The Low Frequency Trading Arms Race:

Machines Versus Delays*

November 18, 2023

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

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

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 This creates delays and drags down the performance of her portfolio as the investment signal becomes outdated.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 make this idea operational, we deviate from the classical method of portfolio construction of Fama and French (1992) 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 investor's size threshold. As the threshold increases, the bid-ask spreads tighten while the number of eligible transactions declines, delaying the trade. In this way, we no longer assume that the trade occurs at the end of the month, right after the investor receives the signal.

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. 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.

We also take into account real-world frictions that an investor would face. If she wants to unwind the existing position but is unable to do so due to delays or lack of volume, she will have to pay the cost of carry to finance the additional unwanted positions in her inventory. In our framework, she may end up holding the existing position for several months after the signal suggests that she needs to unwind the position. As such, our empirical framework imposes a rich structure to provide

[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.

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

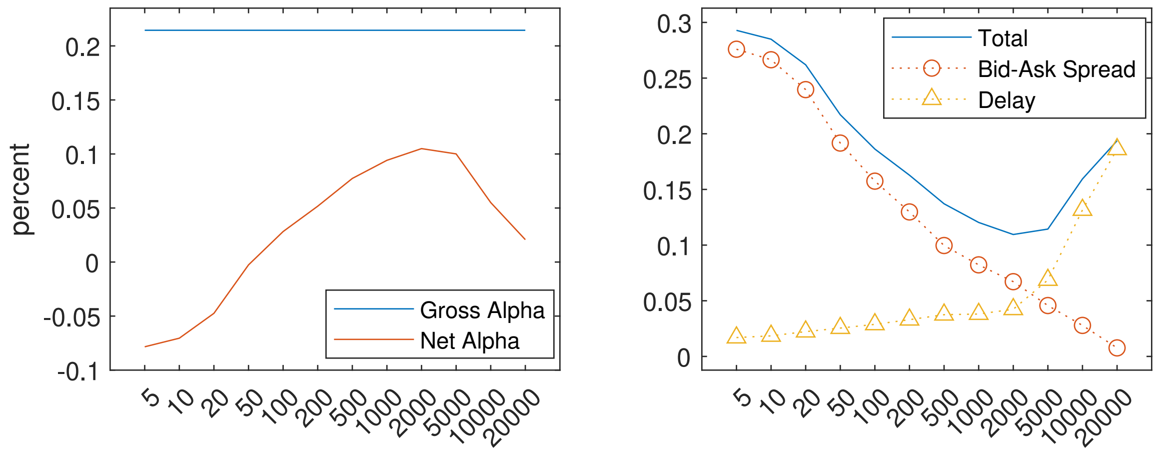

This figure plots the bond CAPM alphas of the long-short strategies based on corporate bonds' credit spreads before and after accounting for transaction costs (left panel). The transaction costs are decomposed into the bid-ask spread costs and delay costs (right panel). Values on the x-axis are the trade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

a more realistic estimate of the actual returns of actively managed corporate bond portfolios.

We show that the cost of delay is substantial. Consider a simple example in which the investor 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. 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 best 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 developed 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 regression, generalized linear models with penalization, dimension reduction via principal components regression (PCR), partial least squares (PLS) and instrumented principal components (IPCA), and regression trees (including boosted and extreme trees and random forests). 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. Most 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 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 at about $2 million per trade for all strategies with an interior solution. For example, the 'COMBO' strategy which averages the expected returns of all ML strategies generates an alpha of 0.25% before transaction costs and 0.01% after costs at the optimum. Of the 0.24% cost, 0.11% is due to delay, while 0.13% 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

[2]The characteristics comprise the 29 stock and bond characteristics used in Kelly, Palhares, and Pruitt (2021).

strategies. 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 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, for a strategy with a monthly turnover rate of 20%, a gross alpha of around 0.1% is sufficient to break even. However, if the turnover rate is 35%, the break-even gross alpha doubles to 0.2%. 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 (44 funds) generated an after-tax 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. These results support the validity of our assessment of ML-based corporate bond strategies.

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; Binsbergen et al. 2023; Dickerson et al. 2023a; 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; Pinter et al. 2021; Choi et al. 2023). 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. (2021), 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, one needs an exogenously specified size that captures the reality to discipline the estimated transaction costs.

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 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 notational 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.

2.1 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,

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

where P^h_{t+1} and P^h_t 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

$$R^{b(v)h}_{t+1} = \left(1 + R^f_{t+1} \times \frac{d}{Days_{t+1}}\right)\left(\frac{P^h_{t+1} + AI_{t+1} + C_{d,t+1}}{P^{b(v)}_{t+1,d} + AI_{t+1,d}}\right) - 1, \tag{2}$$

where $P^{b(v)}_{t+1,d}$ 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 $C_{d,t+1}$ 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 (2) 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 P^h_t) 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

$$R_{t+1}^{hs(v)} = \left(\frac{P_{t+2,d}^{s(v)} + AI_{t+2,d} + C_{t+1} + C_{t+2,d}}{P_t^h + AI_t} \right) \div \left(1 + R_{t+2}^f \times \frac{d}{Days_{t+2}} \right) - 1, \quad (3)$$

where $P_{t+2,d}^{s(v)}$ 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 $R_{t+1}^{hs(v)} = R_{t+1}^{hh}$ 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

$$R_{t+1}^{b(v)s(v)} = \left(1 + R_{t+1}^{f} \times \frac{d_1}{Days_{t+1}}\right)\left(\frac{P_{t+2,d_2}^{s(v)} + AI_{t+2,d_2} + C_{d_1,t+1} + C_{t+2,d_2}}{P_{t+1,d_1}^{b(v)} + AI_{t+1,d_1}}\right) \div \left(1 + R_{t+2}^{f} \times \frac{d_2}{Days_{t+2}}\right) - 1. \tag{4}$$

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

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

2.2 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 2. 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 $f(x_{t+1}, I_{t+1})$.

To concretely describe the set of actions in each month, we consider seven bonds as shown in Table 2. 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.

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

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. (2021). 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 return, which allows us to closely replicate the real-time investor's trading profits.

2.3 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 half the 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.

3 Data

3.1 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 Appendix A. To train the machine learning models, we construct the 29 bond and equity characteristics (otherwise known as model features) used by Kelly et al. (2021), henceforth KPP. This data combines several monthly bond and stock-based characteristics that have been shown the predict one-month ahead future corporate bond excess returns. The database includes 15 bond-based characteristics and 14 equity-based characteristics. Detailed descriptions of the construction of these variables are provided in the Appendix Table A2. All characteristics are cross-sectionally rank demeaned to lie in the interval $[-0.50, 0.50]$. In robustness, we use the publicly available dataset made available by KPP which includes bond returns and the 29 characteristics. Our results are close to identical.[5] Overall, the data used to train the ML models with non-missing data for the 29 stock and bond characteristics comprises 15,483 bonds issued by 1,492 firms over the sample period from January 1998 to December 2022 ($T = 288$).

3.2 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 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 \$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 first eleven cutoff values are set following Edwards et al. (2007) and we add another large value of \$20 million as 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 2.1. 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}).

[5] We thank Bryan Kelly and Seth Pruitt for making this data publicly available on their websites.

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 746,464 bond-month observations from August 2002 to November 2022 ($T = 244$).

Table 3 reports the summary statistics of the panel data for selected transaction sizes of $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 -0.10%, 0.50%, and -0.13% for R^{bh}, R^{hs}, and R^{bs}, respectively. The average for short positions are higher and 0.35%, 0.95%, and 0.85% for R^{sh}, R^{hb}, and R^{sb}, respectively.[6] 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 unexecuted trades increases. This explains the net return's convergence to either the gross return or the risk-free rate.

Table 4 reports the same statistics using duration-adjusted 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

[6]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).

$1 million account for less than 20% of the number of trades. The prevalence of small trades suggests the importance of trade delays and their associated costs. 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 (Wu 2022), leading to the declining share of small trades in the 2020 pandemic crisis. This pattern confirms the increasing importance of trade delays.[7]

4 Performance of the Machine Learning Models

4.1 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:

$$R_{i,t+1} = E_t(R_{i,t+1}) + \epsilon_{i,t+1}, \tag{5}$$

where,

$$E_t(R_{i,t+1}) = g^*(z_{i,\tau}). \tag{6}$$

Bonds are indexed as $i = 1, \dots, N$ and months by $t = 1, \dots, T$. 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. With one exception, all of our model estimates minimize the mean squared prediction errors (MSE). In total, we consider eleven linear and non-linear machine learning models including ordinary least squares (OLS), OLS with the Huber loss function (OLS-Huber); penalized linear regression techniques: LASSO, Ridge and Elastic Net (ENET); dimension reduction methods including principal component analysis (PCA), partial least squares (PLS) and instrumented principal component analysis (IPCA); and non-linear regression tree based methods including random forests (RF), gradient

[7] In Internet Appendix C, we study institutional investors' quarterly position changes in eMAXX to confirm the validity of trade size estimates.

boosted trees (GB) and extreme trees (XTREE). In addition, we form the combination model (COMBO), which is the equally-weighted average across all of the eleven 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 29-dimensional vector of bond characteristics at $t-1$. We perform cross-validation using a 85:15 training-validation split. We then use the vector of characteristics available at time t to produce a forecast of bond excess returns at $t+1$. These forecasts (expected returns) are available to the bond portfolio manager at time t, meaning they can trade on them at the end of the month. Thereafter, all models are re-trained and cross-validated every six months with an expanding window.[8] This model is then used to form predictions each month for the following six months. We provide additional details related to the cross-validation and training of the respective models in Section B of the Online Appendix.

4.2 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 29 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 returns. We then sign the raw characteristics so that, on average, the long position has higher returns than the short position. We compute the average gross/net returns of the high-minus-low portfolio and the turnover rate.

In addition, we estimate CAPM alpha by running time series regressions of the strategies' returns on the market factor:

$$R_t^e = \alpha + \beta MKTB_{Net,t} + \varepsilon_t, \tag{7}$$

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

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 A. 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. (2023a) is 0.36% over the same period. The lower value of the ETF returns suggests that even 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 average excess returns and CAPM alphas before transaction costs, shown in Table 5. Of the 29 characteristics, four generate significantly positive average returns. These are credit spreads (0.44%), six-month changes in credit spreads (6mspread, 0.52%), the ratio of spreads to distance-to-default (sprtod2d, 0.37%), and volatility (0.40%). Many of the variables are related to credit spreads, a finding consistent with Nozawa (2017), who show that credit spreads predict bond returns. Looking at the CAPM alphas, equity momentum (0.14%), 6mspread (0.45%), and sprtod2d (0.24%) generate significant returns after adjusting for market risk. Figure 5 plots the gross average excess returns (Panel A) and the CAPM alphas (Panel B), visualizing the information in Table 5. Overall, these results seem to support the observation of Dick-Nielsen et al. (2023) that many bond characteristics, when used individually, do not predict bond returns.

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 CAPM alphas.

The second column of Table 6 reports the average excess returns of the machine learning-based strategies. We find that ML algorithms significantly improve the return predictability of the underlying signals. Ten of the twelve strategies deliver significantly positive average excess returns. Somewhat surprisingly, relatively simple methods, including OLS and OLS-Huber, generate high average excess returns of 0.37% (t=3.34) and 0.50% (t=3.15), respectively. These results show that the ML techniques we use do indeed extract useful information in predicting bond returns. They are also consistent with Bali et al. (2020), who find that most ML techniques perform equally well in predicting corporate bond returns. Looking at the CAPM alphas, five of the twelve strategies generate significant alphas before transaction costs. These strategies include OLS (0.28%), LASSO (0.23%), Ridge (0.28%), RF (0.22%), and GB (0.24%).

Binsbergen et al. (2023) find that adjusting for corporate bond duration significantly affects 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 7 shows the average duration-adjusted returns and the CAPM alphas. Using duration-adjusted returns improves the performance of the model after adjusting for market exposure. Ten of the twelve algorithms now generate significant CAPM alphas, confirming the value of combining multiple signals to generate reliable return forecasts.

4.3 Impact of Transaction Costs

In this section, we evaluate the impact of transaction costs on the performance of the ML strategies. Since transaction costs depend on trade size, we first use two sizes: $100,000 and the optimal value that maximizes the net CAPM alpha. We contrast the optimal size with $100,000 as a reference point because this is the value used in the prior literature to represent typical institutional trades (e.g., Bessembinder et al. 2008).

The third and sixth columns of Table 6 show the average excess returns and CAPM alphas net of transaction costs, assuming a trade size of $100,000. After taking costs into account, ten of the twelve algorithms generate negative average excess returns, and all of them have negative CAPM alphas. Figure 7 shows the gross and 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 three of them have significantly negative net alphas.

These costs can be reduced by choosing the optimal trade size. The eighth column of Table 6 reports the optimal trade size for each ML algorithm. We find that all but PCA and IPCA have an optimal trade size of $2 million. Since the large trade size reduces the half spreads charged on each trade, the net returns and alphas improve from the case of $100,000.

The fourth and seventh columns of Table 6 show the net returns and alphas at the optimal trade size. At the optimum, the CAPM alpha ranges from -0.09% to 0.04%. While six algorithms produce positive net alphas, none are statistically or economically significant. Figure 8 visualizes these estimates. While the net alphas are less negative than Figure 7, it is a challenge for the ML algorithm to generate significantly above market returns net of costs. Using duration-adjusted returns, reported in Table 7, leads to a similar conclusion.

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 9 shows the average returns of the long-short strategies before and after transaction costs. For example, the left panel of Panel A plots gross and net CAPM alphas using OLS. Before costs, this signal generates an average alpha of 0.28%. 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 part explained by half spreads and the part explained by delays. To do this, we compute alternative net returns using ICE's quotes on the transaction dates provided by TRACE. Thus, this alternative reflects the cost of delays but not half spreads. The difference between gross returns and these alternative net returns gives us the pure effect of delays, and the remainder is accounted for by the half spreads that drive a wedge between ICE quotes and TRACE transaction prices.

The right panels of Figure 9 show the decomposition of costs. Continuing with the OLS example, the effect of half-spreads falls from 0.6% at the volume of $5,000 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 \$5,000 volume to 0.30% at \$2 million, reflecting the cost of missing trading opportunities. As a result, the sum of the two costs exhibits a U-shaped pattern with 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 9. However, this pattern does not hold for an uninformative signal. In Panel F, we repeat this decomposition for PCA, which produces a negative alpha even before transaction costs. In this case, the cost of delay does not increase significantly as size increases. For a size greater than \$10 million, the cost of delay decreases. This is because if a signal is unprofitable, it is better not to execute the trade and avoid incurring bid-ask spreads. Therefore, a longer delay makes the net return less negative, which is interpreted as a benefit of delay. In this case, the total cost monotonically declines with trade size and therefore we cannot find an optimal trade size. In practice, such cases are not relevant in evaluating the role of delays: A more interesting case is strategies with positive alphas before costs. Looking at all panels in Figure 9, ten out of twelve strategies exhibit positive gross alphas and thus we observe U-shaped functions with an interior solution for the optimal trade size.

Table 8 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 0.24% to 0.35% while the delay cost ranges from 0.01% to 0.05%. With the optimal trade size, the half-spread cost is lower, ranging from 0.02% to 0.16%, reflecting the cost savings for large transactions. On the other hand, the delay cost is now higher, ranging from 0.04% to 0.13%. This pattern highlights the key trade-off between half spreads and delays.

Table 9 reports the same decomposition for individual signals. Since the profitability and portfolio turnover rates differ greatly among the signals, we observe a greater variation in transaction costs.

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 other from January 2013 to November 2022. Figure 10 shows the cost decomposition for the COMBO 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.

4.4 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. This method implicitly assumes that the realization of trades above the threshold is the result of optimal investor decisions. An investor could have negotiated with a dealer to trade immediately at size v, but chose not to do so because of the prohibitive cost. In such a case, we do not observe such trades that could have taken place. But if we did observe such trades, the costs would be higher than we observe in our data. Thus, under our assumption, the estimates are the lower bound of the true transaction costs. Even if our assumption of optimal execution does not hold, our estimates of net returns still represent realistic gains from implementable strategies with no approximation.

