Factor Investing with Delays*

Alexander Dickerson† Yoshio Nozawa‡ Cesare Robotti§

September 3, 2024

Abstract

We introduce a novel framework for computing transaction costs of trading strategies using infrequently traded assets. Infrequency leads to delays and missed trading opportunities. By modeling liquidity supply and demand and estimating supply elasticity through an exogenous demand shock, we calculate expected delays and bid-ask spreads that factor investors would experience. Applying this method to a comprehensive library of 341 corporate bond and machine-learning generated factors, we demonstrate that the majority underperform due to high turnover rates. Notably, even advanced machine-learning-based trading strategies yield zero or negative bond CAPM alphas after accounting for transaction costs. Our findings underscore the critical impact of delay costs in illiquid securities and provide valuable insights for factor investing in corporate bond markets.

JEL Classification: G12, G13

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

*We thank Pedro Barroso (discussant), Andrew Chen, Mike Chernov, Oleg Chuprinin, Christian Julliard, Edith Hotchkiss, Grace Hu (discussant), Alexey Ivashchenko (discussant), Jane Li, Hai Lin (discussant), Savina Risova, Konstantin Sokolov (discussant), Sam Wang, Alex Zhou, and the seminar participants at the ABFER, Atlanta Fed, ESSFM 2024 at Gerzensee, Temple University, the University of British Columbia, the University of Georgia, the University of New South Wales, the University of North Carolina Chapel Hill, the University of Technology Sydney, Monash University, the University of Toronto, the University of California (UCLA), the World Symposium for Investment Research, Bristol Financial Markets Conference, University of York Asset Pricing Workshop, and the Wellington Finance Summit for helpful comments. We also thank Julien Vincent for cloud computing support. A part of this research was conducted whilst Nozawa visited the Institute of Economic Research at Hitotsubashi University. The previous version of this paper was circulated under the title "The Low Frequency Trading Arms Race: Machines Versus Delays".

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

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

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

1 Introduction

Factor investing research has a long tradition dating back to Fama and French (1993). Typically, researchers evaluate proposed strategies using historical data, assuming that the representative investor is endowed with abundant trading opportunities and can trade instantaneously, without delay. This assumption may be justified in continuously traded, limit order book-driven equity markets, where hedge funds and other active investors trade frequently. However, this assumption is rendered questionable for illiquid assets, such as corporate bonds, where trading occurs in fragmented, over-the-counter (OTC) markets. These markets are characterized by relatively sparse trade frequency and the resulting trading delays are likely to significantly alter the performance of factor strategies.[1] To date, however, the literature is silent on how to quantify the cost of delays.

The trade delay can be inferred from the probability of a trade occurring in a given time period. Econometricians could simply estimate the delay by observing how often a given security is traded in the data. For example, if the estimated probability of trading per day is 5%, then the expected delay is about 1/0.05=20 days.[2]

However, the fact that a bond trades once every 20 days does not mean that active factor investors would have traded it only once every 20 days. One needs to understand why there were no trades on the remaining 19 days. The lack of trades may be due to a lack of liquidity supply from dealers, for example, due to their inventory capacity constraints, or a lack of liquidity demand from customers, who buy and hold. If a lack of demand is the reason why we observed trades only once every 20 days, then a factor investor, if they existed, could have traded the bond immediately after placing an order. That is, we cannot take the historical data - a single equilibrium outcome - as 'given' to estimate the probability of trades or the associated trade delay.

To overcome this challenge and quantify the delay costs, we use corporate bonds as an example of an illiquid asset with infrequent transactions. We use the transaction data provided by TRACE and infer the trading intensity, or the probability of a trade occurring in a given period, by forecasting transactions using a set of covariates. The fitted value of the forecasts is, in principle, the probability

[1] For example, Andreani et al. (2023) document that corporate bonds trade, on average, only six times in a given month. Many bonds exhibit zero trades within a month, and remain untraded for extended periods of time.

[2] If a delay is exponentially distributed with intensity λ, the mean delay is $1/\lambda$.

of transactions occurring. Under the assumption that transactions are independent over time, the estimated probability can be converted into transaction delays using an exponential distribution.

In order to disentangle the liquidity supply and demand effects, we model the supply and demand function for the probability of trading (or immediacy), which is treated as a service provided by dealers. Investors who purchase immediacy do so by paying higher bid-ask spreads. In the data, we observe an equilibrium in which demand is generated by real investors who are a mix of active (e.g., mutual funds) and inactive (e.g., insurance companies). By modelling the supply and demand functions, we can consider the counterfactual: assuming the supply curve is fixed, what happens if we introduce an active factor investor into the economy and shift the demand curve? The addition of active investors will change the equilibrium probability of trading and bid-ask spreads, providing the trading delay that the representative investor would face, if they existed in the sample. The explicit modelling of the liquidity supply and demand curves avoids simply taking the historical transactions-based data as 'given'.

We estimate the counterfactual in two steps. First, we ignore the slope of the supply curve and estimate the counterfactual reflecting only the positive shift in demand. This provides the lower bound on transaction costs because we ignore a potential increase in bid-ask spreads caused by the increase in demand. To do this, we estimate the demand function where the probability of trading is a linear function of bid-ask spreads and other bond characteristics. As a demand shifter, we add to the function Li and Yu (2023)'s measure of investor composition, $InvComp$, which is the bond-level average of the portfolio turnover rate of bondholders. Since bond ownership is persistent, this variable can be viewed as a Bartik instrument and allows us to reliably quantify the shift in liquidity demand.[3]

We estimate an OLS regression of trading frequency on $InvComp$ and other characteristics and confirm that bonds owned by more active investors are indeed more likely to be traded in the future. A one standard deviation increase in $InvComp$ increases the probability of trading by 1.7% to 3.0%, depending on the size and direction of the trade. We then shift the liquidity demand curve by increasing $InvComp$ to the 99th percentile of the data to capture the idea that factor investors

[3]See Goldsmith-Pinkham et al. (2020) for a detailed overview of Bartik instruments.

need to rebalance their portfolios regularly and trade bonds exactly at the end of the month, which requires high immediacy.

The shift results in a significant increase in the probability of trades. For example, in the data, transactions of \$100,000 and \$2 million occur 88% and 59% of the time within a month. In our counterfactual, they increase to 93% and 67%, respectively.

We can interpret these estimates as follows: For \$2 million trades, 41(=100-59) out of 100 days have no trades in the data. Of these 41 days, 8(=67-59) days are attributable to inactive buy-and-hold investors who do not need to trade. The remaining 33(=100-67) days represent the period when dealers do not provide liquidity. If there were active factor investors in the data, bonds would be traded on 67 out of 100 days.

In the counterfactual, the expected delays implied by the probability of trading shrink substantially. Using the given data, the expected delays are 10 and 24 days for \$100,000 and \$2 million transactions, respectively.[4] The corresponding values in our counterfactual fall to 8 and 19 days, closer to the values estimated in Kargar et al. (2023).

We then consider the shape of the liquidity supply curve. If the supply curve is upward sloping, an outward shift in the demand curve would shift the equilibrium along the supply curve, increasing bid-ask spreads and dampening the increase in the probability of trading. To compute the new equilibrium values of trade probability and bid-ask spreads, we estimate the slope of both the supply and demand functions. To remove the correlation between bid-ask spreads and unobserved drivers of bond trading, we instrument bid-ask spreads with $InvComp$ to estimate the supply faction and dealer capital ratio of He et al. (2017) and dealer bond inventory to estimate the demand function. We show that these instruments are significant drivers of bid-ask spreads: controlling for bond characteristics such as bond age, rating, maturity, and size, $InvComp$ and dealer inventory move bid-ask spreads positively and the dealer capital ratio moves them negatively.

In the new counterfactual, we estimate that bid-ask spreads increase by 11 to 12 basis points (bps) from the data, depending on the size of the trade. This is significant given that the corresponding average spreads in the data are 65 bps and 34 bps, for a \$100,000 or \$2 million trade

[4]This conversion is done using the exponential distribution. Specifically, the expected delay is $-\frac{21}{\log(1-Prob)}$

respectively. The widening of bid-ask spreads dampens the probability of trades demanded by the factor investor. However, in our estimates, the liquidity demand curve is very inelastic and the probability of trades in this counterfactual is similar to those that ignore the increase in bid-ask spreads.

Next, we study how the estimated delays influence factor investing. We calculate the cost of delays by jointly simulating delays and returns. With the simulated data, we construct portfolios and compute factor returns with and without delays. The difference in gross and net returns is the cost of delays, which is determined by the parameters of the data generation process. The delays are due to situations where an investor places an order but cannot execute the trade, resulting in a potentially missed investment opportunity. The signal that informs the investor's decision to buy the security may change before they can execute the order, and this creates an opportunity cost. Similarly, if the investor wants to sell the security, she may end up holding it unintentionally due to execution delays, further dragging down the performance of the portfolio with unintended inventory.

We show that the delay cost depends on three parameters of the data generation process. The first is the probability of trading, which controls the frequency of transactions and thus the magnitude of delays. For this, we use the equilibrium value in our counterfactual described above. The other two parameters are the factor turnover rate, which controls how quickly the signal changes, and gross factor performance, which determines the magnitude of lost profits due to delays. The cost of delay is high if the securities are traded infrequently, and the portfolio turnover rate and gross performance are high. Importantly, the cost of delay cannot be approximated by widening bid-ask spreads. Because it depends on gross performance, it penalizes the factor that performs well before costs more than the bid-ask cost does.

Armed with the estimated costs of delays and bid-ask spreads, we test the potential performance of corporate bond factor investing using 341 characteristics using the institutional grade Intercontinental Exchange (ICE) corporate bond pricing data that are free of market microstructure noise and potential data errors. We make our data library publicly available on `openbondassetpricing.com/data`. The data repository includes bond level data comprising the majority of all characteristics used to either predict bond returns or form factor portfolios and the time-series of the 341 bond factors.

To complement the publicly available data offering, and to aid replicability in the field, all of our results can be reproduced with the newly developed `PyBondLab` package in Python. This package is made available to the public and can be installed on any machine using pip (`pip install PyBondLab`) to replicate our results. The package includes functionality to compute bond portfolios rebalanced over any given horizon, as well as turnover which correctly accounts for endogenous changes in portfolio value. Furthermore, various tools are provided to dampen portfolio turnover including a buy-hold-spread (known as banding) and staggered holding periods. Detailed tutorials, examples and documentation can be found on the `PyBondLab` download webpage.

We compute gross performance and portfolio turnover rates using this package and assign the estimated delay costs to each of the 341 factors. The broad coverage of factors makes our study the first comprehensive study of bond factor performance *net of costs*.

We find that delay costs significantly affect the performance of the factors we study, even when we use the lower bound estimates for transaction costs. The adjustment for delay costs is particularly important for factors which rely on past returns as the signals. Among the 341 factors we study, the factor with the highest CAPM alpha is the factor based on the past one-month stock returns ($\alpha = 0.69\%$). However, this type of factor has the highest delay costs: since its gross performance is very high, it is painful for investors to miss the trading opportunity due to delays, resulting in severe opportunity costs of 0.36% when the trade size is \$100,000 and 0.44% when the trade size is \$2 million. These costs are exacerbated by the fact that the signal changes very quickly, forcing investors to trade frequently. Other reversal/momentum-type strategies, such as those based on three- or six-month changes in credit spreads or changes in option-implied volatility, face similar problems and generate negative CAPM alphas after transaction costs.

