model outperforms (underperforms) the row model.