4.5 Determinants of Optimal Trade Size

The optimal trade size for individual signals significantly varies. The ninth column of Table 5 shows that the optimal size ranges from $500,000 to $20 million (i.e., the corner solution). 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:

$$Turn_p = \frac{1}{T}\sum_{t=1}^{T}\sum_{i\in N_t}|w_{i,p,t+1} - (1+R^{hh}_{i,t+1})w_{i,p,t}| \quad \text{where } p \in \{long, short\}, \tag{8}$$

$$Turn_s = 0.5(Turn_{long} + Turn_{short}). \tag{9}$$

In this exercise, we use the turnover rate when the trade size is the smallest ($5,000) 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 41 strategies (29 based on single signals and 12 based on ML) into three categories based independently on their turnover rate and gross CAPM α, creating nine bins. 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 α/low turnover bin and that for the low gross α/high turnover bin are missing because no strategy falls into these two categories. In Panel A, we report the average of the optimal trade size. The optimal trade size is strongly decreasing in gross α. For the low α category, the optimal volume is $1.4 million, while for the high α category it is $0.2 million. This is because when the signal is profitable, it is better to execute trades as soon as possible and avoid missing the trading opportunity generated by the signal. In such a case, it is optimal to choose a small trade size and reduce the cost of delays. On the other hand, the turnover rate does not significantly affect the optimal size. For strategies with the medium gross α, the optimal size is $0.92 million, $1.00 million, and $0.84 million for low, middle, and high turnover rates, respectively.

Turning to half spread costs (Panel C), they increase in both turnover rate and gross α. The positive relationship between half spread cost and turnover rate is somewhat mechanical, as the cost increases as investors trade more frequently. The cost is also positively correlated with gross α because a highly profitable signal optimally sets a small trade size, leading to higher bid-ask spreads.

The pattern for delay costs is more nuanced. Delay costs are highest when both gross α and turnover are high. When gross α 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 α 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 α and turnover rate, the delay cost is 0.13% on average, higher than other bins (ranging from 0.00% to 0.06%). The total cost, shown in Panel B, is the sum of the half-spread cost and the delay cost.

4.6 How Much Gross α 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 α that achieves the desired level of net α under the assumption that trade size is optimally chosen. This will allow other researchers to calculate the gross α of their strategies and quickly check whether they also generate α net of costs.

Since the relationship between gross and net α is affected by the persistence of the signal, we use the 41 strategies and estimate multivariate regressions of net alpha and the associated t-statistics on the turnover rate and gross alpha:

$$\alpha_{Net,s} = 0.013 - 0.002Turn_s + 0.472\alpha_{Gross,s} - 0.001Turn_s \times \alpha_{Gross,s} + \varepsilon_s, \tag{10}$$

$$t(\alpha_{Net,s}) = 0.177 - 0.017Turn_s + 5.809\alpha_{Gross,s} - 0.040Turn_s \times \alpha_{Gross,s} + \varepsilon_s. \tag{11}$$

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 11 shows the combination of $Turn$ and α_{Gross} required to achieve net

α of 0%, 0.1%, and 0.2% per month. In the figure, a strategy in the northwest region of the graph generates higher net α, while a strategy in the southeast region generates lower net α. 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 we can see in the chart, none of the ML strategies achieve this goal.

The right panel plots the bound using the t statistic of the net α. To achieve the t value above the 10% level, a strategy must be above the red dotted line, while the 5% significance requires being above the yellow solid line.

Achieving statistically significant net CAPM α requires relatively high values of gross α. For example, consider a hypothetical strategy with a turnover rate of 20%. Then, it must generate gross α of 0.36% and 0.43% to achieve a t-statistic of 1.65 and 1.96, respectively. If the turnover rate is 30%, the corresponding required gross α's are 0.43% and 0.50%, respectively. If a strategy's turnover is higher, then it requires higher levels of gross α to be useful in practice. For future reference, we tabulate the net CAPM α as a function of gross α and turnover rate in Table A3 in Appendix.

4.7 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 6. 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 COMBO, the gross and net alphas for the long-short strategy are 0.25% and 0.01%, respectively, while the corresponding values for the long-only strategy are 0.13% and 0.02%. None of the twelve long-only strategies generate significant alphas after transaction costs even at the optimal trade size.

Second, we investigate whether our results are driven by our choice of data, sample period and our construction of the 29 bond and stock characteristics. To verify all three of these potential concerns, we source the publicly available 29 KPP stock and bond characteristics and bond returns based on the WRDS TRACE database from Seth Pruitt's website, and replicate the estimates of net returns and alphas. This exercise not only uses a publicly available data source, but encompasses a different sample period and data source (TRACE is based on bond transactions, not quotes). Due to the signals' availability, the sample for this exercise is from July 2006 to July 2020 (T=169). Table 12 reports the gross and net returns and alphas. We confirm that the ML strategies are profitable before costs and market risk adjustment. However, the net alphas are still insignificantly different from zero in all cases, corroborating our main results.

Third, 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 OLS strategy is 745 and 663, respectively. In practice, investors may sample a smaller number of bonds for ease of implementation. 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 13 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 6. For example, for OLS, the gross returns are now 0.68% (t=3.97), higher than the main result using quintile portfolios (0.37%, t=3.34). 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 is almost doubled relative to holding a larger allocation of bonds in the long and short leg of the position. The high turnover inflates the transaction costs, which leads to an insignificant net CAPM alpha for OLS (0.12%, t=0.71). Looking across strategies, eleven of the twelve ML strategies have insignificant net alphas. The only exception is GB with an optimal trade size: it generates a significant net alpha of 0.46% (t=2.88).

5 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.

5.1 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 12. 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. 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 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.

5.2 Skill and Manager Value Added

Berk and Van 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 Van 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 Van 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.

5.3 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.[9] 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.[10] 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 (γ = 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

[9]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.

[10]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).

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 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.

6 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

Andreani, Martina, Diogo Palhares, and Scott Richardson, 2023, Computing corporate bond returns: a word (or two) of caution, *Review of Accounting Studies* 1–20.

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.

Asness, Clifford S, Andrea Frazzini, and Lasse Heje Pedersen, 2019, Quality minus junk, *Review of Accounting Studies* 24, 34–112.

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.

Bai, Jennie, Turan G. Bali, and Quan Wen, 2019, 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, 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, 2021, Book-to-market, mispricing, and the cross-section of corporate bond returns, Working Paper, University of Warwick.

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.

Bharath, Sreedhar T, and Tyler Shumway, 2008, Forecasting default with the merton distance to default model, *Review of Financial Studies* 21, 1339–1369.

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* .

Campbell, John Y., and Glen B. Taksler, 2003, Equity volatility and corporate bond yields, *Journal of Finance* 58, 2321–2349.

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, 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, Yesol Huh, and Sean Seunghun Shin, 2023, 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.

Correia, Maria, Scott Richardson, and İrem Tuna, 2012, Value investing in credit markets, *Review of Accounting Studies* 17, 572–609.

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.

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

Dickerson, Alexander, Cesare Robotti, and Giulio Rossetti, 2023b, Noisy prices and return-based anomalies in corporate bonds, Working Paper, Warwick Business School.

Duan, Tinghua, Frank Weikai Li, and Quan Wen, 2021, Is carbon risk priced in the cross-section of corporate bond returns?, *Available at SSRN 3709572* .

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.

Gamba, Andrea, and Alessio Saretto, 2013, Firm policies and the cross-section of cds spreads, *WBS Finance Group Research Paper* .

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.

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, 2021, Predicting individual corporate bond returns, *Available at SSRN 4374213* .

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 Richardson, 2018, Common factors in corporate bond returns, *Journal of Investment Management* 16, 17–46.

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

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

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

Li, Sophia Zhengzi, Peixuan Yuan, and Guofu Zhou, 2023, Risk-based momentum of corporate bonds, *Available at SSRN 4374766* .

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.

Nozawa, Yoshio, 2017, What drives the cross-section of credit spreads? a variance decomposition approach, *Journal of Finance* 72, 2045–2072.

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

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.

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: Flow Chart To Compute Net Returns

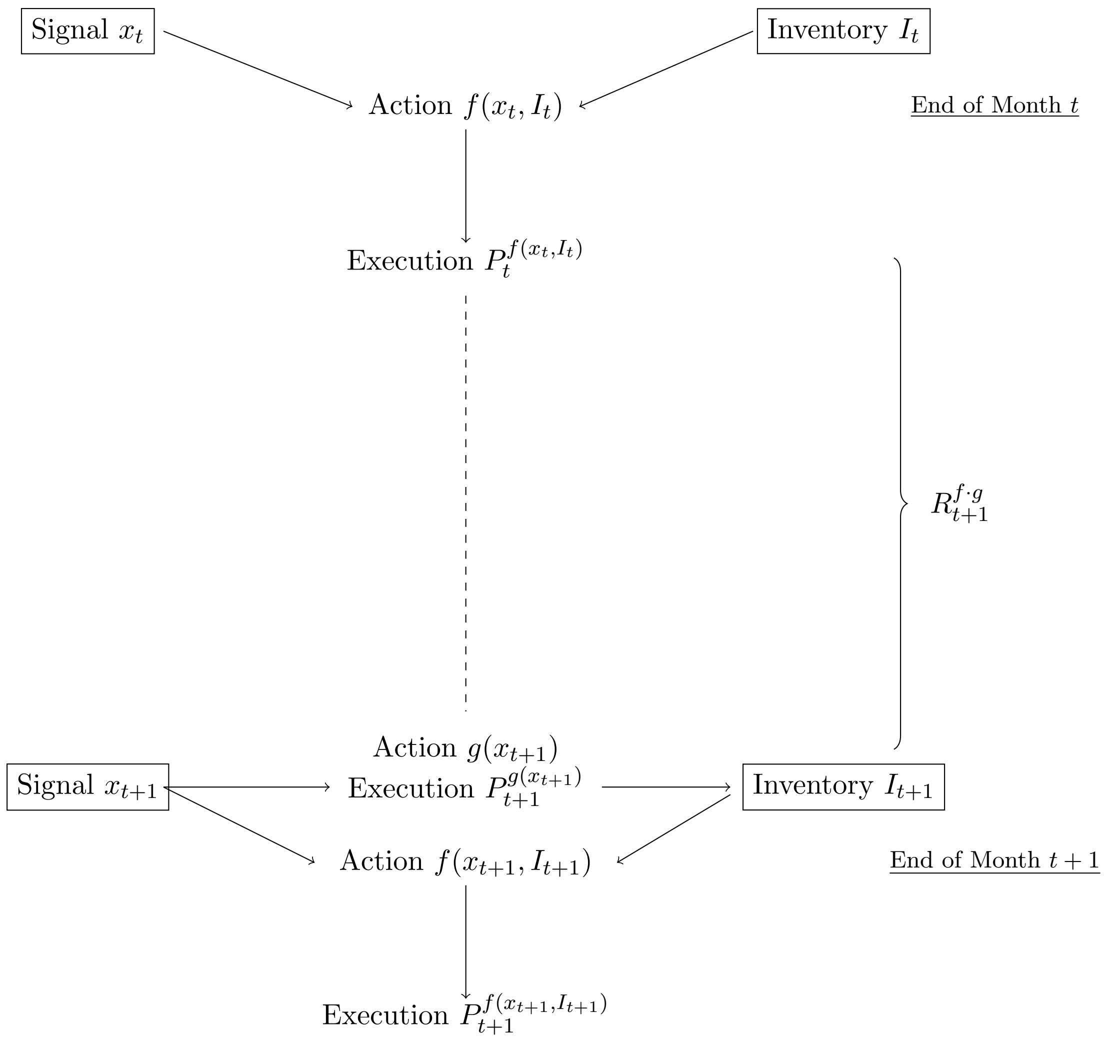

Figure 3: Summary Statistics For Different Trade Size

Panel A. Mean and Median Returns By Size

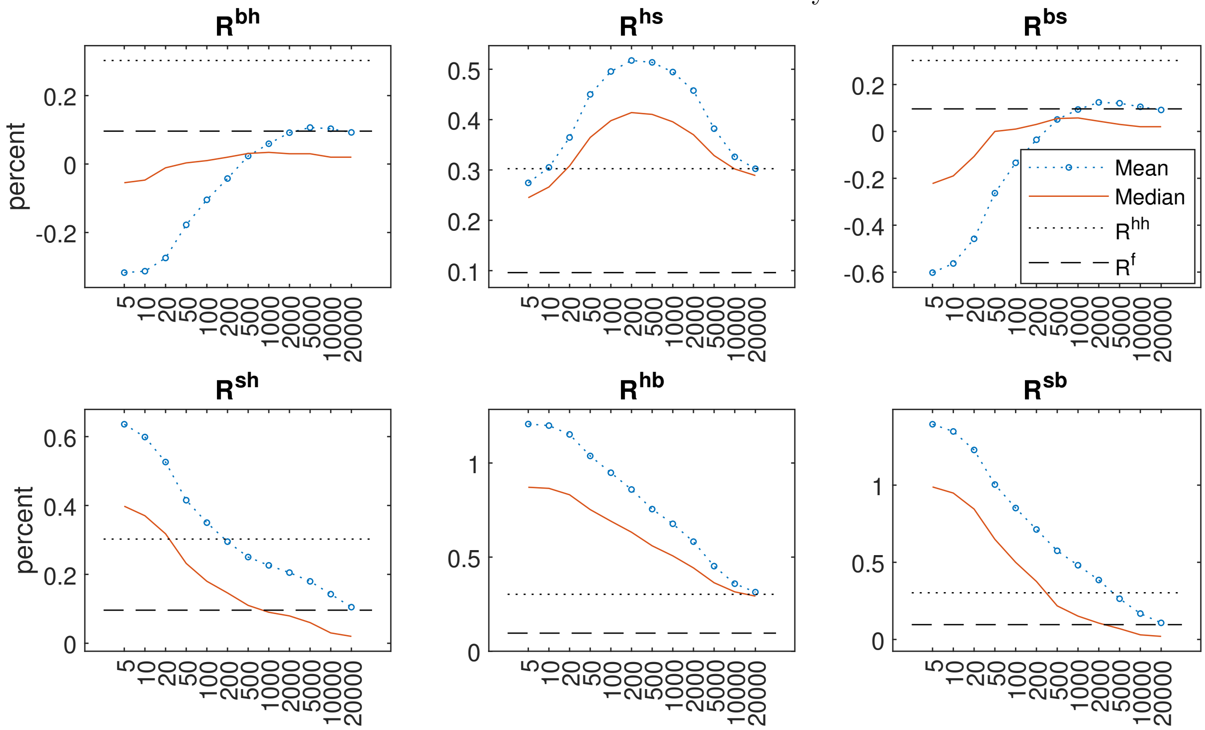

Panel B. Percentage of Observations with No Trade

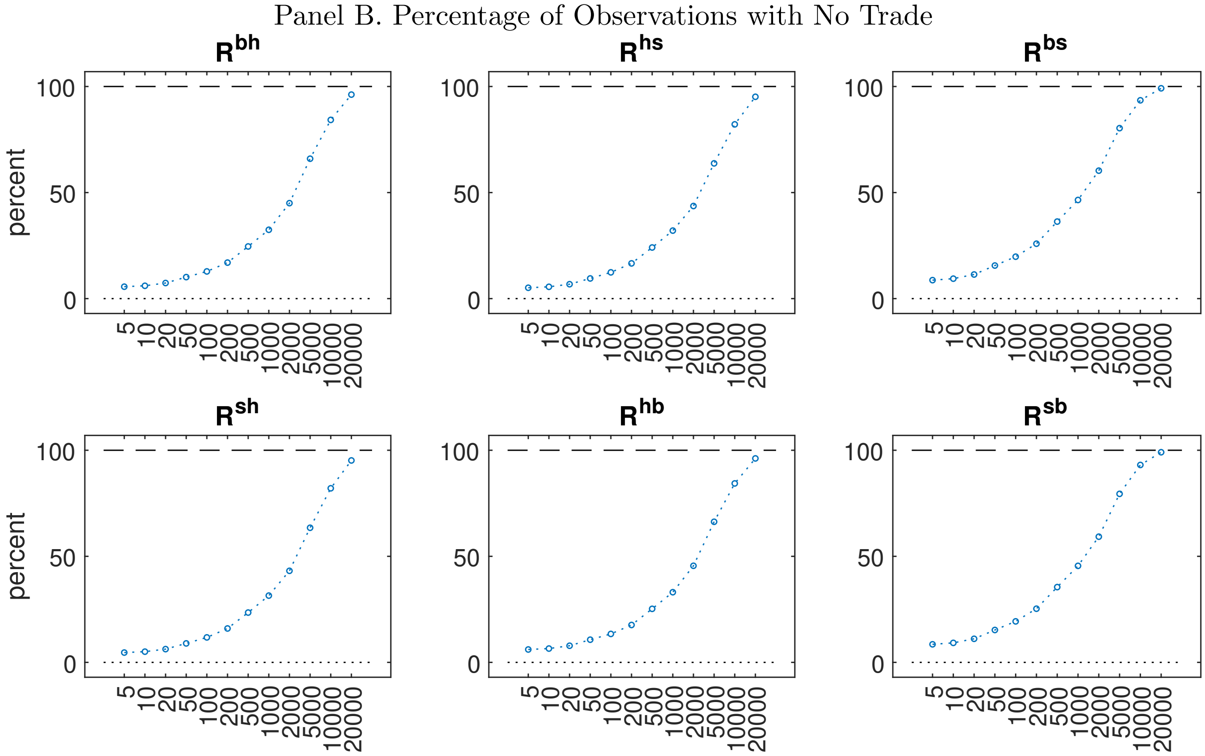

Panel A plots the mean and median net returns for different trade sizes. Panel B plots the percentage of observations where there is no trade to calculate a return in the month. Values on the x-axis are in thousand dollars.