Examining the 341 factors, 56 factors have statistically significant bond CAPM alphas before transaction costs. However, once we account for delay costs and bid-ask spreads, none of the factors have significant alphas. In the most lenient case, where trade size is large and cost estimates are at the lower bound, only five factors have a CAPM alpha greater than 0.2% per month.

In addition to examining the factors individually, we also test whether a combination of factors can outperform the passive bond market index after transaction costs. For this purpose, we use nine machine learning models proposed in the literature and estimate them to generate the highest average returns before transaction costs. While all models generate economically large and significant gross CAPM alphas out of sample, none of them generate positive CAPM alphas after costs. These results suggest that accounting for transaction costs significantly changes the view of which factors are useful in practice.

In summary, this paper fills a gap in the literature by proposing a method to estimate the main impediment to factor investing in illiquid assets – delayed execution. In over-the-counter markets, delay costs due to search frictions are of first order importance. This is in contrast to the stock market with anonymous transactions through limit order books, where price impact is of paramount importance. In dealer-driven markets, large trades have low bid-ask spreads (e.g., see Edwards et al. 2007) but suffer from delays due to the strategic behavior of dealers to pre-arrange transactions (e.g., see Goldstein and Hotchkiss 2020; Choi et al. 2023 for evidence).[5] This observation highlights that considering only bid-ask spreads is not sufficient to evaluate the performance of factors in illiquid markets. Therefore, in this paper, we quantify delay costs and comprehensively evaluate the performance of corporate bond factors.

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; Augustin et al. 2020; Kelly et al. 2021; Sandulescu 2022; Binsbergen et al. 2023; Dickerson et al. 2023; Dick-Nielsen et al. 2023). The paper closest to ours is Ivashchenko and Kosowski (2023), who study the performance of nine factors after accounting for transaction costs. Our paper differs from Ivashchenko and Kosowski (2023) in that we highlight the novel delay costs faced by investors and employ a comprehensive factor data library in testing the performance of factor models spanning the majority of characteristics used to form factor strategies.

This paper also relates to the extensive literature measuring illiquidity and transaction costs in

[5]Goldstein and Hotchkiss (2020) present the empirical evidence of trade delays in the corporate bond market. 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 corporate bond market (e.g., Edwards et al. 2007; Chen et al. 2007; Feldhütter 2010; Bao et al. 2011; Schestag et al. 2016; Dick-Nielsen and Rossi 2018; Pinter et al. 2021; Choi et al. 2023; Kargar 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.[6] 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.

Our paper aims to provide the best practice in accounting for transaction costs in the study of the cross-section of corporate bond returns. Table 1 lists recent papers on this topic. The papers are classified into two groups: the first group of papers does not consider net returns after transaction costs and the second group does so. However, even among the papers in the second group, there is substantial heterogeneity in the transaction cost estimates. For example, Bali et al. (2020), Jostova et al. (2013), and Kelly et al. (2021) report significant trading profits arising from anomalies after accounting for transaction costs while Chordia et al. (2017), Bartram et al. (2023), and Nozawa et al. (2023) report anomalous returns largely disappear net of costs. However, none of the papers account for delay costs.[7]

The remainder of the paper is organized as follows: Section 2 describes our data set; Section 3 provides detailed methods for calculating delay costs using simulated data; Section 4 estimates the trade intensity and expected delays; Section 5 provides the evaluation of the factor performance; Section 6 extends the baseline results with more advanced factor investment strategies; and Section 7 provides concluding remarks.

[6]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).

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

2 Data

2.1 Data for Constructing Bond Factors

Our datasets include daily bond data from Enhanced TRACE (TRACE) and the constituent bonds from the Bank of America (BAML) Investment Grade and High Yield indices as made available via the Intercontinental Exchange (ICE). We source equity and accounting data from CRSP and COMPUSTAT. We filter the data using standard approaches as prescribed by the literature which is explicitly described in Internet Appendix A. To form bond factors and train the machine learning models, we construct commonly used bond and equity variables used in the literature and then merge these to the equity-based characteristics from Chen and Zimmermann (2022), CZ and Jensen, Kelly, and Pedersen (2023), JKP.[8] This data combines several monthly bond and stock-based characteristics that have been shown the predict one-month ahead future corporate bond excess returns. Our data includes 341 characteristics of which 40 are bond-based characteristics and 301 are equity-based, covering the majority of all corporate bond return predictors used in prior research. This data is publicly available for download on `openbondassetpricing.com/data`.

Detailed descriptions for the construction of our 53 custom-made bond and stock characteristics are provided in Table A.4 of the Internet Appendix. Extensive documentation for the CZ and JKP equity characteristics is available on the respective authors' websites. Missing characteristic data is set to its cross-sectional median at each month t. All characteristics are cross-sectionally rank demeaned to lie in the interval $[-1,1]$. Overall, the data used to train the ML models with the 341 stock and bond characteristics comprises 19,768 bonds issued by 2,110 firms over the sample period from January 1998 to December 2022 ($T = 288$).

2.2 Data for Trade Intensity Estimates

We use Enhanced TRACE to estimate trading intensity and half spreads. We filter the TRACE data by removing trades that are i) when issued, ii) with special conditions, iii) locked in, iv) have more than two days of settlement. We also remove observations with extreme price reversals: if

[8] Available on `openassetpricing.com` and via WRDS respectively.

the product of the two consecutive logarithmic price changes is less than -0.25 (e.g., a 50% increase followed by a -50% decrease), then the observations in the middle are dropped. To estimate trade intensity, we use only dealer-customer transactions. Finally, we restrict the bonds to those that appear in both ICE and TRACE. Finally, we use eMAXX for bond ownership information. The sample period is August 2002 through November 2022. Since this sample is shorter than the one used to train the machine learning models, all factors, whether single signal-based or ML-based, start in August 2002 and end in November 2022 for consistency. That is, the earlier sample is used only to generate signals, train and validate the ML models.

3 Simulated Cost of Trade Delays

3.1 Simulation Methods

In this section, we estimate the cost of delays using simulated data. This allows us to cleanly identify the cost without being influenced by the data structure, such as trade size and missing observations. We assume that on a given day, a transaction occurs with an intensity λ and delay τ follows an exponential distribution

$$\tau_{i,t} \sim \lambda e^{-\lambda\tau}, \tag{1}$$

with a mean delay of $1/\lambda$ days.

Now, we consider a stylized model of factor investing. We suppose that a bond i's gross return consists of predictable components and unpredictable shocks

$$r_{i,t+1} = \beta_{i,t} f_{t+1} + u_{i,t+1}, \tag{2}$$

$$\beta_{i,t+1} = \rho\beta_{i,t} + v_{i,t+1}, \tag{3}$$

where $u_{i,t+1}$ and $v_{i,t+1}$ are mean zero shocks and standard deviation σ_v and σ_u, independent of each other. The parameter ρ is common across bonds and over time and determines the speed of variation in the signal. Later, we show that ρ determines the portfolio turnover rate and $\sigma_v\sqrt{1-\rho^2}$ determines the profitability of the long-short portfolios. We have a systematic risk factor f_t in the dynamics whose risk premium determines the Sharpe ratio of the strategy.

We simulate the data using the system (1), (2) and (3) for $i = 1, \ldots, 1,000$ and $T = 244$ months. Every month, we sort bonds into quintiles based on observed $\beta_{i,t}$. Let $\mathbf{Q}_t$ be a set of bonds in the quintile in month t and $\mathbf{P}_t$ be a set of bonds in the portfolio in month t. Then, a net return on the bonds in the portfolio is

$$r_{i,t+1}^{Net} = \begin{cases} r_{i,t+1} & \text{if } i \in \mathbf{P}_{t-1} \text{ and } i \in \mathbf{Q}_t, \\ r_{i,t+1}((21 - \tau_{i,t+1})/21) & \text{if } i \notin \mathbf{P}_{t-1} \text{ and } i \in \mathbf{Q}_t \text{ and } \tau_{i,t} \leq 21, \\ 0 & \text{if } i \notin \mathbf{P}_{t-1} \text{ and } i \in \mathbf{Q}_t \text{ and } \tau_{i,t} > 21, \\ r_{i,t+1} & \text{if } i \in \mathbf{P}_{t-1} \text{ and } i \notin \mathbf{Q}_t \text{ and } \tau_{i,t} > 21. \\ r_{i,t+1}((21 - \tau_{i,t+1})/21) & \text{if } i \in \mathbf{P}_{t-1} \text{ and } i \notin \mathbf{Q}_t \text{ and } \tau_{i,t} \leq 21. \end{cases} \tag{4}$$

The first two lines represent the intended positions for the bonds. These bonds are included in the portfolio either because the investor already holds them or has successfully acquired them. In the case of a bond purchase, a delay of τ days reduces the return. The third line accounts for the scenario of trade failure in initiating a new position, resulting in a zero rate of return. The fourth and fifth lines pertain to bonds present in the portfolio from the previous month, which the investor retains due to unsuccessful sales attempts.

For each simulation, we compute the equal-weighted average of bond returns across the quintile portfolios. The high-minus-low factor for each simulation is then computed as the difference between the fifth ($\mathbf{P_5}$) and first ($\mathbf{P_1}$) portfolios. In this simulation, the composition of bonds in portfolio $\mathbf{P}$ deviates from what the signal suggests ($\mathbf{Q}$) due to changing β and transaction delays. This delay is costly because it leads to missed opportunities to trade.

3.2 Simulation Results

We simulate the data 1,000 times using the following sets of baseline parameters. In this section, we remain agnostic about the true parameters of the model, in particular, about trade intensity, λ. By studying a wide range of parameters, we can better understand the nature and quantity of the delay costs in factor investing. To this end, we vary $\rho = \{0, 3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.95, 0.99, 0.9995\}$

and $\lambda = \{1, 1/7, 1/21, 1/42, 1/252\}$. For each simulation run, we compute the average of gross and net long-short portfolio returns. We then take the average across simulations. The difference between the average gross and net returns is the cost of delay.

Table 2 reports the other parameters we select to match the sample size and volatility of the factor. In each run, we draw $\beta_{i,0}$ from a normal distribution with standard deviation $\sigma_v/\sqrt{1-\rho^2}$ using the parameters of the stationary distribution. In this illustration, we set the half spread to zero to focus on the cost of delay. Later, when we evaluate the net-of-cost performance of bond factors, we include a half spread and subtract it from a return each time a bond trades.

Figure 1 plots the cost of delay as a function of portfolio turnover rate (which negatively depends on ρ) and expected delay ($1/\lambda$ days). In this figure, we fix $\sigma_v = 0.2\%\sqrt{1-\rho^2}$ to roughly match gross average factor returns match the performance of the best 15 performing factors in the data (0.56% per month). As we can see, the cost of delay increases with both the portfolio turnover rate and decreases with the trading intensity, λ. Focusing on the case where a trade occurs on average once a month $1/\lambda = 21$, the figure shows that the cost of delay increases with the turnover rate. The cost of delay is about 0.15% per month when the turnover rate is 50% or less, but increases above 0.3% when the turnover rate is above 100%. This cost is substantial when compared to the gross return of 0.56% per month. For an extremely illiquid bond that trades once a year, the cost of delay increases more steeply with the turnover rate and tapers at the level of gross returns because the delay cost cannot exceed the gross returns. As expected, the bonds that trade on average every day ($\lambda = 1$), the cost of delay is smaller.

In Table 3 Panel A, we provide the cost of delays as a function of λ and turnover rates, providing specific values plotted in Figure 1.

Next, we examine how gross returns affect the cost of delay. To do so, we fix the trading intensity parameter $\lambda = 1/21$ (i.e., a bond is traded on average once a month) and vary the profitability of the strategy $\sigma_v = \{0.05\%, 0.10\%, 0.20\%, 0.30\%\} \times \sqrt{1-\rho^2}$. A higher value of σ_v results in a greater dispersion of expected returns across bonds and thus a higher gross return for the long-short strategy.

Figure 2 shows the cost of delay as a function of portfolio turnover and average gross portfolio returns holding trade intensity fixed at $\lambda = 1/21$. Table 3 provides the corresponding values that are plotted. The cost increases with both turnover and gross returns, suggesting that the interaction between the two contributes to transaction costs. In contrast, the standard measure of transaction costs, which is half spread times the turnover rate, is independent of the factor's gross profits. Our results show that the cost of delay is more important for the strategy that is highly profitable before transaction costs and has a high turnover rate.

The fact that the cost of delay depends on gross returns changes the preference within a menu of factors. As we show later, the momentum-like factors that have high gross returns also have a high turnover rate, while risk-based factors tend to have lower gross returns and turnover rates. The question, then, is which types of strategies perform better after transaction costs. To see this, we plot a contour set, or an indifference curve, of gross returns and turnover rate for a given net return. Figure 3 shows the set. The net return is increasing in the gross return but decreasing in the turnover rate, so the net return is increasing in the northwest direction. For example, any strategy above the black indifference curve has a net return of 0.2% per month or more. If a strategy is on the curve, it has a net return of 0.2% and any increase in the turnover rate must be accompanied by an increase in the gross return to maintain the 0.2% level. Since the curve is convex in turnover rate, a high turnover strategy requires increasingly higher gross returns to offset the increase in costs and remain equally profitable.

As a contrast to the cost of delay, we plot the traditional transaction cost of bid-ask spreads in the figure. These costs are a linear function of half spread h,

$$\text{Net Return} = \text{Gross Return} - h \times \text{Turnover}, \tag{5}$$

which implies a linear contour set for a given level of net return. As an example, in Figure 3, the black dashed line represents the contour set of 0.2% when the transaction costs are bid-ask spreads of 30 bps but with no delay costs.

Comparing the cost of delays (solid curve) and the cost of bid-ask spreads (dashed line), the contour set using the bid-ask spreads is above the set using the cost of delays for a low turnover rate.

This suggests that, for these factors, the net returns appear higher using the cost of delays than using the bid-ask spreads. However, for a high turnover rate, the pattern is reversed. The reversal occurs because the cost of delay increases with gross returns and thus it requires a greater increase in gross returns to offset a given increase in turnover rate to keep the net return constant. Thus, the cost of delay is qualitatively different from the cost of bid-ask spreads. If an investor considers the cost of delay, she will prefer a factor that has a low turnover rate and low gross returns. If, instead, she only considers the bid-ask spreads, she will prefer a factor with a high turnover rate and high gross returns. As we will show, the factors based on past returns tend to have high gross returns and high turnover rate. Therefore, including delay costs dramatically reduces the returns of these factors, and the reduction is much greater than when we only consider bid-ask spreads.

4 Estimating Trade Intensity

To quantify the delay cost, we estimate the trade intensity facing an active investor who wishes to rebalance portfolios every month. We consider monthly rebalancing given that it is standard in the factor investing literature but also extend the analysis by extending the holding period and introducing inertia in rebalancing in Section 6.1.

The probability of trades facing the active investor differs from the observed frequency of trades because the data is generated by real investors including insurance firms who tend to buy and hold. Thus, we consider a counterfactual scenario by introducing a hypothetical, small hedge fund pursuing factor investing strategies. In this counterfactual, both trade intensity and the compensation for the intensity, namely, half spreads will differ from the data. As we introduce a new investor, dealers who provide liquidity will respond to her demand, potentially changing immediacy and half spreads. Therefore, we need to specify liquidity supply and demand curves for this exercise.

In our setup, we view the probability of trades, $Prob_{i,t}[\mathbb{1}_{i,t+1} = 1]$, as a service provided by dealers. To compensate the dealers for the production of this service, customers have to pay half spreads, $h_{i,t}$. More formally, the supply and demand curves for bond i in month t are:

$$Prob_{i,t}[\mathbb{1}_{i,t+1} = 1] = \gamma h_{i,t} + \eta S_{i,t} + \phi X_{i,t}, \tag{6}$$

$$Prob_{i,t}[\mathbb{1}_{i,t+1} = 1] = \alpha h_{i,t} + \delta D_{i,t} + \zeta X_{i,t}. \tag{7}$$

The parameter γ is the sensitivity of trade intensity to half spreads. We expect $\gamma \geq 0$ because dealers require higher half spreads to supply immediacy. $S_{i,t}$ and $D_{i,t}$ are idiosyncratic supply and demand shifters, respectively, which are not captured by the bond characteristics in vector $X_{i,t}$.

We use Eqs. (6) and (7) to estimate the changes in trade intensity and half spreads in response to the shift in the demand curve caused by replacing passive investors in the data with hypothetical active investors. The key to this exercise are the parameters δ, which measures how much the demand shifts when the investor becomes active, and γ, which measures how dealers' liquidity provision reacts to a shift in the demand curve. In the following, we estimate these two parameters.

Estimating the probability of transactions is equivalent to estimating delays. Once we estimate the probability of trading for an active investor, we can use the exponential distribution and back out λ from the average probability of trade, given a trade size. Specifically, we set λ to satisfy the following equation,

$$\widehat{Prob}_{i,t}[\mathbb{1}_{i,t+1} = 1] = 1 - e^{-21\times\lambda}. \tag{8}$$

4.1 Incidence of Transactions and Half Spreads

In Eqs. (6) and (7), we define the trade indicator function $\mathbb{1}_{i,t}$ as a dummy if a trade with a size *above* threshold values (such as \$100,000 or \$2 million) occurs at least once a month.[9] This is motivated by the fact that larger transactions have lower half spreads, indicating that dealers are more willing to accommodate smaller trades if there is a demand.

To confirm the relationship between size and half spreads, we follow O'Hara and Zhou (2021) to compute the half spread for each trade. Let $\bar{P}$ denote the transaction price of the last interdealer trade that occurred before a customer-dealer trade. Then, the cost of the customer-dealer trade

[9] The dummy is defined at the month level rather than the day level. This is because at the day level, trades tend to cluster within a few day window, violating the underlying assumption for the exponential distribution.

for bond i is

$$h_{i,t} = (\log P_{i,t} - \log \bar{P}_{i,t}) \times I_{i,t} \tag{9}$$

where $I_{i,t}$ is 1 if the trade is customer buy and -1 if the trade is customer sell. We allow the reference trade to occur within the five-business day window preceding the trade. If there are no interdealer trades within this window, the half spread for the trade is treated as missing. We winsorize the half spreads at the 0.5 and 99.5 percentiles to reduce the influence of outliers.

Figure 4 plots average half spreads as a function of size. Consistent with Edwards et al. (2007), half spreads decline sharply with trade size. These are explained by the bargaining power of large customers (Duffie et al. 2005) and the strategic behavior of dealers to pre-arrange large trades (e.g., Goldstein and Hotchkiss 2020, Choi et al. 2023). In Section 6.3, we separate these two effects and confirm that both contribute to the declining pattern in halfspreads.[10] Therefore, from the client's perspective, when dealers are willing to trade a large size, it is easy to execute all trades smaller than that. Thus, we use the indicator variable that equals one when a minimum size trade is observed and zero otherwise.

4.2 Estimating Demand Shift Due to Changes in Investor Base

To quantify the shift in the demand curve, $D_{i,t}$, we use the investor composition measure of Li and Yu (2023), denoted *InvComp*. This variable measures the average activeness of the bond holders and is constructed as the portfolio turnover rate of investors averaged at the bond level.[11] The idea is that the frequency with which investors receive a liquidity shock is persistent, and thus

[10] Our results show that halfspreads decline with investor size after controlling for their identity. This contrasts with Pinter et al. (2021), who find that in the U.K. bond market, half spreads increase with trade size after controlling for investor identity. The difference arises because U.S. dealers behave differently from the rest of the world due to regulations such as the Volcker Rule and the TRACE disclosure requirement. The unique constraints imposed on U.S. dealers make them behave more strategically to control their balance sheets.

[11] Specifically, for investor j in quarter q, the net transaction is

$$nt_{j,q} = \frac{|\sum_i holding_{i,j,q} - \sum_i holding_{i,j,q-1}|}{\sum_i holding_{i,j,q-1}},$$

and we use the four-quarter average to capture the investor's preference,

$$NT_{j,q} = \frac{1}{4}\sum_{k=0}^{3} nt_{j,q-k}.$$

the investor whose portfolio turnover is high in the past will have a higher liquidity demand in the future.

To estimate a shift in (7), we run a forecasting regression of transaction in month $t+1$,

$$\mathbb{1}_{i,t+1} = \alpha_0 + \alpha h_{i,t} + \delta InvComp_{i,t} + \zeta X_{i,t} + e_{i,t+1}, \tag{10}$$

and compute a fitted value $\widehat{Prob}_{i,t}[\mathbb{1}_{i,t+1} = 1]|_{InvComp}$ for a given value of $InvComp$.

For the unbiased estimate for δ, we need

$$E[e_{i,t+1}|InvComp_{i,t}, h_{i,t}, X_{i,t}] = 0. \tag{11}$$

However, clearly, a liquidity price $h_{i,t}$ is an equilibrium object that is correlated with both $InvComp_{i,t}$ and forecast error $e_{i,t+1}$.[12] To eliminate the correlation, we use two proxies for liquidity supply shocks as instruments for $h_{i,t}$ including the dealer capital ratio of He et al. (2017), denoted $CAP_{i,t}$, and the 28-day inventory changes to capture the inventory pressure on dealers, denoted $\Delta Inventory_{i,t}$. Since we are interested in estimating δ but not α, we replace $h_{i,t}$ in Eq. (7) with $CAP_{i,t}, \Delta Inventory_{i,t}$ to estimate δ.

Table 4 Panel A reports the coefficient estimates, adjusted R^2, and the associated t-statistics. The right-hand-side variables are standardized for ease of interpretation. Standard errors are clustered at the bond level. We find that the measure of investor composition positively predicts the incidence of bond transactions, suggesting that investors with a history of frequent trading are

Finally, we aggregate at the bond level to obtain the investor composition measure,

$$InvComp_{i,q} = \frac{\sum_j holding_{i,j,q} \times NT_{j,q}}{\sum_j holding_{i,j,q}}.$$

We use $InvComp_{i,q}$ to predict transactions in each month in quarter $q+1$.

[12] If there are two right-hand-side variables, h and $InvComp$ that are mean zero, then we can show that OLS estimate for δ is

$$\begin{aligned}\hat{\delta} &= \delta + \frac{1}{|D|}(E[h^2]E[InvComp \cdot e] - E[h \cdot InvComp]E[h \cdot e]).\\ &= \delta + \frac{1}{|D|}E[z^2]E[InvComp^* \cdot e]\end{aligned}$$

Therefore, if $E[InvComp \cdot e] = 0$ and $E[h \cdot InvComp] = 0$ or $E[h \cdot e] = 0$, then $\hat{\delta} = \delta$. If $InvComp$ is correlated with h and h is correlated with the error term, then $\hat{\delta}$ is biased even if $E[InvComp \cdot e] = 0$.

more likely to demand liquidity. The dealer capital ratio predicts transactions positively, but only for larger trade sizes. The changes in the dealer inventory predict customer sells negatively and buys positively. The coefficients on the control variables also make sense: a bond with a lower credit rating (higher numerical value of the rating), shorter maturity, and larger size tends to trade more frequently.