Figure 4: Distribution of Transaction Volume in TRACE: July 2002-December 2022

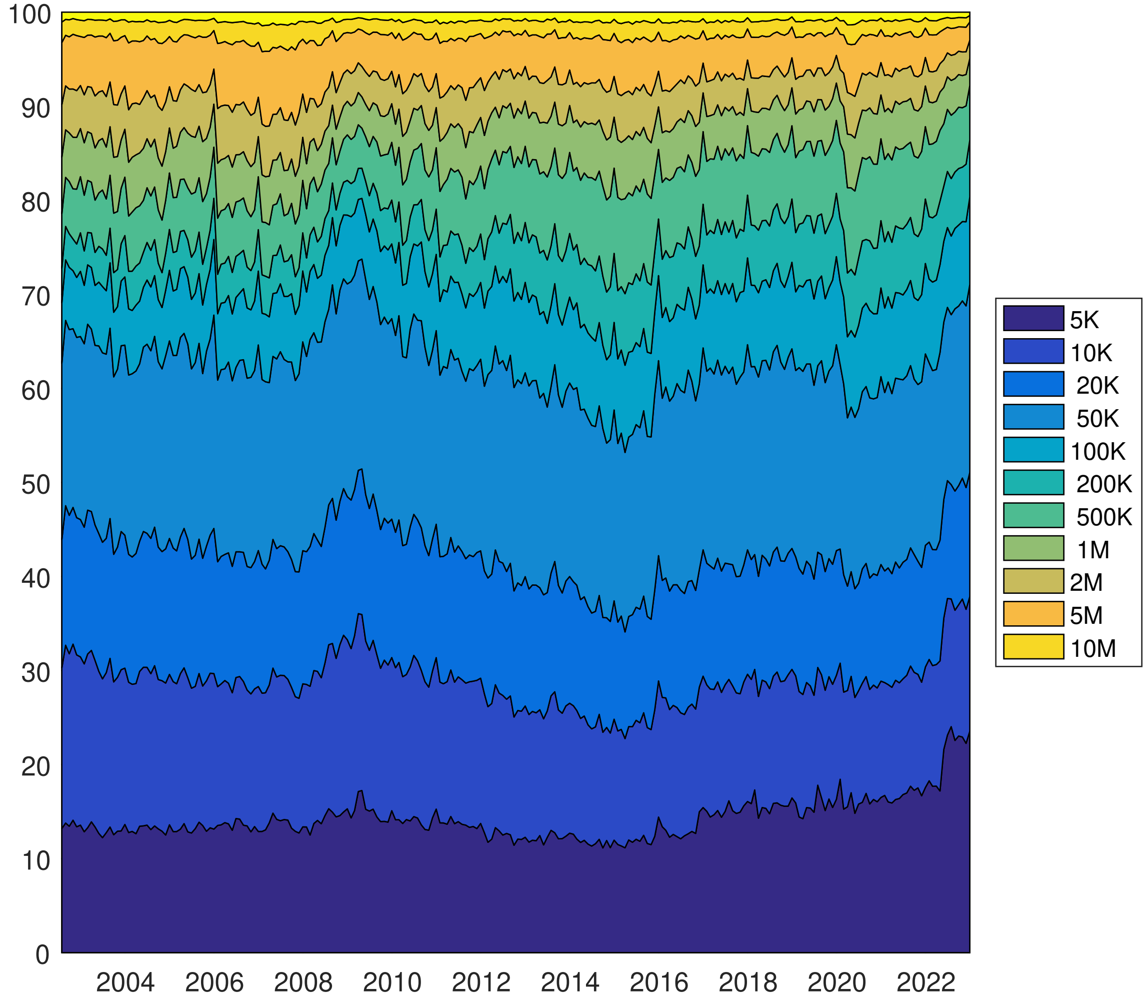

This figure plots the cumulative frequency of transaction size in TRACE. For example, the area below 10K represents transactions with a size below $10,000. The sample is from July 2002 to December 2022 and includes only dealer-customer trades.

Figure 5: Gross and Net Returns For Each Characteristic at $100,000 Volume Threshold

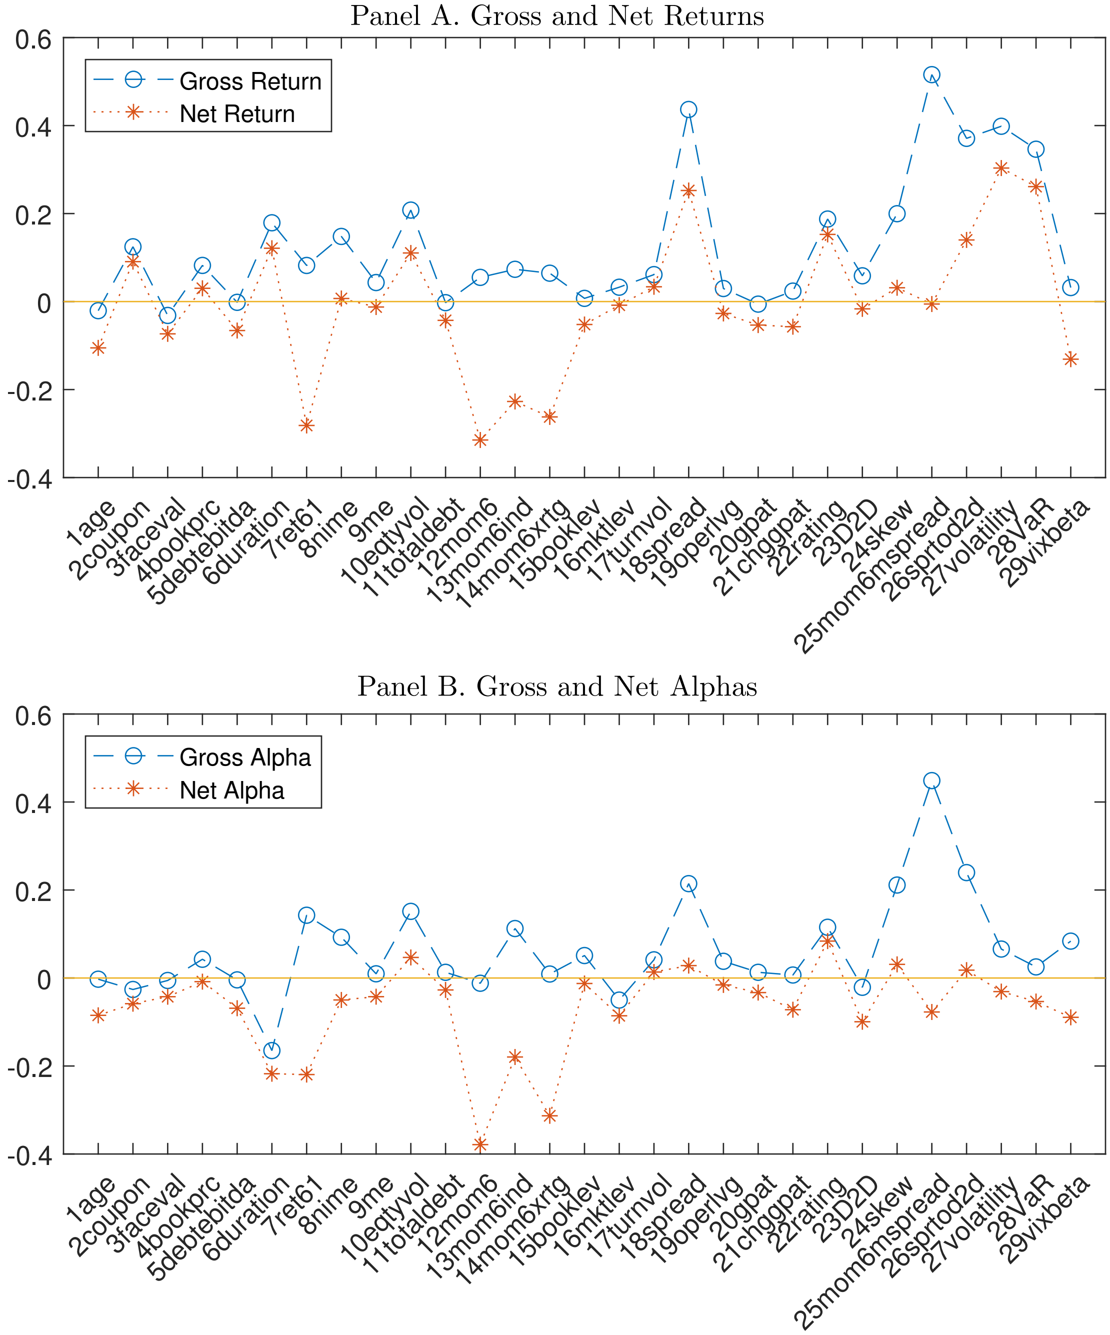

These figures plot the gross/net returns and CAPM α for each underlying signal that forms the basis of machine-learning algorithm. To compute net returns and α, we use the fixed trade size of $100,000.

Figure 6: Gross and Net Returns For Each Characteristic at the Optimal Volume Threshold

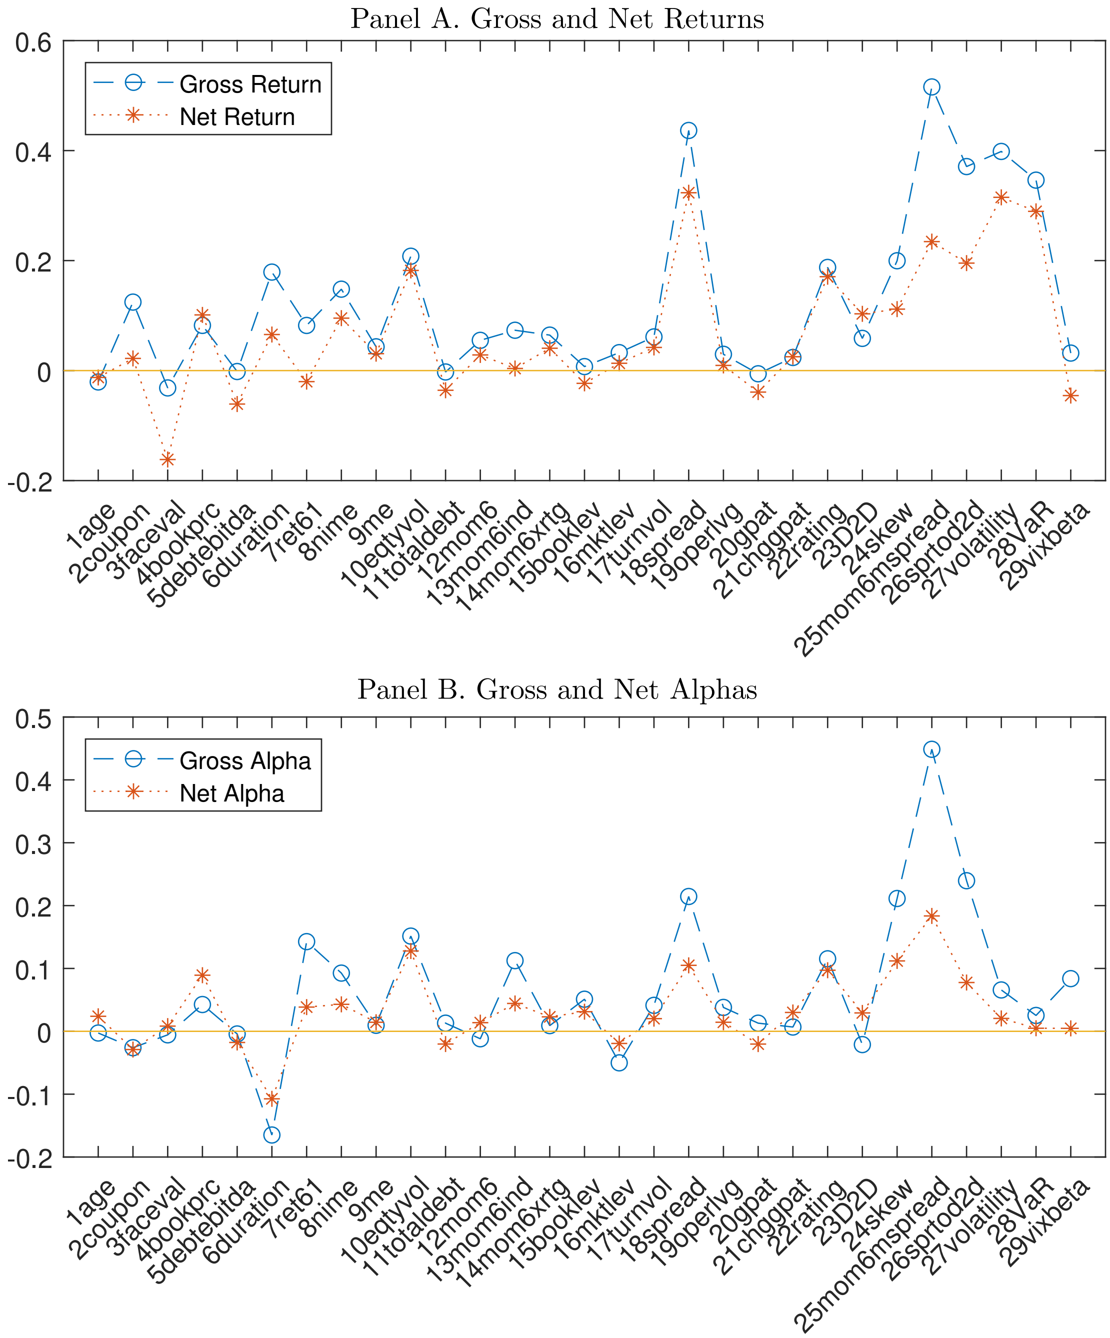

These figures plot the gross/net returns and CAPM α for each underlying signal that forms the basis of machine-learning algorithm. To compute net returns and α, we use the optimal value that maximizes the net α.