Table 4 Panel B reports the regression-implied probability of trading. To obtain this value, we compute the fitted values of the regression for each observation and average them. The average probability of trading is higher for small transactions of \$100,000, 88% and 87% for customer sells and buys, respectively. This probability corresponds to trade intensity λ of 0.100 and 0.097 and the expected delay $(1/\lambda)$ of 10.0 and 10.4 days. For larger transactions of \$2 million, the probability drops to 59% and 58% for sells and buys. This corresponds to intensity $\lambda = 0.042, 0.041$ and an expected delay of 23.7 and 24.4 days, respectively.

We now consider the probability that an active investor would experience. In Figure 5, the demand curve shifts to the right, reflecting the fact that we introduce active investors in the market who requires greater liquidity than the actual investors in the data. We aim to quantify the magnitude of the shift by comparing the equilibrium in the data, E^{Data}, with the counterfactual, E^{LB}, as if the supply curve is a vertical line.

Since this investor does not buy and hold and requires liquidity at the end of each month, we set the liquidity demand variable, $InvComp$, at the 99th percentile. We then compute the fitted value of the regression (10), keeping the coefficients unchanged, and report the average "hypothetical" probability of trading at the bottom of Table 4. The hypothetical probability of sell, estimated at 95.3% and 70.9% for \$100,000 and \$2 million, respectively, is higher than the observed frequency of trade (87.7% and 58.8%), which reduces delays. Converting the monthly probability of trades into trade intensity using the exponential distribution, the expected delay for \$100,000 transactions shortens to 6.9 days ($\lambda = 0.145$) for the hypothetical investor. For trades with a size of \$2 million, the expected delay decreases from 23.7 days ($\lambda = 0.042$) to 17.0 days ($\lambda = 0.059$). Therefore, for a factor investor, the transaction delays are less than what we observe in the data.

The estimates above ignore dealers' reaction to the increased demand for immediacy. Thus, the estimated costs should be treated as the lower bound of transaction costs. While this procedure has the advantage of needing only a simple regression, it has the drawback of not capturing dealers' strategic behavior that provides greater immediacy in exchange for higher half spreads. That is, in practice, dealers' liquidity supply function is upward-sloping, and their behavior increases the half spreads and reduces the counterfactual probability of trades, both of which would increase transaction costs from the level documented above.

4.3 Estimating Supply and Demand Elasticity

Next, we consider how h and λ react if there is a positive shift in a demand curve and dealers are allowed to react by raising half spreads. Such a reaction is modelled as an equilibrium moving along the liquidity supply curve, shown in point E^{CF} in Figure 5. To quantify the change in half spreads and the probability of transactions, we need to estimate the supply and demand sensitivity parameters, γ and α. As before, an empirical challenge arises due to unobservable bond characteristics driving both $\lambda_{i,t}$ and $e_{i,t}$, implying $E[e_{i,t+1}|X_{i,t}, h_{i,t}] \neq 0$. To tackle this problem, we use the two-stage least square estimates. First, we regress half spreads on a set of instruments,

$$h_{i,t} = a + bZ_{i,t} + cX_{i,t} + u_{i,t}, \tag{12}$$

and obtain the fitted value of the regression, $\hat{h}_{i,t}$. As before, for the supply function, the instrument $Z_{i,t}$ is the investor composition measure, and for the demand function, it includes the dealer inventory measure and dealer capital ratio.

Next, we replace $h_{i,t}$ in Equation (10) with $\hat{h}_{i,t}$ to estimate the demand sensitivity parameter α. We repeat the same procedure for the supply function to estimate γ. Based on these parameters, we can compute the counterfactual changes in half spreads and the trade probability by

$$Prob^{CF} - Prob^{Data} = \frac{Prob^{LB} - Prob^{Data}}{1 - \alpha/\gamma}, \tag{13}$$

$$h^{CF} - h^{Data} = \frac{Prob^{CF} - Prob^{Data}}{\gamma}, \tag{14}$$

where $Prob^{LB}$ is the probability of transactions assuming no reactions by the dealers.

Table 5 reports the estimates of the first-stage regression in (12). We find that the intermediary capital ratio predicts half spreads negatively, while the 28-day dealer inventories and the log investor composition predicts them positively. Large t-statistics indicate that they are valid instruments that generate variation in half spreads after controlling for known determinants of bond illiquidity such as credit ratings, size, and age. In this regression, we pool all transactions regardless of trade size to maximize the number of observations. Still, the number is less than the entire sample used to predict trade incidence because if there is no trade in a month, the observation of half spreads is missing. Thus, in the second stage, we use a fitted value of half spreads, $\hat{h}_{i,t}$, for all observations including those with missing $h_{i,t}$. This is necessary because if we drop observations with missing $h_{i,t}$, then the probability of transactions will be severely upward biased.

Using the instrumented half spreads, now we estimate the second-stage regression and report the results in Table 6. Panel A reports the estimated slope coefficients, α and γ. Looking at the slope of the demand curve α for customer sell, the estimates are -0.008 and -0.075 for the trade size of $100,000 and $2M, respectively. The corresponding values for customer buys are 0.004 and -0.040, respectively. These estimates suggest that the liquidity demand are relatively insensitive to price changes. In the data, when a customer wants to trade, then she does so regardless of variations in half spreads charged by dealers.

Turning to the supply curve, the slope coefficients γ are estimated at 0.447 and 0.710 for $100,000 and $2M customer sells, respectively. The values for customer buys are highly similar. The positive slope implies that dealers are willing to increase liquidity provision for a given increase in half spreads. Thus, it is important to account for dealers' reaction to the increased customer activities. If we introduce hypothetical, active investor engaging in factor investing, then dealers will increase the half spreads from the level observed in the data.

Panel B of Table 6 reports the probability of transactions and half spreads in the new equilibrium. We find that the equilibrium level of the probability of transactions is highly similar to the counterfactual documented in Table 4. Because the demand curve is relatively flat, its parallel shift leads to the almost same increase in the equilibrium quantity. The new equilibrium will have higher half spreads than what is observed in the data as well. We estimate that, an increased liquidity demand leads to an increase in half spreads by 10 to 12 bps. Therefore, once we account for dealers' reaction function, the factor investor experiences not only a delay but also an increase in half spreads.

4.4 Delay Costs Based on the Counterfactual Trade Intensity

Now, we estimate the cost of delays based on the estimated trade intensity, λ =0.042, 0.058, 0.100,0.145. The first two values correspond to the estimates for the trade size of \$2 million, and the last two are those of \$100,000. As before, we vary the profitability of the strategy $\sigma_v = \{0.05\%, 0.10\%, 0.20\%, 0.30\%\} \times \sqrt{1-\rho^2}$ and plot the relationship between the cost of delay and portfolio turnover rate. The other parameters remain the same as in Figure 2. For this exercise, we also add to the plot the cost arising from half spreads defined in (9).

Figure 6 shows the cost of delays and bid-ask spreads. In Panels (a) and (b), we plot the cost when trade size is \$100,000. In this case, the bid-ask spread cost dominates the cost of delay because delays are short and the spread is large for small transactions. In Panels (c) and (d), we plot the cost for a trade size of \$2 million. As trade intensity falls, the cost of delays becomes substantial. In both Panels (c) and (d), the delay cost exceeds the bid-ask spread costs when the gross average return of the strategy is set to replicate the performance of the top 15 bond factors (0.56% per month).

Figure 7 shows the net returns of the simulated strategies. Panels (a) and (b) correspond to the cases with a trade size of \$100,000. In this case, net returns fall almost linearly with the turnover rate because the bid-ask spread cost linearly increases with turnover. Panels (c) and (d) correspond to the case with a trade size of \$2 million. In this case, the net return decreases more steeply with the turnover rate when its gross return is higher (e.g. 0.84%) than when the gross return is lower (e.g. 0.14%). This reflects the fact that the delay cost increases with the gross return.

We can infer the effect of estimation bias in λ by comparing Panels (a)/(c) with (b)/(d) in Figures 6 and 7. For more direct analysis, in Figure 8, we plot the cost of delay before and after accounting for the bias corresponding to the lack of liquidity demand. For this plot, we fix $\sigma_v = 0.20\%$. We compute the cost of delays using the average λ^{Data} in the data and the counterfactual λ^{LB} by setting $InvComp$ to the 99th percentile.

In Figure 8, we find that accounting for the liquidity demand effect reduces the cost of delay. From the perspective of active investors, the bond market is more active than the average frequency observed in the data, and thus transactions are more likely to occur, which is reflected in a higher value of λ.

4.5 Validations of the Estimates

Estimates by Credit Rating

In the results above, we find that lower credit quality (i.e., higher value of credit rating) is associated with a higher probability of transactions. To assess a potential bias driving this result, we repeat the analysis using the subsample of IG bonds and HY bonds. We focus on customer sells because the results for customer buys are highly similar. Table 7 reports the estimates by credit rating. In the data, HY bonds are traded more frequently than IG bonds. Thus, the expected delay in the data is 10.6 days for $100,000 trades for IG bonds and 8.6 days for HY bonds. However, this leads to the counter-intuitive conclusion that HY bonds are easier to trade with shorter delays, reminding us that we cannot take the data as given because some investors optimally choose not to trade bonds. In the data, IG bonds are more likely to be held by buy-and-hold insurance companies for regulatory reasons. This pattern gives the impression that IG bonds are illiquid.

The counterfactual trading intensity obtained from the regression accounts for this bias. If an investor tries to trade actively then the expected delay for IG bonds would be 7.1 and 16.3 days with a trade size of $100,000 and $2M, respectively. These values are very similar to the corresponding values for HY bonds, which are 6.2 and 15.6 days, respectively. Therefore, the apparent ‘liquidity’ of HY bonds is fully explained by the different investor compositions in the data and adjusting for it helps avoid the conclusion that HY bonds are easier to trade than IG bonds.

Time-Series of Delays

For further confirmation, Figure 9 plots the estimated delay based on the regression in Table 4 averaged over bonds by month. There is a break in the data in January 2005, which corresponds to a change in the ICE index criterion that removes small bonds from the sample. Since the amount outstanding is one of the most important predictors of a trade, an increase in the amount outstanding significantly reduces the delay. Thereafter, delays slowly decrease, reaching 8.9 days and 18.5 days for $100,000 and $2 million trades, respectively, by the end of the sample.

The decreasing trend may reflect the introduction of electronic trading platforms as studied in O'Hara and Zhou (2021). However, Kargar et al. (2023) shows that the volume share of the electronic platform is still small, at 16.4% for the period January 2017 to March 2021.

Compared to the results of Kargar et al. (2023), our estimate of delays for $100,000 transactions is comparable: In their Figure 6, the delay for voice odd-lot trades by the most connected customer is about 10 to 12 days,[13] while that for an electronic platform is 4 days. Our estimate of around 7-day delay for $100,000 trades lies between these two estimates. Our estimate of the delay for a $2 million trade is longer than that of Kargar et al. (2023). In their Figure 6, the delay for voice round lot trades is about 10 to 12 days while our estimates in Table 4 is around 19 days. The difference likely reflects the difference in the sample – their voice trade is still matched to orders first placed on the electronic platform, which represents 18.1% of the total volume in TRACE. Since our sample includes all bonds in TRACE, our delay estimate for large trades is longer.