Figure 7: Average Excess Returns and CAPM Alphas of the ML Strategies: Volume of $100K

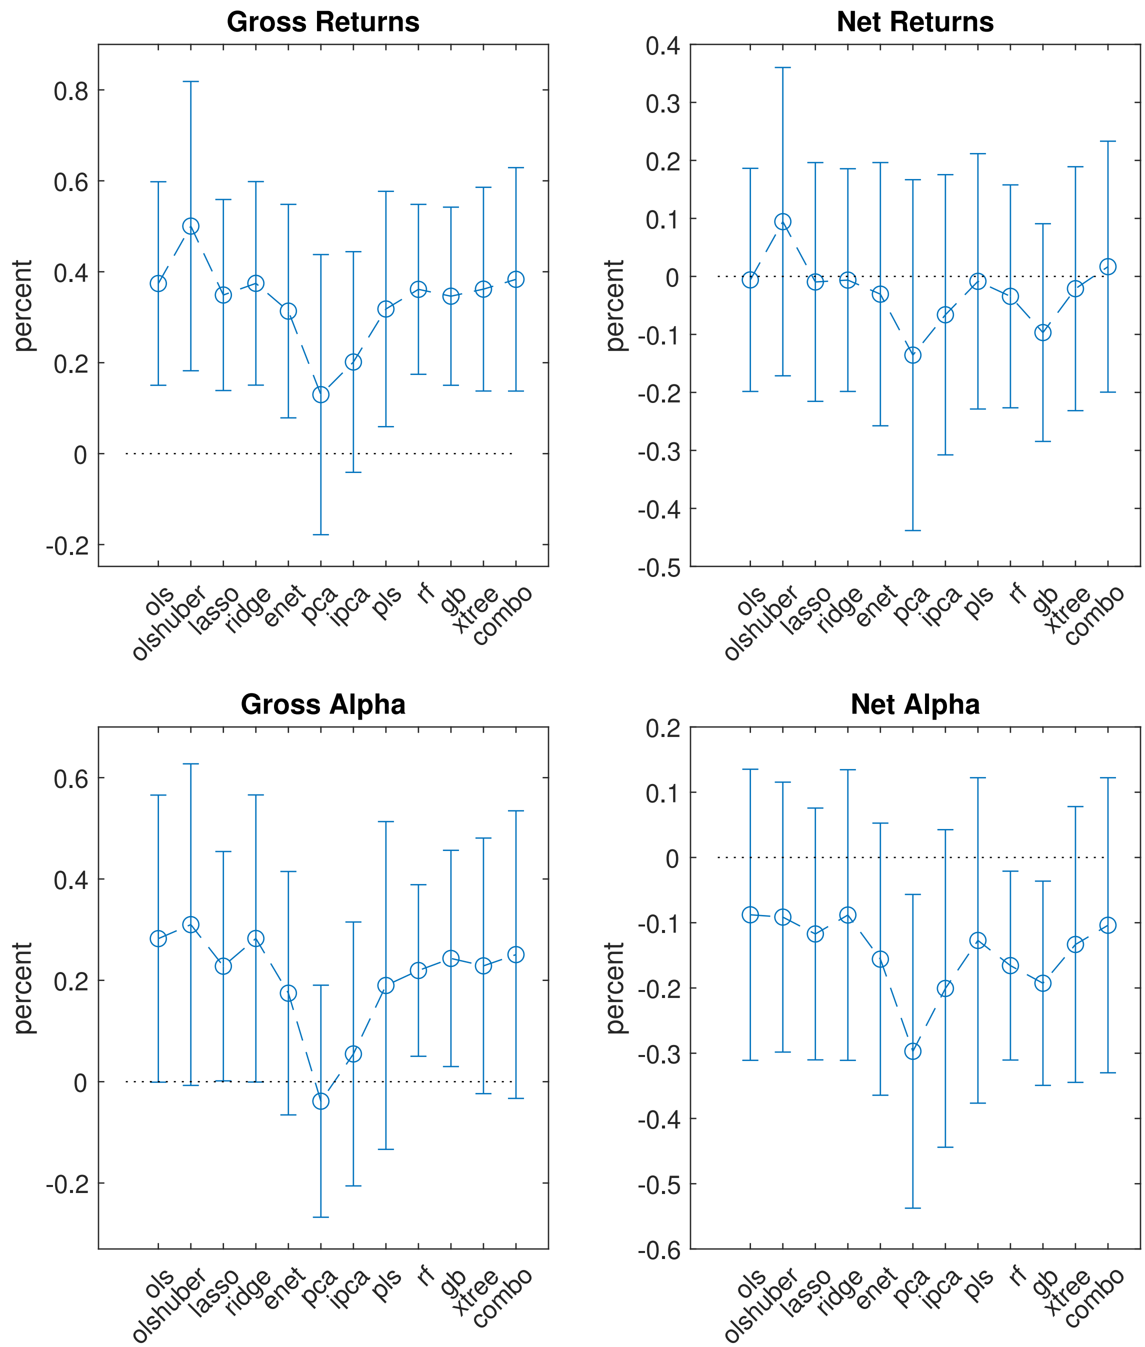

This figure shows the point estimate and associated two standard error bars for the long-short portfolios based on the expected returns generated by the machine learning algorithms. Gross returns and alphas are before transaction costs, and net returns and alphas are after costs. Transaction costs are calculated using the minimum volume threshold of $100,000.

Figure 8: Average Excess Returns and CAPM Alphas of the ML Strategies: Optimal Volume

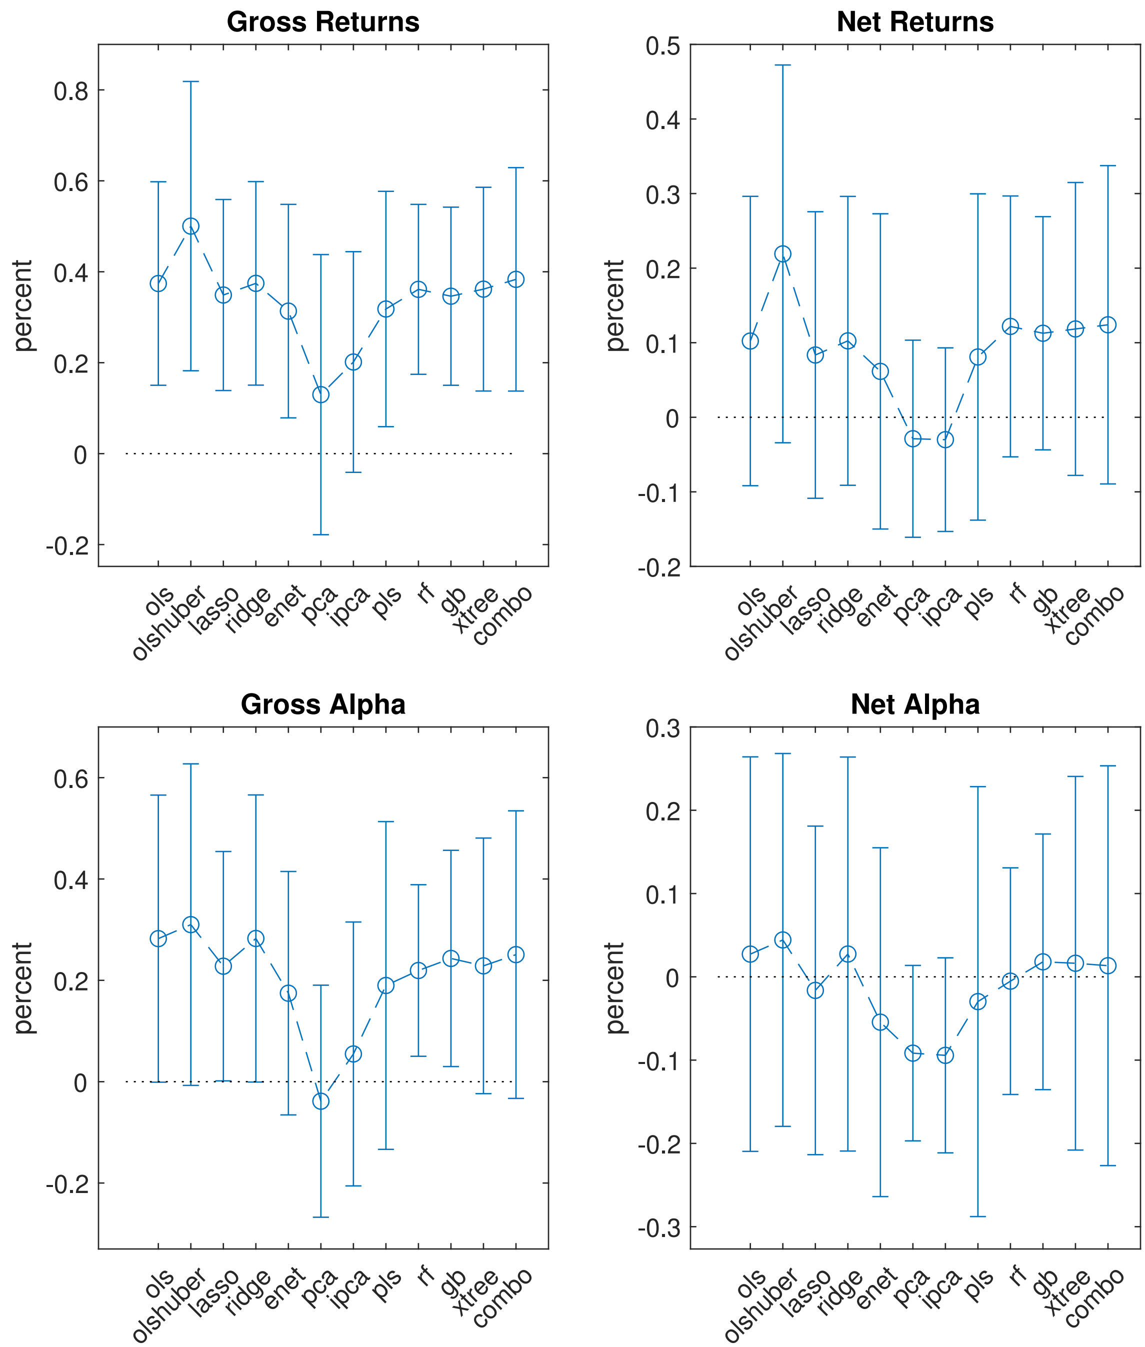

This figure shows the point estimate and associated two standard error bars for the long-short portfolios based on the expected returns generated by the machine learning algorithms. Gross returns and alphas are before transaction costs, and net returns and alphas are after costs. Transaction costs are calculated using the optimal threshold that maximizes the net alpha.

Figure 9: Effect of Transaction Costs: ML Strategies

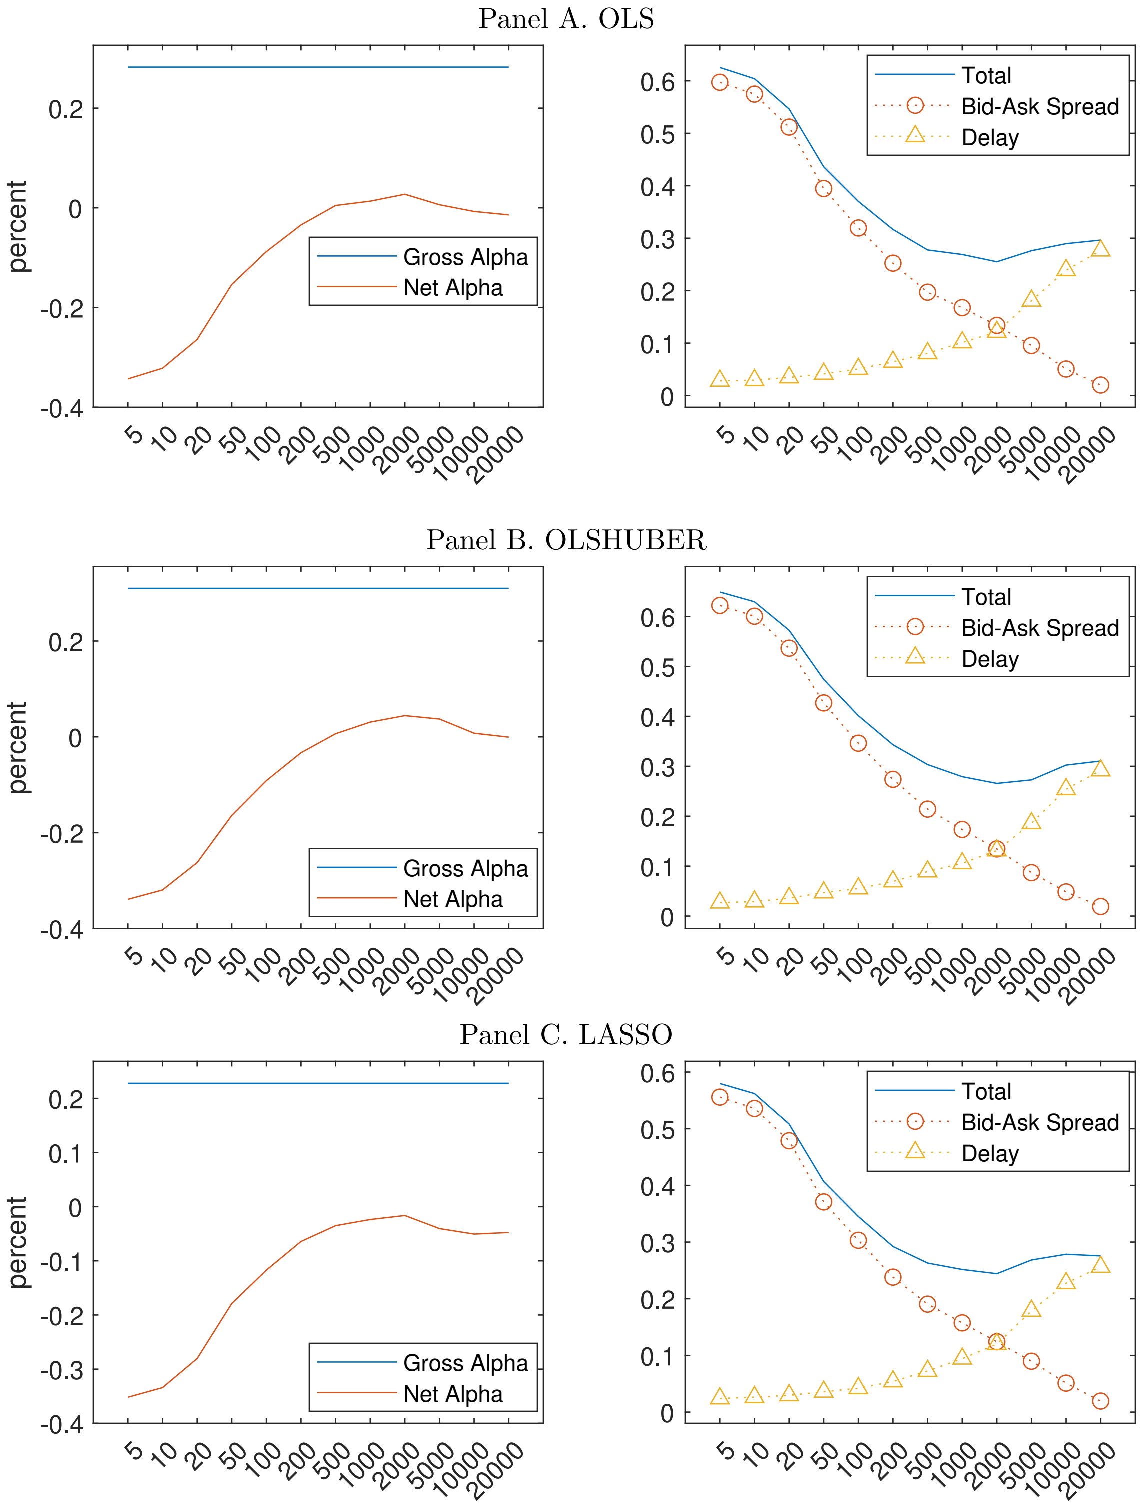

This figure plots the bond CAPM alphas on the long-short strategies before and after accounting for transaction costs (left panels). The transaction costs are decomposed into bid-ask spreads and delays (right panels). Values on the x-axis are in thousand dollars.

Figure 9, Continued

Panel D. RIDGE

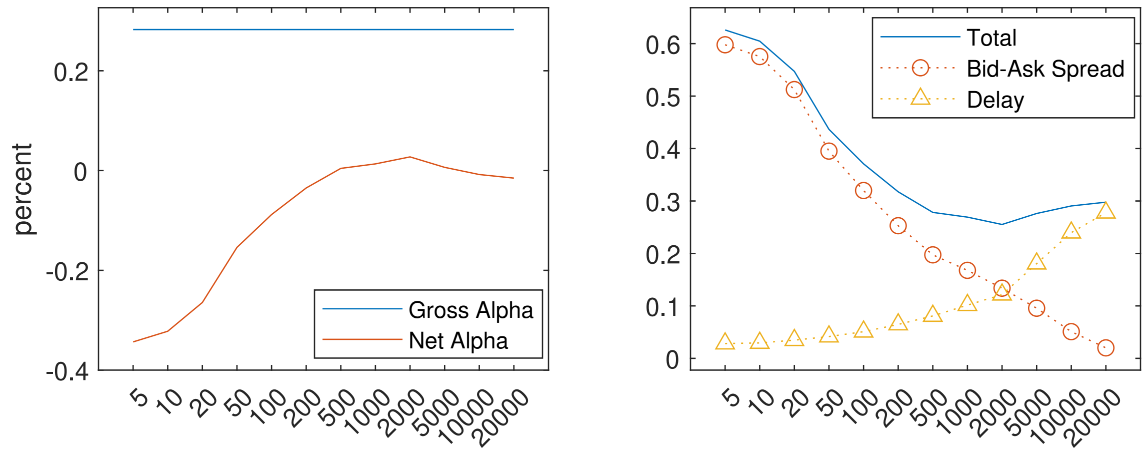

Panel E. ENET

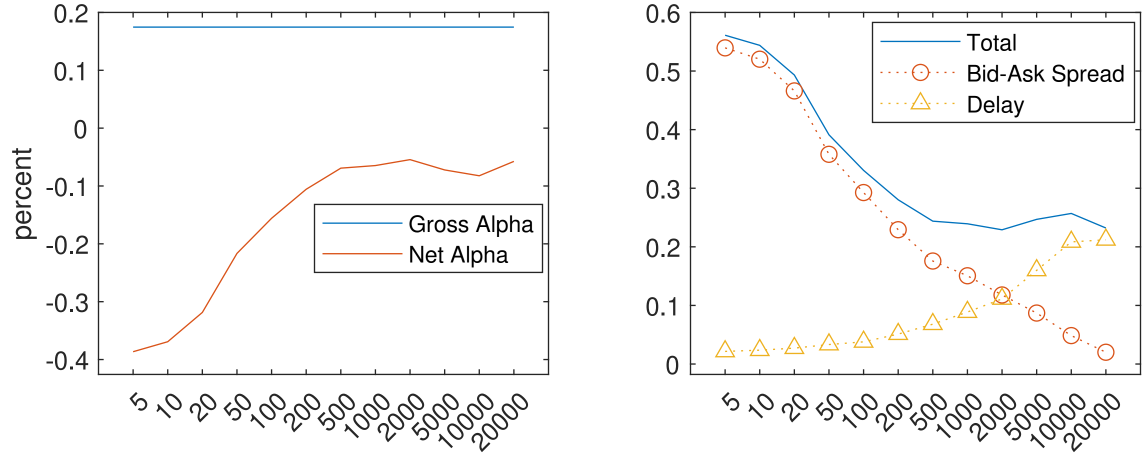

Panel F. PCA

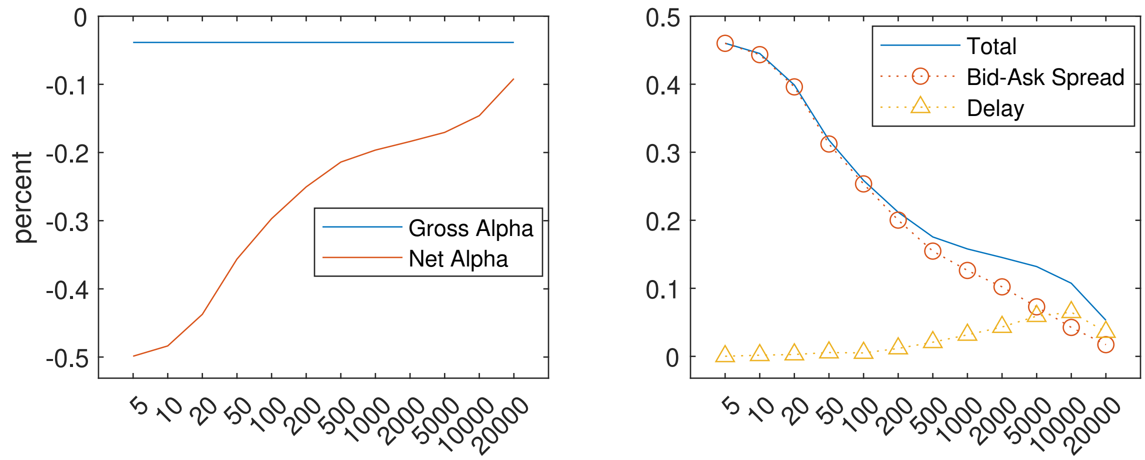

Figure 9, Continued

Panel G. IPCA

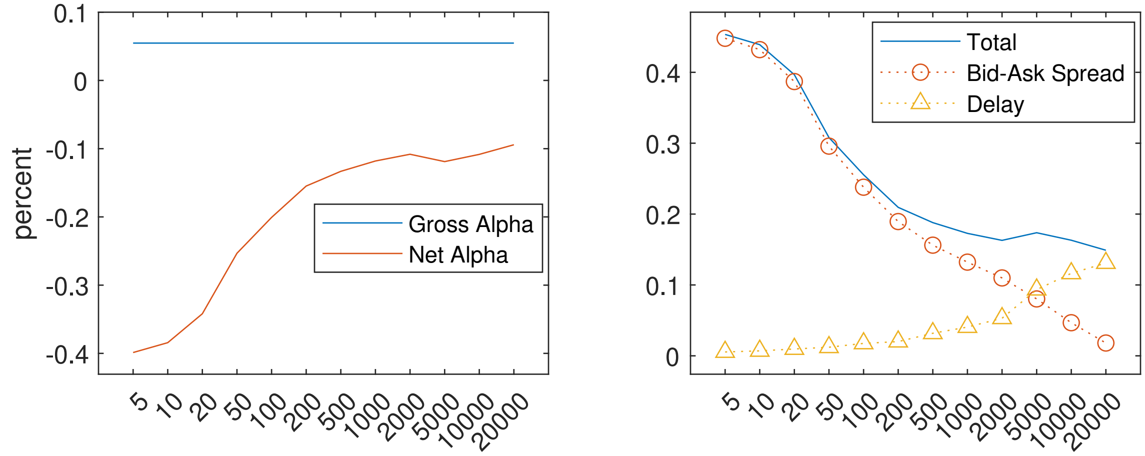

Panel H. PLS

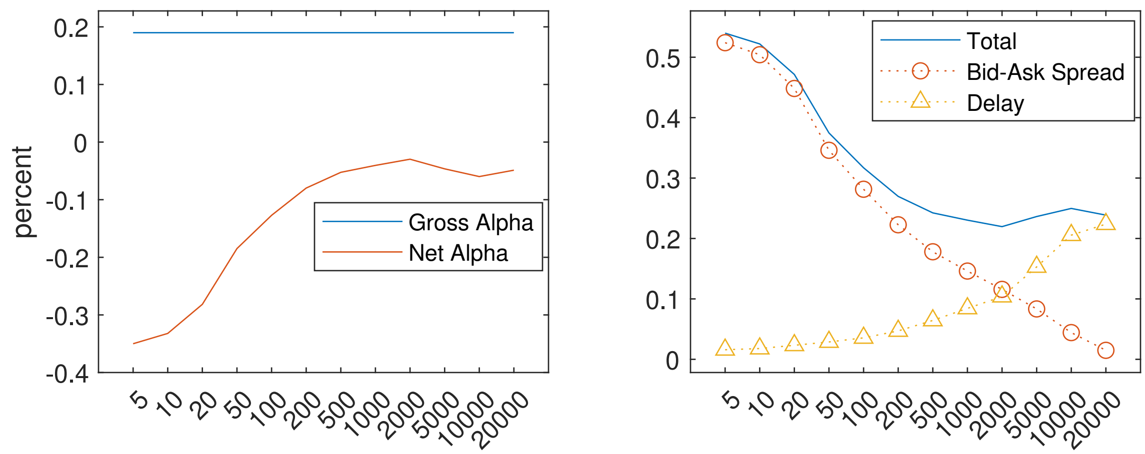

Panel I. RF

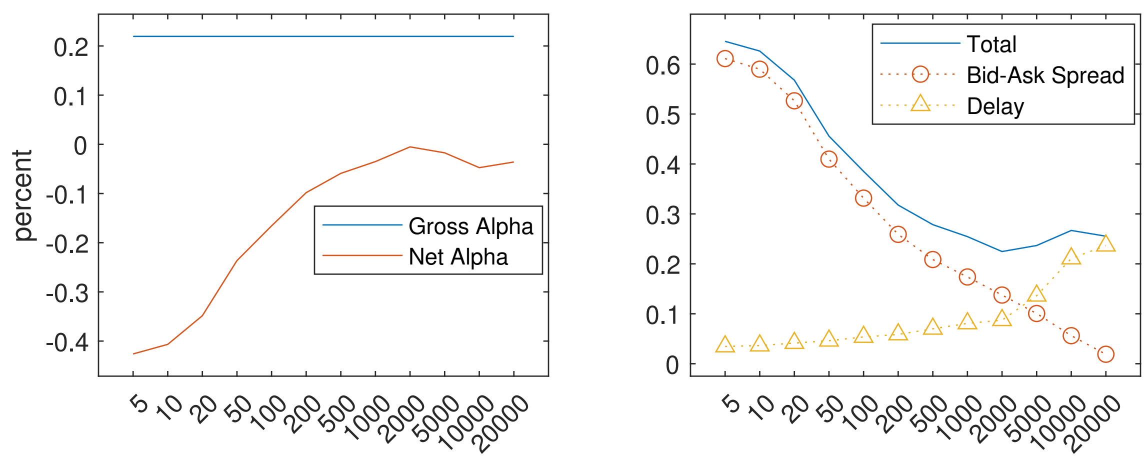

Figure 9, Continued

Panel J. GB

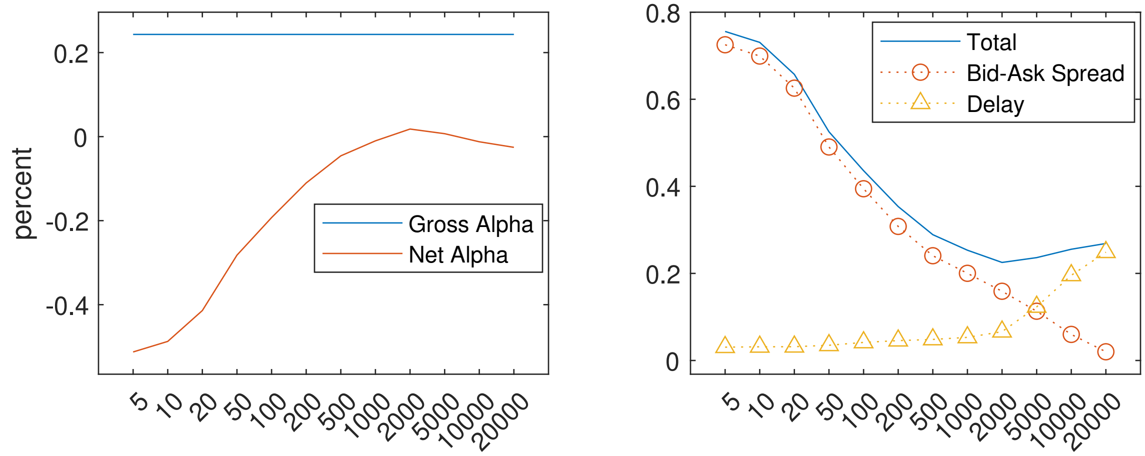

Panel K. XTREE

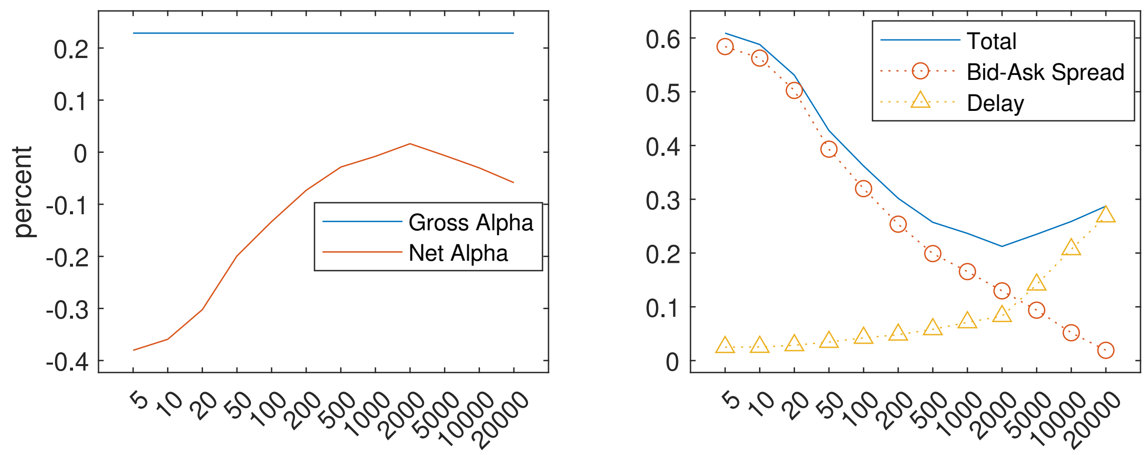

Panel L. COMBO

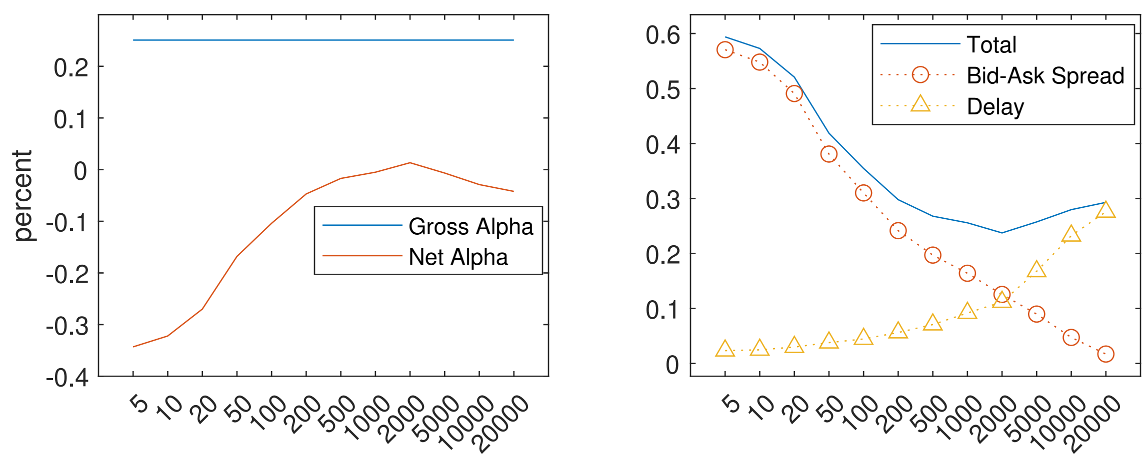

Figure 10: Subperiod Analysis for COMBO

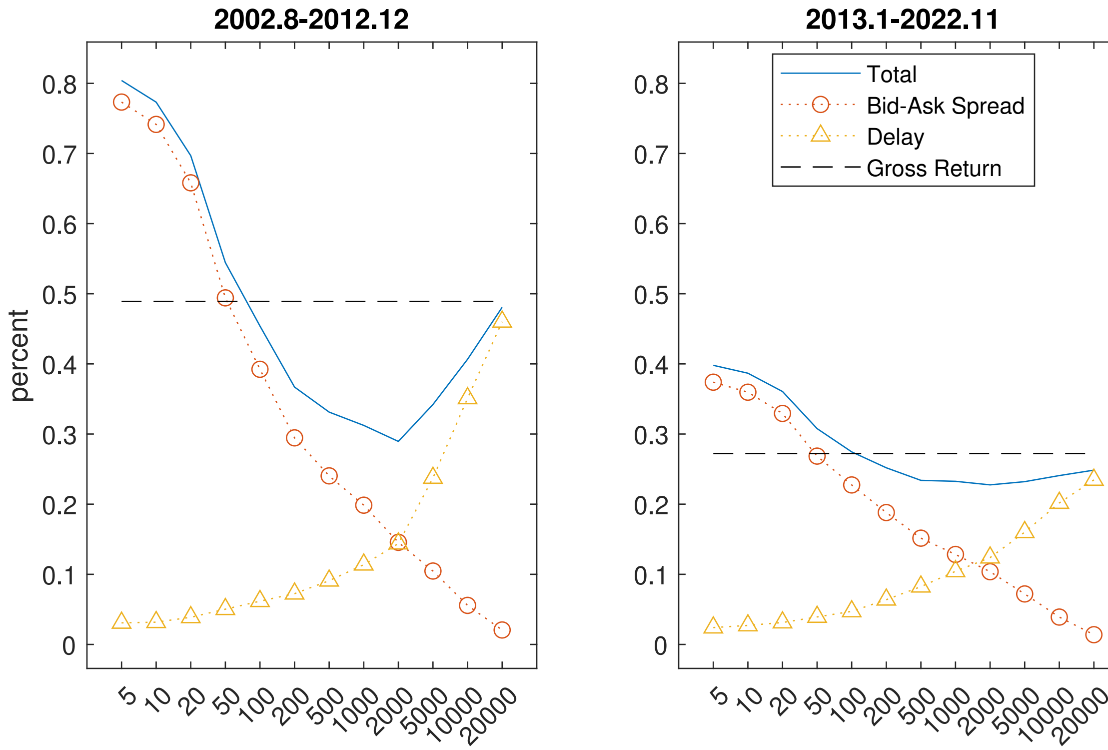

This figure shows the decomposition of transaction costs for COMBO strategies using the two subperiods. In this figure, total costs are the difference between gross and net average returns. The bid-ask spread cost is the difference between the gross average return and an alternative net average return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