Robustness

Our results are not sensitive to the specification of the investor liquidity demand variable. In Internet Appendix Tables A.1, A.2 and A.3, we conduct robustness tests by running logit regressions instead of linear regressions, allowing for the nonlinear dependence of the probability of trading on investor composition, and introducing the interaction between investor composition and bond characteristics. The resulting counterfactual trade intensities are not significantly different from

[13] They show delays as a function of the number of failed inquiries, and we have to integrate over them to get the average delays. Since we do not know the distribution of requests, we only know the range of delays.

those in Table 4.

5 Performance of the Corporate Bond Factors

In this section, we comprehensively examine the performance of corporate bond factors after transaction costs.

5.1 Analytical Framework

To generate the 341 corporate bond factors, we compute monthly bond total returns in the standard manner: $R_{i,t+1} = \frac{P_{i,t+1}+AI_{i,t+1}+C_{i,t+1}}{P_{i,t}+AI_{i,t}} - 1$, where $P_{i,t+1}$ is the clean price of bond i in month $t+1$, $AI_{i,t+1}$ is the accrued interest, and $C_{i,t+1}$ is the coupon payment, if any.

In the main results, the factors are assumed to be rebalanced monthly (at the end of the month), based on characteristic-based signals observed to the factor investor at month-end. Bond excess returns, r_{t+1}, are defined as the bond total return minus the one-month risk-free rate of return from Kenneth French's webpage. Following the literature (Novy-Marx and Velikov 2015, Chen and Velikov 2023) who evaluate hundreds of equity factors/anomalies, we form the high-minus-low factor strategy using deciles with (bond) market capitalization as the weights. The long (short) position is assumed to be decile ten (one) such that all factors are sign-corrected to be increasing in expected returns.

We compute monthly turnover as follows,

$$\begin{aligned} TO_{p,t+1} &= \sum_{i \in N_t} \left| w_{i,p,t+1} - \frac{1+R_{i,t+1}}{1+R_{p,t+1}} w_{i,p,t} \right|, \\ TO_{f,t+1} &= \frac{1}{2} \sum_{p \in 1,10} TO_{p,t+1} \end{aligned} \tag{15}$$

where $TO_{f,t+1}$ is the monthly turnover for factor f. Equation (15) follows Novy-Marx and Velikov (2015) and accounts for endogenous changes in portfolio weights month-to-month due to changes in bond value.

Risk-Adjustments To compute factor alphas, we use both the gross bond market factor and the net bond market factor computed with real net return data. We estimate the net bond CAPM (CAPMB) alpha by running time series regressions of the factor excess returns on the corporate bond market factor:

$$r_t^f = \alpha + \beta MKTB_{Net,t} + \varepsilon_t, \tag{16}$$

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

Transaction Cost Mitigation We implement a suite of rule-based cost mitigation techniques including fewer portfolio break-points (quintiles), staggered rebalancing and buy-hold spreads (banding) following Novy-Marx and Velikov (2019). These methods reduce turnover at the cost of potentially reduced expected returns and are discussed in detail later on.

Open Source Bond Asset Pricing To aid replicability, all out-of-sample factor returns and associated turnovers are computed with the `PyBondLab` Python software package. The software is also used to implement the transaction-cost mitigation methodologies. Detailed user guides and additional information can be located here. The software can be used in conjunction with the publicly available factor characteristic data and bond returns that we make available on `openbondassetpricing.com/data`. In addition, the time-series of our factor returns (341 of them) and monthly turnovers that are used in the main results are available for download here.

5.2 Empirical Results: Net of Cost Performance

We examine the performance of the 341 corporate bond factors using simulated transaction costs. The delay and bid-ask spread costs are assigned to each factor based on its turnover rate and gross returns. Based on the counterfactual trading intensity for a given trade size, we construct a grid of transaction costs in the same manner as in Table 3 Panel B. We then linearly interpolate the grid to find the corresponding value for each factor. We take the gross returns and α in the data as given (i.e., we do not simulate them) and subtract the transaction costs obtained from the simulation.

In Table 8, we report the distribution of the CAPM alphas across the 341 bond factors. Panel A reports the distribution before transaction costs. The average and median gross CAPM alpha are 9 bps and 6 bps per month. The number of factors that earn non-trivial alphas in the range of $[0.1\%, 0.2\%)$ and $[0.2\%, \infty)$ is 69 and 50, respectively. The number of factors with statistically significant alphas is 56. In a nutshell, using `PyBondLab`, we successfully identify a significant number of factors that reliably outperform the bond market.

In Panel B, we report the net-of-cost α using the lower bound estimates for transaction costs with trade size of \$100,000 and \$2 million. With transaction costs, the number of factors with non-trivial alphas decline dramatically. The number of factors in buckets $[0.1\%, 0.2\%)$ and $[0.2\%, \infty)$ is 3 and 0 with a trade size of \$100,000 and 6 and 5 with a trade size of \$2 million. None of the factors earn significantly positive alphas.

To understand the dramatic changes in CAPM α, in Table 9, we report the gross and net CAPM α of the top-performing factors for trade sizes \$100,000 and \$2 million. Specifically, we sort them by gross CAPM alpha and select the top 15 characteristics. We find that the profitable factors before costs fall into two groups. First, there are factors based on past returns, including past 1-month stock returns (ret10), past three- and six-month changes in credit spreads (mom3mspread and mom6mspread), and transformed stock returns (trend factor, Han et al. 2016). These factors tend to have high gross alpha and turnover rates because the signal changes quickly.

Second, there are credit spread-related characteristics. Unlike return-based measures, these factors tend to have lower turnover rates because the price of a bond does not change dramatically. For example, bond price, a close relative to the bond book-to-market ratio of Bartram et al. (2023) has the second best gross CAPM alpha of 0.63% per month with the turnover rate of 30%.

Examining the net α for a trade size of \$100,000, the net α in Table 9 is generally low: 10 out of the 15 characteristics have negative alphas, and none of them earns a positively significant alpha. This is mostly due to the bid-ask spread costs. When the trade size is small, the cost of delay tends to be small but the bid-ask spreads are quite high and dominate the gross alpha. Net alpha is particularly low for signals that are based on past returns due to their high turnover rate.

When trade size is increased to \$2 million, the bid-ask spread shrinks dramatically. On the other hand, the delay cost increases and partly offsets the benefit of the lower spread. For example, the strategy based on past 1-month stock returns has the delay cost of 0.44% per month and the bid-ask spread costs of 0.65%.

Figure 10 visualizes the relationship between turnover rate, gross CAPM alphas, and net alphas. In the plot, each point represents a factor plotted against its gross return and turnover rate. As before, the factors above the curve are those that achieve the given level of net return after delay and bid-ask spread costs. Using the trade size of \$100,000 (Panel A), most of the factors are below the 10 bps threshold, meaning that they earn less than 0.1% per month after costs. Since the curve is steeply upward sloping, there is almost no chance for a factor with a turnover rate of 50% or more to generate meaningful net returns.

Figure 10 Panel B shows the performance of the factors when the trade size is \$2 million. Thanks to lower bid-ask spread costs, the curve is flatter than in Panel A, and the number of factors with meaningful net returns increases. Now, eight factors have net returns above 0.1%, and several factors have net returns above 0.2%. Still, the vast majority of factors have turnover rates way too high to be profitable after transaction costs.

Since the net-of-cost factor performance is quite weak even when the cost is estimated using the lower bound, their performance is even worse if we use the counterfactual equilibrium where dealers require higher bid-ask spreads in response to the increased liquidity demand. Panel C of

Table 8 reports the distribution of the CAPM alphas when these counterfactual estimates are used. The number of factors in buckets [0.1%,0.2%) and [0.2%,∞) is 1 and 0 with a trade size of \$100,000 and 8 and 1 with a trade size of \$2 million. Again, none of the factors earn significantly positive alphas.

6 Extensions of Our Framework

6.1 Different Investment Horizons and Banding

The factors outlined in prior sections, which utilize the high-minus-low decile sort prevalent in the academic literature, may overestimate the true trading costs associated with these factors. In reality, investment managers could employ a variety of techniques to reduce turnover whilst attempting to minimize the loss in expected returns. To address this, we use several rules-based cost mitigation techniques including staggered rebalancing where only a fraction of the position is traded each month (Jegadeesh and Titman, 1993) and "banding" (also known as a buy/hold spread) introduced by Novy-Marx and Velikov (2015). The buy/hold spread technique implements thresholds for portfolio rebalancing, imposing a more stringent criterion for initiating new positions compared to maintaining existing ones. By introducing this inertia into the trading process, the method effectively creates a "band" around current positions, potentially preserving strategy performance while significantly reducing trading frequency. For both the staggered rebalancing and banding turnover reduction method, we utilize the `PyBondLab` open source bond portfolio software.

Staggered partial rebalancing This technique rebalances a fraction of the portfolio holdings each month and was first used by Jegadeesh and Titman (1993). In any month $t+1$, the factor return is the equally weighted average month $t+1$ return of the factor strategy implemented in the prior month and up to H months earlier, where H is the staggered holding period. For example, if $H = 3$, this mechanically limits the one-sided monthly portfolio turnover to be $\leq 33.33\%$.

Table 10 shows the distribution of CAPM alphas when the holding period is extended to three months. As expected, the slower rebalancing reduces the efficiency of the factor performance before transaction costs. Comparing the results with a one-month holding period in Table 8, the average gross alpha decreases from 9 bps to 8 bps and the number of significant factors decreases from 56 to 36. However, net performance improves slightly due to a sharp reduction in the turnover rate. For example, the average net alpha for trade sizes of $100,000 and $2 million with the lower bound cost estimates increases from -41 bps and -14 bps to -22 bps and -6 bps, respectively. Thus, it is in fact more profitable for bond investors to rebalance portfolios infrequently.

In Table 11, we list the top 15 factors with the highest gross alphas using the lower-bound cost estimates. The table confirms that the turnover rates of these factors are much lower than Table 9, which reduces both delay and bid-ask spread costs.

Banding Novy-Marx and Velikov (2015) and Novy-Marx and Velikov (2019) demonstrate that banding is the most successful method for reducing portfolio turnover in the equity market. We follow their suggestion and create a band with a width of two around the first and tenth portfolios. This means that, once a bond is included in the portfolios one or ten, then it stays in the same portfolio as long as the bond's signal remains in the first three or last three deciles. The idea is that as long as the signal is "strong enough" then an investor keeps the position even after observing small changes in the signal. This is another way of reducing portfolio turnover rates.

Table 12 reports the distribution of the net CAPM alphas with banding. Unlike Novy-Marx and Velikov (2015), banding does not improve the average net CAPM alphas relative to the strategy with longer holding periods. The average net CAPM alphas are -24 bps and -7 bps, which are slightly below the values reported in Table 10. However, banding improves the right tail of the distribution: now, the number of factors with net alphas more than 0.1% per month increases to 12 (trade size of $100,000) and 23 (trade size of $2 million) from the corresponding values of 5 and 16 in Table 10, respectively. One of the 341 factors has statistically significant CAPM alpha with banding. For completeness, Table 13 reports the top 15 factors with banding.

In summary, both methods of reducing portfolio turnover are effective in improving net CAPM alphas. This suggests that in future bond research it is important to study the performance of factors with a holding period longer than one month.