Figure 11: Turnover Rate, Gross and Net CAPM α

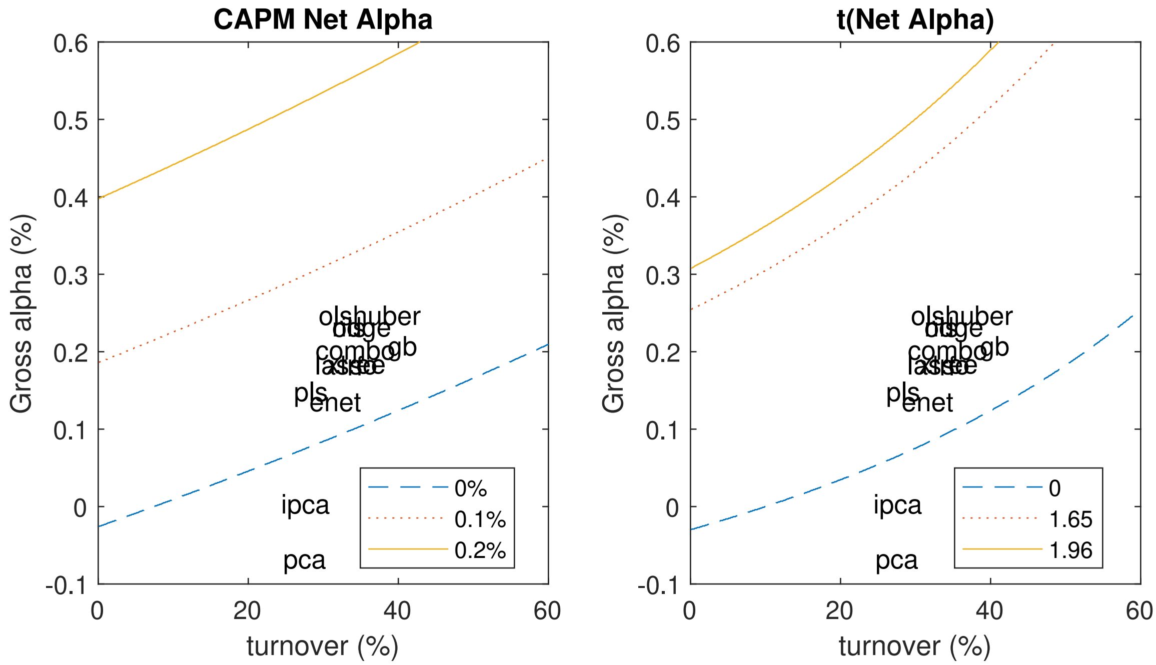

The figures plot the combination of gross CAPM α and portfolio turnover rate that matches the target values of net CAPM α and the associated t-statistics. The boundaries are estimated by regressing the net CAPM α's and t-statistics on gross α, portfolio turnover rate, and the product of the two. The regression uses 41 strategies including 29 individual signals and 12 machine learning algorithms.

Figure 12: Mutual Fund Alphas

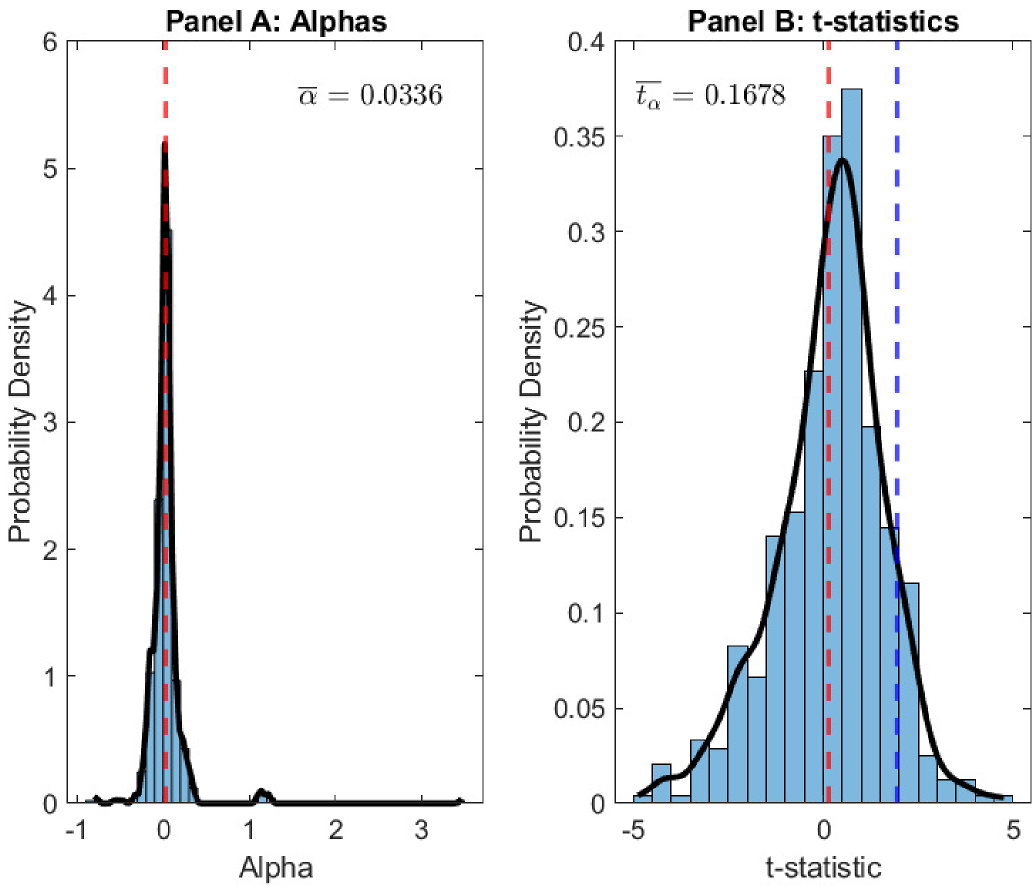

This figure plots the cross-sectional distribution of the corporate bond net of costs single-factor $MKTB_{Net}$ alphas and associated t-statistics. The dashed red lines indicate the mean values for the alphas (t-statistics). The dashed blue line represents the cut-off value for the 95% level of significance ($t = 1.96$). 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)	
Duan, Li, and Wen (2021)	
Friewald and Nagler (2016)	
Gebhardt, Hvidkjaer, and Swaminathan (2005a)	
Gebhardt, Hvidkjaer, and Swaminathan (2005b)	
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)
Bartram, Grinblatt, and Nozawa (2021)	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 (2021)	Fixed at 20 to 80bps
Houweling and Van Zundert (2017)	Maturity-rating, following Chen et al. (2007)
Israel, Palhares, and Richardson (2017)	Maturity-rating, following Chen et al. (2007)
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: 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	$-R^{sh}$	$-R^{sb}$	$-R^{hb} + R^f$	$-R^{hh} + R^f$

Table 3: Summary Statistics for Returns: Volume of $100K, $1M, and $10M

Variable	N	Mean	Std.	p1	p10	p50	p90	p99	NoTrade(%)
R^{hh}	746,498	0.30	3.65	-9.54	-2.36	0.29	2.91	9.32	
Panel A. $100K									
R^{bh}	746,498	-0.10	3.82	-8.73	-2.36	0.01	1.94	7.23	12.82
R^{hs}	746,498	0.50	4.20	-10.45	-2.59	0.40	3.67	11.14	12.38
R^{bs}	746,498	-0.13	4.13	-9.72	-2.68	0.01	2.14	8.19	19.77
R^{sh}	746,498	0.35	4.32	-7.88	-1.80	0.18	2.60	8.53	11.78
R^{hb}	746,498	0.95	4.40	-9.52	-2.10	0.69	4.23	12.33	13.42
R^{sb}	746,498	0.85	4.83	-7.83	-1.54	0.50	3.54	10.79	19.33
Panel B. $1M									
R^{bh}	746,498	0.06	3.07	-7.00	-1.48	0.03	1.58	6.61	32.46
R^{hs}	746,498	0.49	4.25	-10.70	-2.57	0.40	3.64	11.33	32.06
R^{bs}	746,498	0.09	3.48	-8.11	-1.75	0.06	1.90	7.91	46.50
R^{sh}	746,498	0.23	3.44	-6.69	-1.24	0.09	1.88	7.12	31.48
R^{hb}	746,498	0.68	4.40	-10.16	-2.36	0.51	3.86	11.81	33.08
R^{sb}	746,498	0.48	3.99	-7.31	-1.20	0.15	2.59	8.98	45.51
Panel C. $10M									
R^{bh}	746,498	0.10	1.75	-2.84	0.00	0.02	0.37	3.24	84.25
R^{hs}	746,498	0.33	3.98	-10.23	-2.47	0.30	3.11	10.05	82.15
R^{bs}	746,498	0.11	1.98	-3.28	0.00	0.02	0.37	3.65	93.48
R^{sh}	746,498	0.14	2.09	-2.75	0.00	0.03	0.40	3.63	82.07
R^{hb}	746,498	0.36	3.98	-10.08	-2.43	0.32	3.14	10.12	84.36
R^{sb}	746,498	0.17	2.34	-3.04	0.00	0.03	0.41	4.26	93.06

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, and C correspond to the statistics with volume thresholds of $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 4: Summary Statistics for Returns: Volume of $100K, $1M, and $10M: Duration-Adjusted Returns

Variable	N	Mean	Std.	p1	p10	p50	p90	p99	NoTrade(%)
R^{hh}	746,464	0.20	3.42	-9.33	-1.47	0.13	2.04	8.68	
Panel A. $100K									
R^{bh}	746,464	-0.28	3.83	-8.77	-2.46	-0.05	1.55	7.38	12.82
R^{hs}	746,464	0.32	4.26	-12.18	-2.42	0.24	3.33	10.64	12.38
R^{bs}	746,464	-0.31	4.29	-11.15	-2.94	-0.01	2.04	8.51	19.77
R^{sh}	746,464	0.16	4.32	-8.01	-1.87	0.08	2.17	8.71	11.78
R^{hb}	746,464	0.77	4.42	-10.78	-1.96	0.52	3.86	11.84	13.42
R^{sb}	746,464	0.53	4.94	-9.25	-1.96	0.25	3.21	10.88	19.34
Panel B. $1M									
R^{bh}	746,464	-0.10	3.05	-7.08	-1.52	0.00	1.02	6.49	32.46
R^{hs}	746,464	0.38	4.22	-11.31	-2.32	0.27	3.26	10.86	32.07
R^{bs}	746,464	-0.05	3.52	-8.58	-1.80	0.01	1.47	7.95	46.50
R^{sh}	746,464	0.06	3.42	-6.85	-1.24	0.03	1.30	7.05	31.48
R^{hb}	746,464	0.57	4.37	-10.77	-2.11	0.39	3.50	11.34	33.08
R^{sb}	746,464	0.25	4.02	-8.04	-1.33	0.09	2.02	8.93	45.51
Panel C. $10M									
R^{bh}	746,464	0.06	1.73	-2.81	0.00	0.02	0.32	2.64	84.25
R^{hs}	746,464	0.30	3.98	-10.25	-2.37	0.27	2.99	9.93	82.15
R^{bs}	746,464	0.06	1.97	-3.13	0.00	0.02	0.32	3.10	93.48
R^{sh}	746,464	0.09	2.06	-2.68	0.00	0.02	0.34	3.05	82.07
R^{hb}	746,464	0.34	3.97	-10.09	-2.34	0.29	3.03	10.02	84.36
R^{sb}	746,464	0.11	2.31	-2.92	0.00	0.02	0.36	3.57	93.06

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, and C corresponds to the statistics with the volume threshold of $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 5: Performance of Underlying Signals

Signal	Sign	Excess Returns Gross	Excess Returns Net $100K	Excess Returns Net Optimal	CAPM α Gross	CAPM α Net $100K	CAPM α Net Optimal	Optimal Volume	Turnover (%)
age	+	-0.021	-0.105	-0.013	-0.003	-0.085	0.024	20000	6.84
		(-0.98)	(-4.50)	(-0.46)	(-0.09)	(-3.09)	(1.13)		
coupon	+	0.125	0.091	0.022	-0.026	-0.058	-0.029	20000	5.57
		(1.29)	(0.95)	(0.49)	(-0.38)	(-0.83)	(-0.72)		
faceval	-	-0.032	-0.073	-0.162	-0.006	-0.043	0.008	20000	4.96
		(-0.74)	(-1.64)	(-1.99)	(-0.10)	(-0.77)	(0.21)		
bookprc	+	0.082	0.030	0.101	0.043	-0.008	0.089	20000	7.72
		(0.91)	(0.34)	(2.10)	(0.47)	(-0.09)	(1.65)		
debtebitda	+	-0.002	-0.066	-0.061	-0.004	-0.069	-0.018	20000	7.76
		(-0.03)	(-1.21)	(-1.37)	(-0.07)	(-1.09)	(-0.36)		
duration	+	0.179	0.121	0.066	-0.165	-0.218	-0.107	20000	6.78
		(1.07)	(0.73)	(0.70)	(-1.73)	(-2.22)	(-1.66)		
ret61	+	0.082	-0.282	-0.020	0.143	-0.219	0.039	10000	29.11
		(1.21)	(-3.06)	(-0.39)	(1.98)	(-2.53)	(0.86)		
nime	-	0.148	0.007	0.095	0.093	-0.050	0.043	5000	14.11
		(1.89)	(0.08)	(1.33)	(1.30)	(-0.56)	(0.61)		
me	-	0.043	-0.012	0.031	0.010	-0.042	0.014	5000	7.62
		(0.36)	(-0.10)	(0.31)	(0.08)	(-0.33)	(0.13)		
eqtyvol	+	0.208	0.110	0.182	0.152	0.047	0.128	2000	12.17
		(1.48)	(0.78)	(1.40)	(1.06)	(0.32)	(1.00)		
totaldebt	-	-0.003	-0.043	-0.036	0.013	-0.027	-0.020	500	5.86
		(-0.04)	(-0.65)	(-0.55)	(0.16)	(-0.34)	(-0.25)		
mom6	-	0.055	-0.315	0.028	-0.012	-0.379	0.014	10000	29.99
		(0.40)	(-3.06)	(0.41)	(-0.09)	(-3.50)	(0.22)		
mom6ind	+	0.073	-0.227	0.004	0.112	-0.180	0.044	10000	28.14
		(1.31)	(-2.87)	(0.12)	(1.74)	(-2.21)	(1.12)		
mom6xrtg	-	0.065	-0.262	0.041	0.009	-0.313	0.023	10000	28.01
		(0.54)	(-2.89)	(0.61)	(0.07)	(-3.12)	(0.37)		
booklev	-	0.007	-0.052	-0.023	0.051	-0.013	0.031	20000	7.27
		(0.20)	(-1.35)	(-0.66)	(1.37)	(-0.33)	(1.17)		
mktlev	+	0.033	-0.008	0.013	-0.050	-0.086	-0.019	20000	7.41
		(0.37)	(-0.10)	(0.23)	(-0.59)	(-1.01)	(-0.32)		
turnvol	+	0.061	0.034	0.042	0.041	0.014	0.020	2000	5.34
		(1.23)	(0.67)	(0.83)	(0.76)	(0.25)	(0.38)		

Table 5, Continued

Signal	Sign	Excess Returns			CAPM α			Optimal Volume	Turnover (%)
		Gross	Net $100K	Net Optimal	Gross	Net $100K	Net Optimal		
spread	+	0.437	0.252	0.323	0.215	0.028	0.105	2000	13.67
		(2.37)	(1.38)	(1.88)	(1.72)	(0.23)	(0.92)		
operlvg	+	0.030	-0.027	0.009	0.038	-0.016	0.014	2000	8.63
		(0.73)	(-0.59)	(0.22)	(0.85)	(-0.32)	(0.32)		
gpat	+	-0.006	-0.053	-0.039	0.013	-0.033	-0.020	1000	6.80
		(-0.08)	(-0.75)	(-0.55)	(0.17)	(-0.45)	(-0.27)		
chggpat	-	0.024	-0.057	0.025	0.007	-0.072	0.030	20000	9.10
		(0.54)	(-1.34)	(1.24)	(0.17)	(-1.76)	(1.58)		
rating	+	0.188	0.153	0.171	0.115	0.084	0.097	2000	6.30
		(1.32)	(1.07)	(1.22)	(0.77)	(0.56)	(0.66)		
D2D	-	0.059	-0.017	0.103	-0.021	-0.099	0.029	5000	12.14
		(0.45)	(-0.13)	(0.90)	(-0.15)	(-0.70)	(0.25)		
skew	-	0.200	0.031	0.112	0.211	0.031	0.112	2000	13.03
		(1.79)	(0.42)	(1.42)	(1.67)	(0.38)	(1.29)		
6mspread	-	0.516	-0.006	0.235	0.449	-0.077	0.184	5000	37.03
		(3.68)	(-0.06)	(2.35)	(3.00)	(-0.78)	(1.69)		
sprtod2d	+	0.371	0.140	0.195	0.240	0.018	0.078	2000	16.57
		(3.30)	(1.39)	(2.00)	(2.01)	(0.18)	(0.77)		
volatility	+	0.398	0.304	0.315	0.066	-0.031	0.021	5000	10.51
		(2.05)	(1.56)	(1.87)	(0.86)	(-0.43)	(0.33)		
VaR	+	0.346	0.261	0.290	0.025	-0.054	0.005	5000	9.48
		(1.82)	(1.36)	(1.70)	(0.28)	(-0.56)	(0.06)		
vixbeta	-	0.032	-0.131	-0.045	0.084	-0.089	0.005	20000	11.10
		(0.46)	(-2.02)	(-1.16)	(1.12)	(-1.39)	(0.16)		

This table reports the average excess returns and CAPM alphas of the long-short portfolios built on bond characteristics. Bond characteristics are defined in Table A2. If 'Sign' is '-', then we use the negative of the characteristic. Each month, we select the top and bottom 20% of bonds in terms of bond characteristics and form a long-short strategy. Gross returns and alphas are before transaction costs and net returns are after costs. Net costs are calculated using both a transaction volume of at least $100,000 and the optimal values reported in the "Optimal Volume (thousand dollars)" column. 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 through November 2022.

Table 6: Performance of ML Strategies

Signal	Excess Returns			CAPM α			Optimal Volume	Turnover (%)
	Gross	Net		Gross	Net			
		\$100K	Optimal		\$100K	Optimal		
OLS	0.374	-0.006	0.102	0.282	-0.088	0.027	2000	27.43
	(3.34)	(-0.06)	(1.05)	(1.99)	(-0.79)	(0.23)		
OLSHUBER	0.500	0.094	0.219	0.310	-0.091	0.044	2000	25.75
	(3.15)	(0.71)	(1.73)	(1.95)	(-0.88)	(0.40)		
LASSO	0.349	-0.010	0.084	0.228	-0.117	-0.016	2000	25.65
	(3.32)	(-0.09)	(0.87)	(2.01)	(-1.22)	(-0.17)		
RIDGE	0.375	-0.006	0.103	0.283	-0.088	0.027	2000	27.44
	(3.35)	(-0.07)	(1.06)	(2.00)	(-0.79)	(0.23)		
ENET	0.313	-0.031	0.062	0.175	-0.156	-0.054	2000	24.91
	(2.67)	(-0.27)	(0.58)	(1.45)	(-1.49)	(-0.52)		
PCA	0.130	-0.136	-0.029	-0.039	-0.297	-0.092	20000	16.58
	(0.84)	(-0.90)	(-0.44)	(-0.34)	(-2.47)	(-1.74)		
IPCA	0.202	-0.066	-0.030	0.055	-0.201	-0.094	20000	16.19
	(1.66)	(-0.55)	(-0.49)	(0.42)	(-1.65)	(-1.61)		
PLS	0.318	-0.009	0.081	0.190	-0.127	-0.030	2000	23.27
	(2.46)	(-0.08)	(0.74)	(1.17)	(-1.02)	(-0.23)		
RF	0.361	-0.034	0.122	0.219	-0.166	-0.005	2000	29.26
	(3.87)	(-0.36)	(1.39)	(2.59)	(-2.29)	(-0.08)		
GB	0.346	-0.097	0.113	0.243	-0.193	0.018	2000	33.83
	(3.54)	(-1.03)	(1.44)	(2.28)	(-2.46)	(0.23)		
XTREE	0.362	-0.021	0.118	0.229	-0.133	0.016	2000	27.03
	(3.23)	(-0.20)	(1.21)	(1.81)	(-1.26)	(0.15)		
COMBO	0.383	0.017	0.124	0.251	-0.104	0.013	2000	25.53
	(3.12)	(0.16)	(1.16)	(1.77)	(-0.92)	(0.11)		

This table reports the average excess returns and CAPM 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 are after costs. Net costs are calculated using both a transaction volume of at least \$100,000 and the optimal values reported in the "Optimal Volume (thousand dollars)" column. 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 through November 2022.

Table 7: Performance of ML Strategies: Duration-Adjusted Returns

Signal	Duration-Adjusted Returns			CAPM α			Optimal	Turnover
	Gross	Net		Gross	Net		Volume	(%)
		\$100K	Optimal		\$100K	Optimal		
OLS	0.365	-0.013	0.080	0.344	-0.026	0.061	2000	27.43
	(3.65)	(-0.15)	(0.90)	(2.79)	(-0.24)	(0.55)		
OLSHUBER	0.490	0.089	0.193	0.392	-0.007	0.096	2000	25.75
	(3.24)	(0.76)	(1.69)	(2.48)	(-0.05)	(0.80)		
LASSO	0.363	0.002	0.084	0.316	-0.037	0.038	2000	25.65
	(3.84)	(0.02)	(0.94)	(2.88)	(-0.34)	(0.37)		
RIDGE	0.365	-0.014	0.080	0.344	-0.027	0.061	2000	27.44
	(3.65)	(-0.16)	(0.90)	(2.79)	(-0.25)	(0.56)		
ENET	0.334	-0.015	0.067	0.280	-0.060	0.015	2000	24.91
	(3.30)	(-0.14)	(0.71)	(2.47)	(-0.53)	(0.14)		
PCA	0.170	-0.119	0.004	0.077	-0.208	-0.035	20000	16.58
	(1.21)	(-0.85)	(0.06)	(0.58)	(-1.46)	(-0.60)		
IPCA	0.136	-0.134	-0.041	0.101	-0.161	-0.063	20000	16.19
	(1.69)	(-1.51)	(-0.93)	(1.11)	(-1.74)	(-1.42)		
PLS	0.314	-0.024	0.059	0.272	-0.058	0.018	2000	23.27
	(2.81)	(-0.25)	(0.62)	(1.96)	(-0.50)	(0.15)		
RF	0.396	0.021	0.130	0.318	-0.050	0.057	2000	29.26
	(4.40)	(0.23)	(1.57)	(3.09)	(-0.46)	(0.61)		
GB	0.381	-0.053	0.118	0.339	-0.091	0.073	2000	33.83
	(4.52)	(-0.65)	(1.67)	(3.39)	(-1.00)	(0.89)		
XTREE	0.358	-0.030	0.096	0.307	-0.065	0.056	2000	27.03
	(3.50)	(-0.29)	(1.03)	(2.49)	(-0.54)	(0.49)		
COMBO	0.384	0.014	0.102	0.339	-0.022	0.061	2000	25.53
	(3.64)	(0.14)	(1.09)	(2.71)	(-0.20)	(0.54)		

This table reports the average excess returns and CAPM 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 are after costs. Net costs are calculated using both a transaction volume of at least \$100,000 and the optimal values reported in the "Optimal Volume (thousand dollars)" column. 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 through November 2022.

Table 8: Decomposition of Transaction Costs

	$100K			Optimal Volume		
Signal	Total	BidAsk	Delay	Total	BidAsk	Delay
OLS	0.370	0.320	0.051	0.255	0.134	0.121
OLSHUBER	0.401	0.346	0.055	0.266	0.134	0.131
LASSO	0.345	0.303	0.042	0.244	0.124	0.120
RIDGE	0.371	0.320	0.051	0.255	0.134	0.121
ENET	0.331	0.293	0.038	0.229	0.118	0.111
PCA	0.259	0.254	0.005	0.053	0.017	0.036
IPCA	0.256	0.238	0.018	0.149	0.018	0.131
PLS	0.317	0.282	0.036	0.220	0.116	0.104
RF	0.385	0.332	0.054	0.225	0.138	0.087
GB	0.436	0.395	0.042	0.225	0.159	0.066
XTREE	0.362	0.320	0.042	0.212	0.130	0.083
COMBO	0.355	0.310	0.045	0.237	0.126	0.112

This table reports the two components of transaction costs, which 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.

Table 9: Decomposition of Transaction Costs: Individual Signals

	$100K			Optimal Volume		
Signal	Total	BidAsk	Delay	Total	BidAsk	Delay
age	0.082	0.053	0.030	-0.027	0.011	-0.037
coupon	0.032	0.029	0.004	0.003	0.006	-0.004
faceval	0.037	0.029	0.008	-0.014	0.004	-0.018
bookprc	0.051	0.069	-0.018	-0.047	0.006	-0.053
debtebitda	0.065	0.064	0.001	0.014	0.008	0.006
duration	0.053	0.056	-0.004	-0.058	0.011	-0.068
ret61	0.362	0.349	0.013	0.104	0.050	0.054
nime	0.143	0.134	0.008	0.050	0.047	0.003
me	0.052	0.056	-0.004	-0.005	0.020	-0.025
eqtyvol	0.105	0.102	0.003	0.024	0.052	-0.028
totaldebt	0.040	0.036	0.004	0.033	0.026	0.008
mom6	0.367	0.385	-0.018	-0.026	0.064	-0.090
mom6ind	0.292	0.280	0.012	0.068	0.049	0.019
mom6xrtg	0.322	0.349	-0.026	-0.014	0.062	-0.076
booklev	0.064	0.062	0.001	0.020	0.007	0.013
mktlev	0.036	0.049	-0.013	-0.031	0.007	-0.038
turnvol	0.028	0.024	0.004	0.021	0.013	0.008
spread	0.186	0.157	0.029	0.110	0.067	0.042
operlvg	0.054	0.057	-0.003	0.024	0.028	-0.004
gpat	0.046	0.040	0.006	0.033	0.024	0.010
chggpat	0.079	0.081	-0.002	-0.023	0.010	-0.033
rating	0.031	0.034	-0.002	0.018	0.017	0.001
D2D	0.078	0.099	-0.021	-0.050	0.028	-0.079
skew	0.180	0.151	0.029	0.099	0.060	0.039
6mspread	0.526	0.458	0.068	0.265	0.112	0.154
sprtod2d	0.222	0.191	0.031	0.162	0.081	0.081
volatility	0.097	0.096	0.001	0.045	0.033	0.012
VaR	0.079	0.081	-0.003	0.020	0.028	-0.008
vixbeta	0.173	0.150	0.023	0.079	0.009	0.071

This table reports the two components of transaction costs, which 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.

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	1.41	0.92		0.06	0.01		7.97
	Middle	1.40	1.00	0.20	0.08	0.07	0.13	16.28
	High		0.84	0.23		0.09	0.26	34.22
		Panel C. Half-Spread Cost (%)			Panel D. Delay Cost (%)			
Turnover	Low	0.01	0.01		0.04	0.00		7.97
	Middle	0.02	0.03	0.08	0.06	0.04	0.05	16.28
	High		0.07	0.13		0.03	0.13	34.22
Average α		-0.02	0.07	0.21	-0.02	0.07	0.21	

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 41 strategies based on individual signals and ML. The signals/strategies are classified into three groups based independently on their turnover rate (calculated using threshold = $5K) and gross CAPM α. The empty cells do not have any signal/strategy that falls into the category.

Table 11: Robustness: Long-Only Strategies

Signal	Excess Returns			CAPM α			Optimal Volume	Turnover (%)
	Gross	Net		Gross	Net			
		\$100K	Optimal		\$100K	Optimal		
OLS	0.451	0.291	0.320	0.129	-0.015	0.026	500	29.67
	(2.69)	(1.86)	(2.08)	(1.99)	(-0.29)	(0.48)		
OLSHUBER	0.530	0.342	0.381	0.173	0.004	0.060	1000	27.16
	(2.64)	(1.87)	(2.13)	(1.74)	(0.05)	(0.76)		
LASSO	0.440	0.283	0.300	0.115	-0.023	0.011	1000	26.98
	(2.56)	(1.75)	(1.93)	(2.01)	(-0.44)	(0.21)		
RIDGE	0.452	0.292	0.320	0.129	-0.015	0.026	500	29.67
	(2.69)	(1.86)	(2.08)	(2.01)	(-0.28)	(0.49)		
ENET	0.429	0.277	0.306	0.094	-0.039	0.000	500	26.88
	(2.39)	(1.65)	(1.85)	(1.49)	(-0.71)	(-0.00)		
PCA	0.306	0.193	0.218	-0.035	-0.134	-0.086	2000	22.60
	(1.54)	(1.03)	(1.24)	(-0.44)	(-1.70)	(-1.14)		
IPCA	0.374	0.262	0.273	0.016	-0.080	-0.041	2000	21.84
	(2.17)	(1.60)	(1.78)	(0.28)	(-1.46)	(-0.75)		
PLS	0.399	0.259	0.268	0.065	-0.060	-0.021	2000	23.27
	(2.24)	(1.56)	(1.73)	(0.85)	(-0.96)	(-0.31)		
RF	0.453	0.299	0.342	0.102	-0.035	0.040	2000	29.26
	(2.36)	(1.63)	(2.03)	(1.62)	(-0.57)	(0.71)		
GB	0.434	0.251	0.322	0.096	-0.071	0.032	2000	33.83
	(2.36)	(1.45)	(2.01)	(1.55)	(-1.29)	(0.62)		
XTREE	0.465	0.286	0.323	0.120	-0.039	0.029	2000	27.03
	(2.53)	(1.66)	(2.02)	(1.79)	(-0.63)	(0.49)		
COMBO	0.472	0.311	0.318	0.130	-0.015	0.024	2000	25.53
	(2.62)	(1.85)	(2.04)	(1.89)	(-0.27)	(0.42)		

This table reports the average excess returns (in excess of T-bills) and CAPM 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 are after costs. Net costs are calculated using both a transaction volume of at least \$100,000 and the optimal values reported in the "Optimal Volume (thousand dollars)" column. 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 through November 2022 (244 Months).