6.2 Combining Characteristics

In the previous section, we have shown that taken individually, the transaction costs including delays dwarf the performance of bond factors in most cases. However, this does not mean that the optimal combination of factors is dominated by transaction costs. Machine learning models are known to provide an excellent ex-ante combination of factors without inducing look-ahead bias. Thus, we revisit the standard machine learning models employed in the literature (e.g. Bali et al. 2023) and study the effect of transaction costs on these models.

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

For the first estimation as of July 2002, we source the last 55 months of data back to January 1998, and estimate the respective ML model. We measure excess returns at t and the 341-dimensional vector of bond characteristics at $t-1$. We perform cross-validation using a 70:30 training-validation split which preserves the temporal ordering of the panel data. We then use the vector of characteristics available at time t to produce a forecast of bond excess returns for $t+1$. These forecasts (expected returns) are available to the bond portfolio manager at time t, meaning they can trade on them at the end of the month. Thereafter, all models are re-trained every 12-months and cross-validated every five-years with an expanding window.[14] We provide additional details related to the cross-validation and training of the respective models in Section C of the Internet Appendix.

Performance of Machine-Learning Models After Costs

We turn to the performance of the machine learning models relative to the simulated transaction costs. We take the gross returns and alpha in the data as given (i.e., we do not simulate them) and subtract the transaction costs obtained from the simulation.

In Table 14, we report the gross and net CAPM α of the ML strategies. When we use the lower bound cost with a trade size of $2 million, the net alpha, or the difference between the gross alpha and the simulated total cost, ranges from -0.15% to -0.52%, which is much lower than the gross alpha. The net alphas are much lower than gross alphas because the models are estimated to maximize average gross returns without taking into account portfolio turnover rates. As a result, the models place a large weight on rapidly changing factors such as the short-term reversal factor. Therefore, when estimating machine learning models with illiquid assets, it is essential to penalize high turnover rates. In this regard, the methodology of Bredendiek et al. (2019) is a fruitful attempt to address the issue of transaction costs.

6.3 Half Spreads By Investor Size in the U.S. Corporate Bond Markets

In this section, we revisit the relationship between bid-ask spreads and trade size and confirm that bid-ask spreads decrease with trade size. This result implies that bid-ask spreads are insufficient measure of transaction costs.

To this end, we examine insurance firms' transactions in eMAXX. We add to TRACE the information about the insurance firm's size for each trade. Using eMAXX's firm id, we aggregate all corporate bond holdings in quarter q and merge the data to all transactions by the firm in quarter $q+1$.

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

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

Table 15 reports the mean and median half spreads in each category. We find that the half spreads decline in both investor and trade sizes. Controlling for trade size, the average and median spread declines as we move from the small investor category to the large investor category. This finding is consistent with O'Hara et al. (2018) and confirms the prediction of Duffie et al. (2005). However, controlling for investor size, the spreads also decline from small to large trades.

The number of observations shows that investor size and trade size are positively correlated, as expected. However, there are plenty of observations in which large investors trade with a small size. For example, there are 17,489 observations where the top 20% insurance firms trade with a size less than \$100,000. The smallest insurance firms still trade with a size of more than \$1 million: there are 55, 636 of such instances in our data.

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

$$\mathrm{c}_{i,k,t} = \beta_v \times \log(\text{Volume})_{i,t} + \tau_t + \alpha_i + \lambda_b + \mathbf{Y}'_{k,t}\delta_C + \varepsilon_v, \tag{17}$$

where $\mathrm{c}_{i,k,t}$ is the trading cost as computed in Equation (9) for client i trading bond k on day t, $\log(\text{Volume})_{i,k,t}$ is the natural logarithm of the given trade's notional, τ_t are trade-day fixed effects, α_i (λ_b) are client (broker) fixed effects respectively. The vector $\mathbf{Y}_{k,t}$ captures bond-level characteristics that we use as controls including bond rating, coupon, maturity and the log of the issue size (amount outstanding). The key coefficient of interest is the estimated value of β_v: if the half-spread continues to decrease after controlling for the client base, we would expect β_v to remain negative and statistically different from zero.

In columns (1) and (4) of Table 16, we confirm the results from prior research documenting a negative relationship between half-spreads and trade size (volume) without and with bond-level controls. In columns (2) and (5), after including client-level fixed effects, the sign and magnitude of the β_v coefficient remains unchanged, reinforcing our results from the double sort above. It would appear that the 'Size Penalty' does not exist in the U.S. corporate bond market. As an additional test, we also include broker fixed effects in columns (3) and (6). β_v remains robustly negative.

The economic significance of the size premium (captured by β_v) is sizeable. A one standard deviation increase in the log of trading volume decreases half-spreads by about 0.104% in column (1).[15] Relative to the average half-spread of 0.179%, this change is economically large. Once including client fixed effects in column (2), the effect remains sizeable and results in a 0.07% decrease in half-spreads.

The evidence from the U.S. insurance firms' transactions shows that both investors' bargaining power and dealers' strategic behavior are at work in explaining the observed transaction costs. The second channel suggests that for an investor, there is a hidden cost of trading with a large size that is missed by the observed half spread. Observed spreads for large trades are low due to the dealers' strategic behavior to pre-arrange trades for risky transactions, creating delays for investors. In sum, these results suggest that it is essential to include the delay costs to accurately capture the transaction costs borne by bond investors, especially when trade size is large.

7 Conclusion

In this article, we propose a novel method to quantify the cost of delays that factor investors would face in an illiquid, OTC-based market. The key is to disentangle the drivers of trade occurrence and estimate the probability of trades with a counterfactual in which the liquidity demand curve is shifted outward by the introduction of active factor investors. We propose instrumental variable estimates for the supply and demand curve to pin down the new equilibrium.

The estimated probability of trades are used as inputs to the simulation, which generates the cost of delays. The cost is a function of both portfolio turnover rates and factor gross performance. Unlike the transaction costs arising from bid-ask spreads, the delay cost is more severe for factors with higher gross performance due to increased opportunity costs.

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

We create a publicly available data library of corporate bond factors, spanning 341 characteristics. The data library comprises both bond-level panel data, including all underlying characteristics to aid replicability in the field, and the time-series of our 341 out-of-sample factor returns. In addition, we introduce the `PyBondLab` package, which enables researchers to generate factors with various conditions in terms of the number of factors, holding periods, transaction cost mitigation techniques, and portfolio weights. We use this package to generate factors and assess their gross and net performance.

Empirically, the transaction costs are large. They are more pronounced for momentum-type factors in which the signal is based on past asset returns. Since these signals move rapidly and are highly profitable before costs, the delay cost tends to be high, eliminating net bond CAPM alphas.

We test several cost mitigation techniques proposed in the literature and confirm that they work well in reducing costs. Still, it is a challenge to find corporate bond factors that reliably outperform the market after transaction costs.

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.

Augustin, Patrick, Linxiao F. Cong, Ricardo L. Aliouchkin, and Roméo Tédongap, 2020, Downside risk and the cross-section of corporate bond returns, Working Paper.

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

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

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

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

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

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

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

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

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

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

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

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

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

Bredendiek, Maximilian, Giorgio Ottonello, and Rossen Valkanov, 2019, Corporate bond portfolios and asset-specific information, Working Paper, Vienna Graduate School of Finance.

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

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

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

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

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

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

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

Choi, Jaewon, 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.

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

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

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

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

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

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

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

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

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

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

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

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

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

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

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

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

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, 2020, Bartik instruments: What, when, why, and how, *American Economic Review* 110, 2586–2624.

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

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

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

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

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

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.

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

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

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Jensen, Theis Ingerslev, Bryan Kelly, and Lasse Heje Pedersen, 2023, Is there a replication crisis in finance?, *Journal of Finance* 78, 2465–2518.

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

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

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

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

Li, Jian, and Haiyue Yu, 2023, Investor composition and the liquidity component in the u.s. corporate bond market, *Journal of Finance* forthcoming.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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 1: Cost of Delay (Percent Per Month)

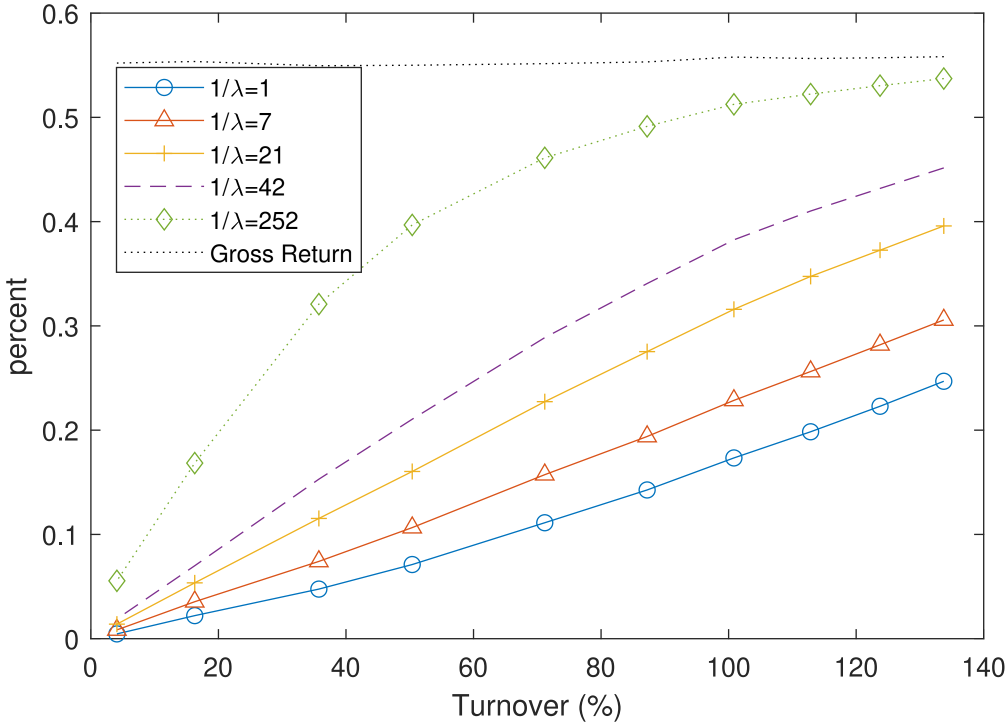

We plot the difference in average returns between the gross factors that do not have delays and the net factors that are subject to delays. Delays follow an exponential distribution with parameter λ, which measures the trade intensity on a day. We vary the signal persistence parameter ρ to change the portfolio turnover rate.