Table 12: Robustness: WRDS Data for Signals

Signal	Excess Returns			CAPM α			Optimal Volume	Turnover (%)
	Gross	Net		Gross	Net			
		\$100K	Optimal		\$100K	Optimal		
OLS	0.252	-0.096	0.141	0.259	-0.113	0.133	5000	34.19
	(3.47)	(-0.97)	(2.31)	(2.20)	(-1.38)	(1.55)		
OLSHUBER	0.262	-0.028	0.130	0.181	-0.126	0.022	5000	26.11
	(2.95)	(-0.23)	(1.42)	(1.57)	(-1.33)	(0.27)		
LASSO	0.303	-0.045	0.188	0.299	-0.072	0.176	5000	34.55
	(3.49)	(-0.45)	(2.70)	(1.71)	(-0.58)	(1.33)		
RIDGE	0.251	-0.098	0.143	0.259	-0.114	0.135	5000	34.29
	(3.45)	(-1.01)	(2.38)	(2.17)	(-1.35)	(1.56)		
ENET	0.306	-0.031	0.195	0.302	-0.059	0.183	5000	33.19
	(3.54)	(-0.29)	(2.71)	(1.79)	(-0.50)	(1.38)		
PCA	0.306	0.023	0.175	0.216	-0.088	0.081	5000	29.11
	(3.36)	(0.23)	(2.34)	(1.12)	(-0.67)	(0.70)		
IPCA	0.076	-0.238	-0.081	0.161	-0.165	-0.020	5000	30.49
	(0.69)	(-2.19)	(-1.03)	(1.69)	(-2.03)	(-0.40)		
PLS	0.310	0.011	0.195	0.257	-0.067	0.144	5000	30.08
	(3.67)	(0.10)	(2.85)	(1.34)	(-0.50)	(1.11)		
RF	0.245	-0.157	0.099	0.273	-0.149	0.100	5000	35.29
	(2.85)	(-1.86)	(1.41)	(1.78)	(-1.31)	(0.77)		
GB	0.134	-0.341	0.003	0.169	-0.336	0.007	10000	35.12
	(1.64)	(-3.85)	(0.05)	(1.38)	(-3.89)	(0.09)		
XTREE	0.245	-0.123	0.115	0.261	-0.125	0.105	5000	34.55
	(2.85)	(-1.17)	(1.48)	(1.56)	(-0.97)	(0.85)		
COMBO	0.305	-0.029	0.181	0.292	-0.065	0.150	5000	31.69
	(3.58)	(-0.31)	(2.67)	(1.69)	(-0.52)	(1.18)		

This table reports the average excess returns and CAPM 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 are after costs. Net costs are calculated using both a transaction volume of at least \$100,000 and the optimal values reported in the "Optimal Volume (thousand dollars)" column. 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 July 2006 through July 2020 (169 Months).

Table 13: Robustness: Smaller Number of Bonds

Signal	Excess Returns			CAPM α			Optimal	Turnover
	Gross	Net		Gross	Net			
		\$100K	Optimal		\$100K	Optimal	Volume	(%)
OLS	0.679	-0.077	0.201	0.566	-0.169	0.120	2000	46.15
	(3.97)	(-0.45)	(1.38)	(3.03)	(-1.04)	(0.71)		
OLSHUBER	0.789	-0.016	0.252	0.555	-0.225	0.091	5000	42.38
	(3.50)	(-0.08)	(1.80)	(2.77)	(-1.62)	(0.74)		
LASSO	0.694	0.012	0.235	0.545	-0.125	0.116	2000	47.53
	(3.95)	(0.06)	(1.52)	(3.29)	(-0.72)	(0.80)		
RIDGE	0.680	-0.076	0.201	0.566	-0.170	0.119	2000	46.13
	(3.94)	(-0.45)	(1.37)	(3.01)	(-1.04)	(0.70)		
ENET	0.566	-0.114	0.132	0.399	-0.268	-0.005	2000	46.47
	(2.90)	(-0.52)	(0.74)	(2.31)	(-1.39)	(-0.03)		
PCA	0.012	-0.485	-0.071	-0.197	-0.684	-0.147	20000	25.47
	(0.04)	(-1.82)	(-0.74)	(-0.84)	(-2.78)	(-1.76)		
IPCA	0.483	0.071	0.192	0.365	-0.027	0.087	1000	36.75
	(2.24)	(0.29)	(0.84)	(1.38)	(-0.10)	(0.35)		
PLS	0.463	-0.184	0.089	0.312	-0.306	-0.018	2000	44.48
	(2.35)	(-1.07)	(0.54)	(1.40)	(-1.75)	(-0.09)		
RF	0.657	-0.114	0.290	0.587	-0.165	0.227	2000	45.88
	(4.56)	(-0.88)	(2.40)	(3.91)	(-1.29)	(1.70)		
GB	0.889	-0.043	0.433	0.924	-0.008	0.459	2000	52.64
	(4.21)	(-0.27)	(3.02)	(3.81)	(-0.05)	(2.88)		
XTREE	0.579	-0.170	0.177	0.471	-0.278	0.060	2000	45.00
	(3.22)	(-0.84)	(1.11)	(2.60)	(-1.40)	(0.37)		
COMBO	0.654	-0.078	0.200	0.565	-0.150	0.128	1000	47.70
	(3.19)	(-0.42)	(1.15)	(2.35)	(-0.81)	(0.65)		

This table reports the average excess returns and CAPM 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 are after costs. Net costs are calculated using both a transaction volume of at least \$100,000 and the optimal values reported in the "Optimal Volume (thousand dollars)" column. 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 through November 2022 (244 Months).

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 ($\gamma = 0.05$, 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.

Appendix

A. 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,

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

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 A1: 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.

B. Variable definitions

Table A2: List of the stock and corporate bond characteristics.

Num.	ID	Characteristic name and description	Reference	Source
		Bond characteristics		
1.	age	Bond age. The number of years the bond has been in issuance.	Israel et al. (2018)	BAML/ICE
2.	coupon	Bond coupon. The annualised bond coupon payment in percent (%).	Chung et al. (2019)	BAML/ICE
3.	faceval	Face value. The bond amount outstanding in units.	Israel et al. (2018)	BAML/ICE
4.	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/COMPUSTAT
5.	debtebitda	Debt-to-EBITDA. Total debt (DLTTQ + DLCQ) divided by EBITDA (SALEQ − COGSQ − XSGAQ).	Kelly et al. (2021)	CRSP/COMPUSTAT
6.	duration	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
7.	ret61	Equity momentum. The sum of the last 6-months of equity returns minus the prior month.	Gebhardt et al. (2005b)	CRSP
8.	nime	Earnings-to-price. Net income (NIQ) divided by market equity.	Correia et al. (2012)	CRSP/COMPUSTAT
9.	me	Equity market capitalization.	Choi and Kim (2018)	CRSP
10.	eqtyvol	Equity volatility defined as the month-end value from a 180-day rolling-period.	Campbell and Taksler (2003)	CRSP
11.	totaldebt	Total firm debt (DLTTQ + DLCQ).	Kelly et al. (2021)	COMPUSTAT
12.	mom6	Corporate bond momentum. The sum of the last 6-months of bond returns minus the prior month.	Gebhardt et al. (2005b)	BAML/ICE
13.	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
14.	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
15.	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)	COMPUSTAT
16.	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/COMPUSTAT
17.	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/COMPUSTAT
18.	spread	Bond option adjusted credit spread. The option adjusted spread of the bond provided by ICE.	Kelly et al. (2021)	BAML/ICE
19.	operlvg	Operating leverage. Sales (SALEQ) minus EBITDA (SALEQ − COGSQ − XSGAQ), divided by EBITDA.	Gamba and Saretto (2013)	COMPUSTAT
20.	gpat	Profitability. Sales (REVTQ) minus cost-of-goods-sold (COGSQ), divided by assets (ATQ).	Choi and Kim (2018)	COMPUSTAT
21.	chggpat	Profitability change. The 5-year change in gross profitability.	Asness et al. (2019)	COMPUSTAT
22.	rating	Bond S&P rating. Bond numerical rating. AAA = 1, ... , D = 22.	Kelly et al. (2021)	BAML/ICE
23.	D2D	Distance-to-default. Computed as	Bharath and Shumway (2008)	CRSP/COMPUSTAT
24.	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)	CRSP/COMPUSTAT

25.	6msspread	Mom. 6m log(Spread). The log of the spread 6 months earlier minus current log spread.	Kelly et al. (2021)	CRSP/COMPUSTAT
26.	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/COMPUSTAT
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
28.	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
29.	vixbeta	VIX beta. Rolling 60-month regression of bond returns on the Fama French 3-factors ($Mkt\text{-}RF$,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

This table presents information on the 29 characteristics we use to form our predictions from the various machine learning (ML) models we employ. All bond related variables are computed using the BAML/ICE database. The equity characteristics use CRSP and COMPUSTAT.

Table A3: Net CAPM α

Turnover Rate (%)	Gross α (%) 0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
5	0.004	0.051	0.098	0.145	0.191	0.238	0.285	0.331	0.378	0.425	0.472
10	-0.004	0.042	0.089	0.135	0.181	0.227	0.274	0.320	0.366	0.412	0.459
15	-0.012	0.034	0.079	0.125	0.171	0.217	0.262	0.308	0.354	0.400	0.446
20	-0.021	0.025	0.070	0.115	0.161	0.206	0.251	0.297	0.342	0.387	0.432
25	-0.029	0.016	0.061	0.106	0.151	0.195	0.240	0.285	0.330	0.375	0.419
30	-0.037	0.007	0.052	0.096	0.140	0.185	0.229	0.273	0.318	0.362	0.406
35	-0.045	-0.002	0.042	0.086	0.130	0.174	0.218	0.262	0.306	0.349	0.393
40	-0.054	-0.010	0.033	0.077	0.120	0.163	0.207	0.250	0.293	0.337	0.380
45	-0.062	-0.019	0.024	0.067	0.110	0.153	0.196	0.238	0.281	0.324	0.367
50	-0.070	-0.028	0.015	0.057	0.100	0.142	0.184	0.227	0.269	0.312	0.354
55	-0.078	-0.037	0.006	0.047	0.089	0.131	0.173	0.215	0.257	0.299	0.341
60	-0.087	-0.045	-0.004	0.038	0.079	0.121	0.162	0.204	0.245	0.286	0.328
65	-0.095	-0.054	-0.013	0.028	0.069	0.110	0.151	0.192	0.233	0.274	0.315
70	-0.103	-0.063	-0.022	0.018	0.059	0.099	0.140	0.180	0.221	0.261	0.302
75	-0.111	-0.071	-0.031	0.009	0.049	0.089	0.129	0.169	0.209	0.249	0.289
80	-0.120	-0.080	-0.041	-0.001	0.038	0.078	0.117	0.157	0.197	0.236	0.276
85	-0.128	-0.089	-0.050	-0.011	0.028	0.067	0.106	0.145	0.184	0.223	0.263
90	-0.136	-0.098	-0.059	-0.021	0.018	0.057	0.095	0.134	0.172	0.211	0.249
95	-0.145	-0.106	-0.068	-0.030	0.008	0.046	0.084	0.122	0.160	0.198	0.236
100	-0.153	-0.115	-0.078	-0.040	-0.002	0.035	0.073	0.111	0.148	0.186	0.223

The table reports the net CAPM α as a function of gross α and portfolio turnover rate. The values are estimated by regressing the net CAPM α's on gross α, portfolio turnover rate, and the product of the two. The regression uses 41 strategies including 29 individual signals and 12 machine learning algorithms. 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.

WRDS Bond Database

The Wharton Research Data Services (WRDS) Bond Database is a pre-processed monthly bond database that uses the Enhanced Trade Reporting and Compliance Engine (TRACE) and Mergent FISD bond databases. It was introduced by WRDS in April 2017. The data is publicly available (requires a valid subscription to WRDS). After logging in to WRDS, the data is available here. We use the version of the WRDS dataset that has a sample end date of 2022:09.

WRDS bond returns. The WRDS data team provides us with three different bond return variables: `RET_EOM` (returns are computed using bond prices that land on any day of the month), `RET_L5M` (a bond must trade on the last five days of the month), and `RET_LDM` (a bond must trade on the last day of the month). For the results based on the WRDS Bond Database, we always use `RET_L5M`, i.e., a return is valid if the bond trades on the last five days of month t and month $t-1$. However, the publicly available data we use from WRDS, imposes a data filter which sets any bond return that is greater than 100% to 100%, i.e., the data is truncated/trimmed at this level. Although this does not make any material difference whatsoever to the main results, we carefully address the issue below.

WRDS bond returns truncation correction. We carefully adjust for the truncation of bonds with returns greater than +100% imposed by WRDS, by setting any bond return which is truncated to the return observed in the ICE database, i.e., if the WRDS bond return is equal to 100% (truncated), we set this value to the bond return from ICE as the 'true' bond return. If the ICE return is missing, we set the value to the return computed from the TRACE data itself. These adjustments do not make any material difference to the robustness results. In total we identify only 94 cases where the truncation occurs, and we are able to address 91 of them. The remaining 3 cases are removed.

WRDS bond filters. To align the data to the Bank of America Merrill Lynch (BAML) corporate bond database provided by the Intercontinental Exchange (ICE), we follow Andreani, Palhares, and Richardson (2023) and use the following filters (all using data provided by WRDS),

1. Remove investment (IG) rated bonds that have less than USD 150 million outstanding prior to, and including, November 2004, and less than USD 250 million after November 2004.

2. Remove non-investment grade (HY) rated bonds that have less than USD 100 million outstanding prior to, and including, September 2016, and less than USD 250 million after September 2016.

3. Remove bonds which are classified as zero-coupon, `bond_type == 'CMTZ'`.

4. Remove bonds which are classified as convertible, `conv == 'N'`.

We merge the WRDS data to the Mergent FISD database (also publicly available via the WRDS data platform) and apply the following filters already discussed above. This procedure delivers a transaction-based TRACE dataset that closely aligns to the quote-based ICE data.

A.2 Correcting price-based TRACE characteristics for microstructure noise

As first emphasized by Bartram, Grinblatt, and Nozawa (2021), BGN, price measurement error shared by a month-end transaction 'price-based' signal and the subsequent return generates correlation between the two. This affects, for example signals based on bond credit spreads, yields and size (market capitalization) for the TRACE (WRDS) database. We follow the methodology of BGN and Dickerson, Robotti, and Rossetti (2023b), DRR, and define the 'month-end' price-based signal to use a transaction-price at least one-business day before the price used to compute a month-end ex ante return. This methodology dampens the transmission mechanism of market microstructure noise (MMN) inherent in the price-based TRACE price signals. DRR show that by accounting for the transmission of the measurement error in this manner, the out-of-sample TRACE price-based anomalies are aligned to those observed when using the ICE dataset.

A.3 Kelly, Palhares, and Pruitt 2021 (KPP) data library

For robustness purposes, we utilize the publicly available KPP data library available for download on Seth Pruitt's personal website here. The KPP data is based on the WRDS TRACE dataset and contains bond excess returns and the 29 stock and bond characteristics required to train the ML models. For the price-based variables, we apply the market microstructure adjustments as described above.

B Machine learning model estimation and cross-validation

For all of our machine learning models, we re-estimate *and* cross-validate the model parameters every 6-months with an expanding window. Within each window we perform the cross-validation with a 85:15 training-validation split. For example, if we have window of 1,000 observations, 850 are used to train the model and the remaining 150 are used for validation. We graphically depict the sample splitting strategy for the training and cross-validation in Figure A.1.

After re-estimation, we hold the estimated model with the optimised parameters fixed for the following 6-months, and then re-estimate and cross-validate with an additional month of data. For all models we utilize the `sklearn` Python package (Pedregosa et al., 2011).

C 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

Figure A.1: Sample splitting strategy for cross-validation.

This figure shows the sample splitting used for cross-validation of the hyper-parameters of the penalized regressions, i.e., lasso, elastic net, ridge, and the regression tree ensembles for a given window. 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 85% of the observations while the latter consists of the final 15% 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.

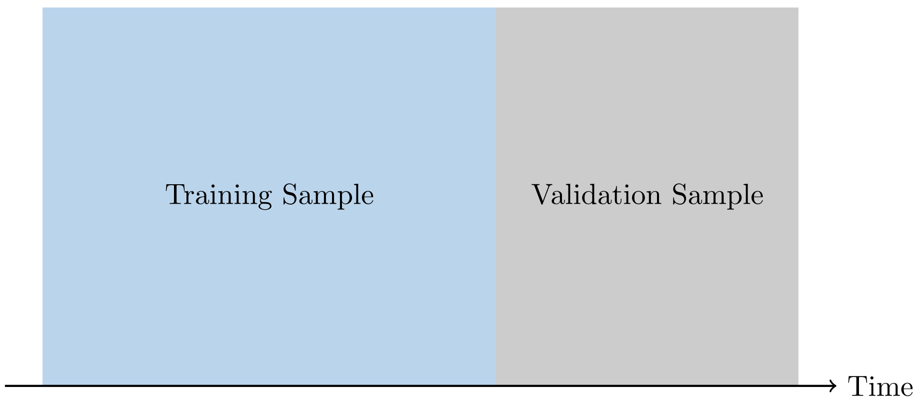

institutional holding data which provides the institutional ownership of corporate bonds at the quarter end from 1998Q2 to 2021Q2.

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.

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