Figure 2: Cost of Delay (Percent Per Month)

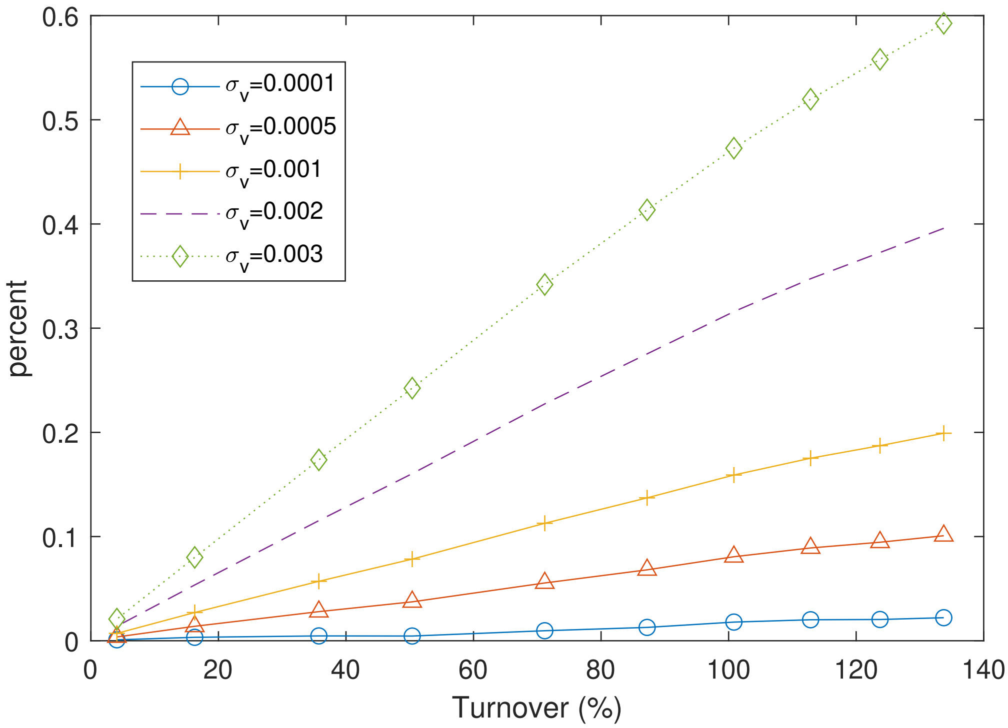

We plot the difference in average returns between the gross factors that do not have delays and the net factors that are subject to delays. Delays follow an exponential distribution with parameter λ, which measures the trade intensity on a given day. We vary the signal persistence parameter ρ to change the portfolio turnover rate.

Figure 3: Indifference Curve For Different Net Returns

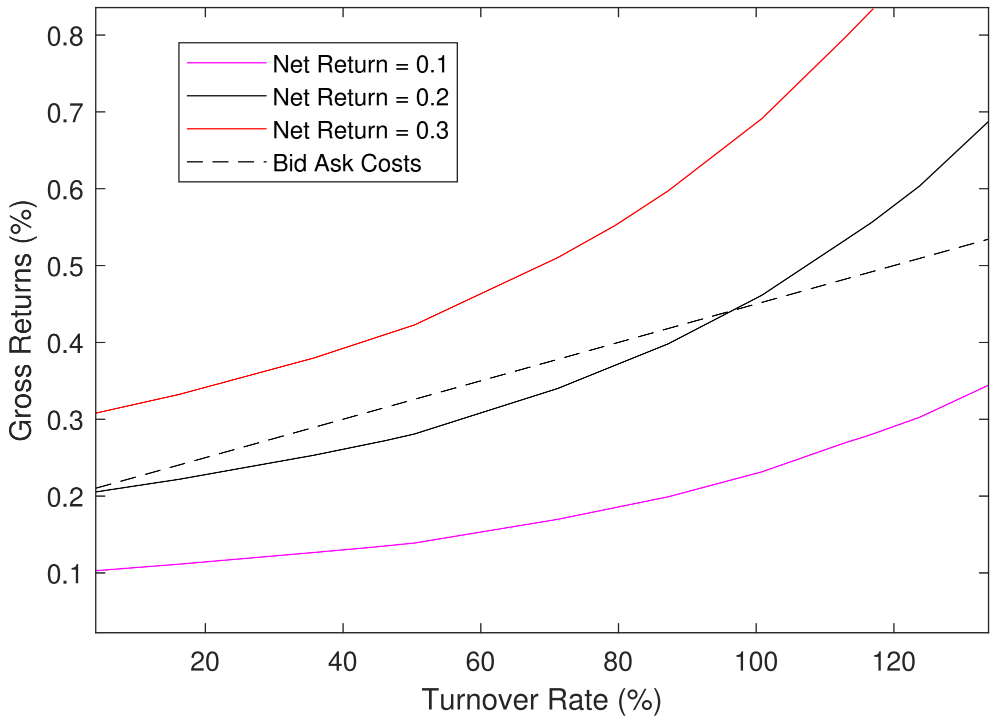

We plot the combination of the monthly portfolio turnover rate and the gross returns that produces a given level of net returns. Each line represents the indifference curve with a target level of net returns, ranging from 0.1% per month to 0.3% per month. The plot is based on the assumption that the average delay is 21 days ($\lambda = 1/21$). We vary ρ and σ_v to generate different values of turnover rates and gross returns. Bid-Ask Costs represent a line for (Gross Returns) = (Target Net Returns of 0.20%) + 0.19 * (Turnover Rate).

Figure 4: Median Half Spreads By Trade Size

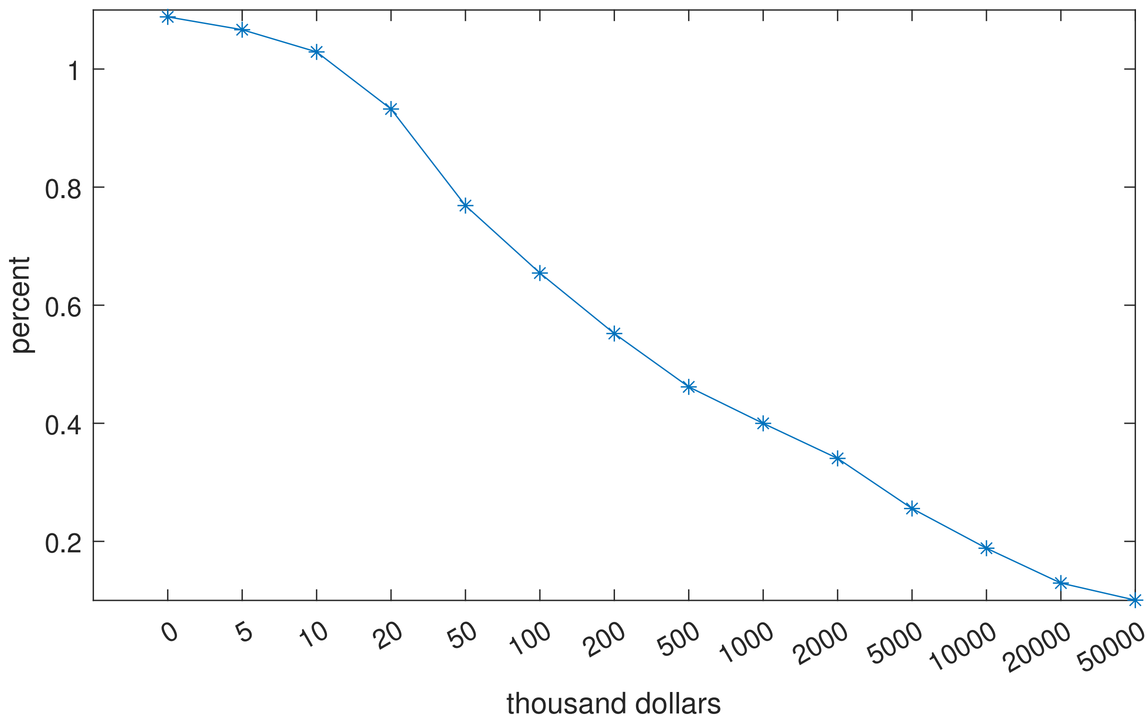

This figure plots the median half-spreads for bond transactions of given trade size. Half spreads are defined in Eq. (9). Each day, we compute the volume-weighted average of all transactions above a trade-size threshold and then take the average within the month. The figure plots the median values in the panel data given trade size.

Figure 5: Liquidity Supply and Demand

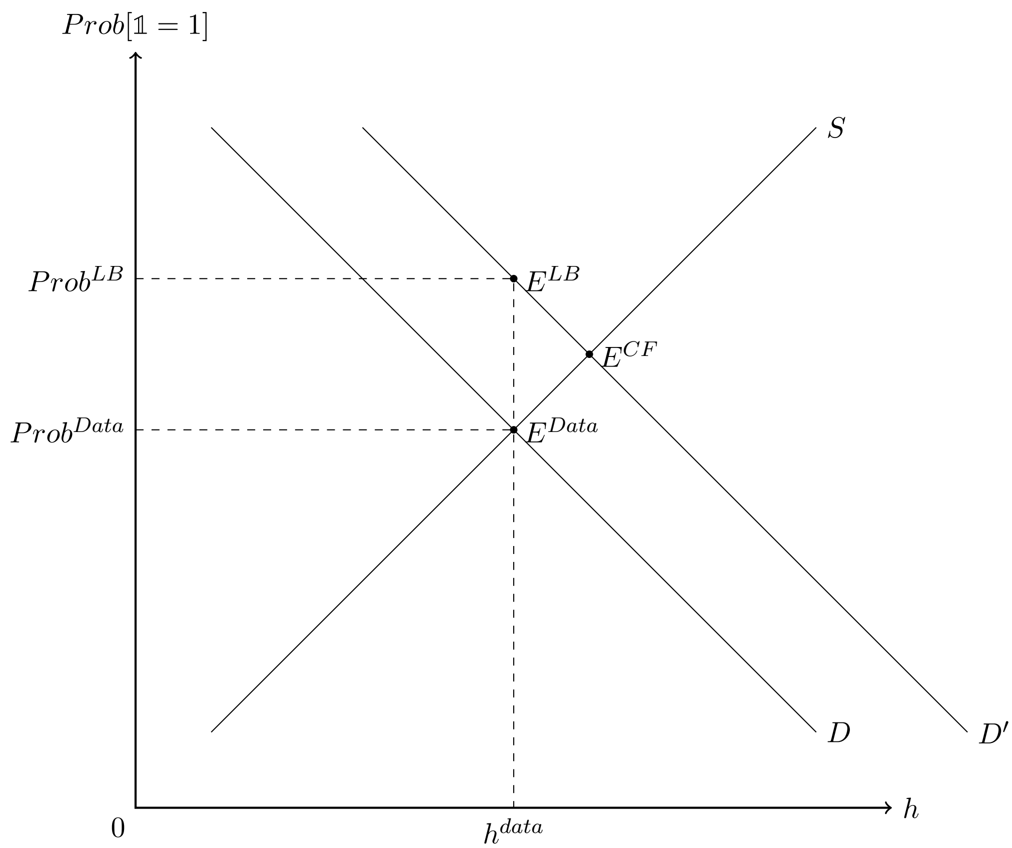

This figure illustrates the supply and demand curves for corporate bond liquidity, where the quantity is measured by the probability of transactions, $Prob[\mathbb{1} = 1]$, and the price is by half spreads charged by dealers, h. E^{Data} denotes the equilibrium observed in the data, which is the intersection between the supply curve S and the demand curve D. D' denotes the hypothetical demand curve when an active investor is introduced and E^{LB} is the hypothetical equilibrium if the supply curve is flat. E^{CF} is the counterfactual with an active investor and dealers responding to the introduction of the active investor.

Figure 6: Cost of Delay (Percent Per Month) by Trade Intensity

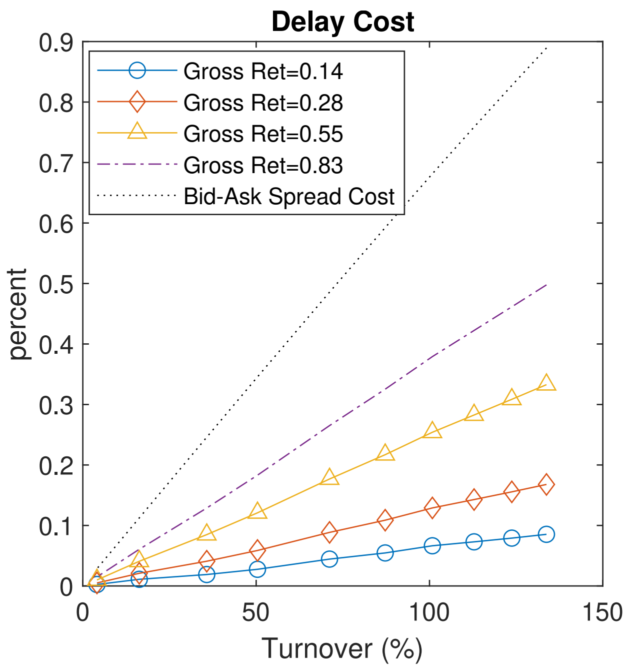

(a) $\lambda = 0.100$

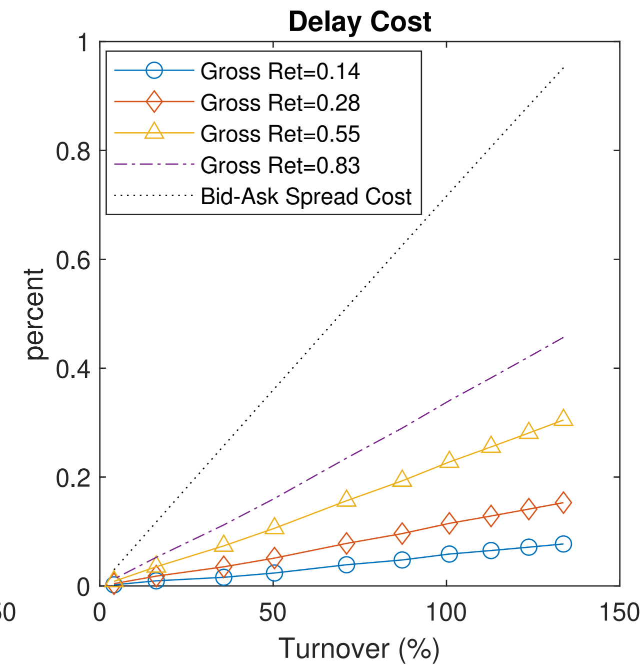

(b) $\lambda = 0.145$

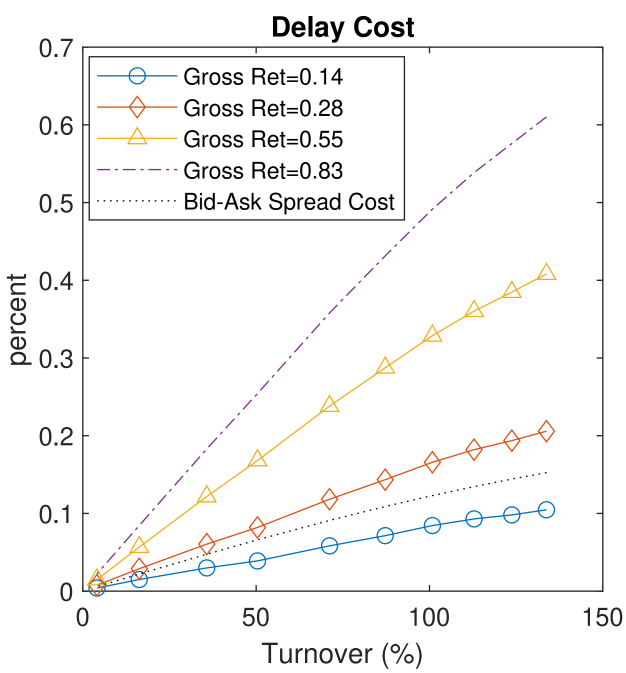

(c) $\lambda = 0.042$

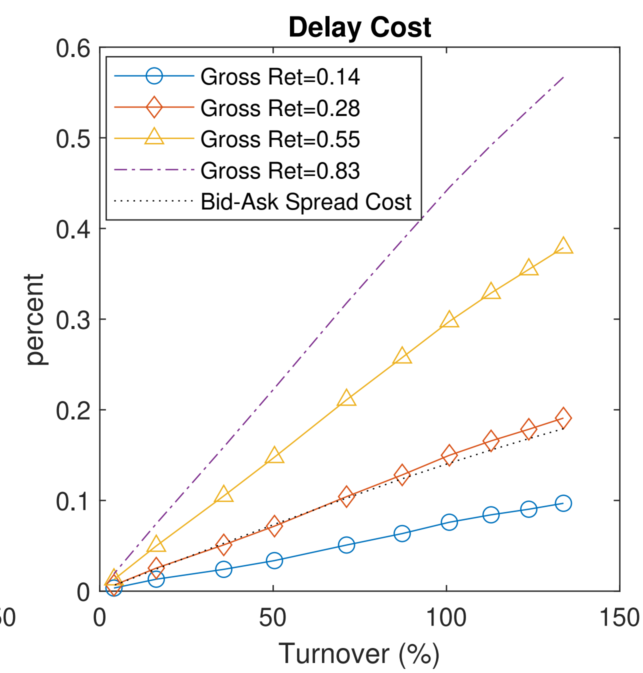

(d) $\lambda = 0.058$

We plot the difference in average returns between the gross factors that do not have delays and the net factors that are subject to delays. Delays follow an exponential distribution with parameter λ, which measures the trade intensity on a day. We vary signal persistence parameter ρ to change the portfolio turnover rate.

Figure 7: Net Returns (Percent Per Month) by Trade Intensity

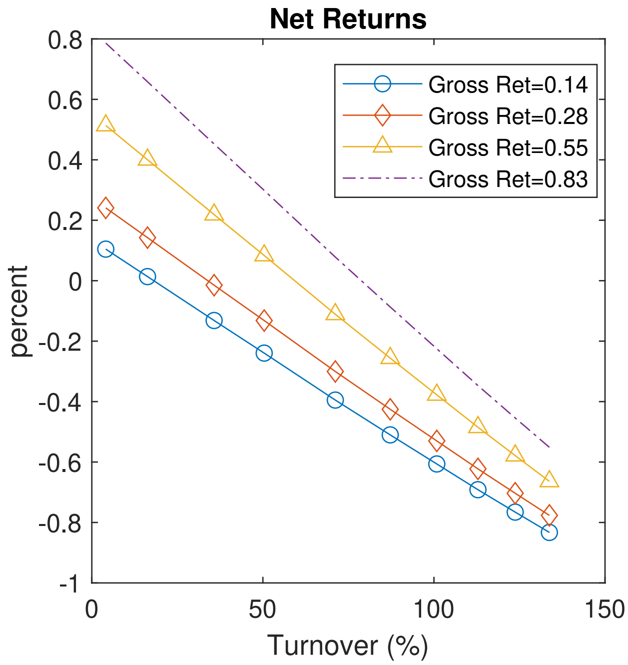

(a) $\lambda = 0.100$

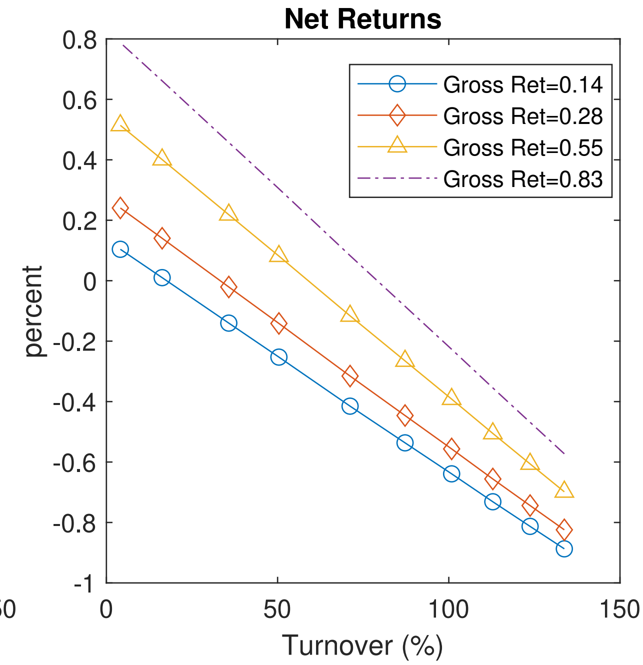

(b) $\lambda = 0.145$

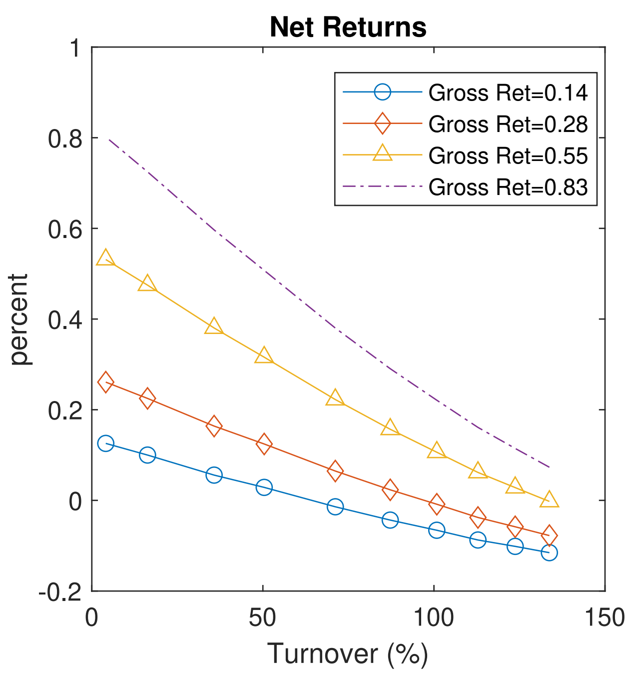

(c) $\lambda = 0.042$

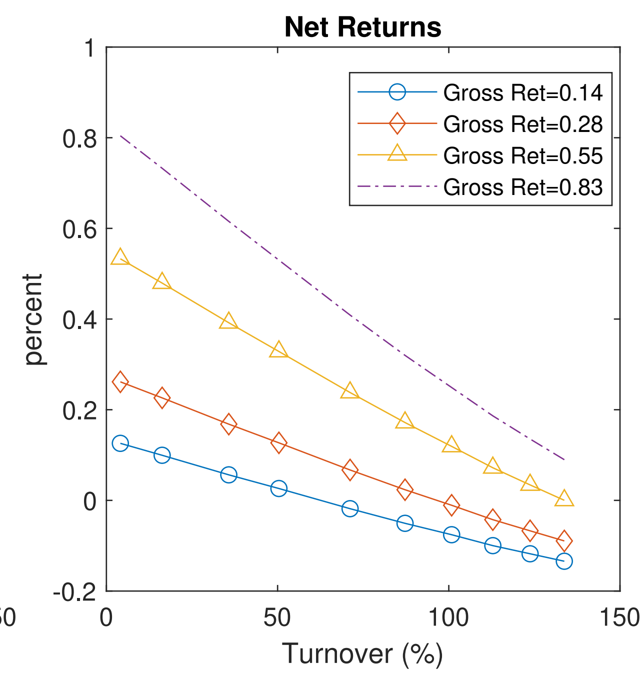

(d) $\lambda = 0.058$

We plot the net factors that are subject to delays and half spreads. Delays follow an exponential distribution with parameter λ, which measures the trade intensity on a day. We vary signal persistence parameter ρ to change the portfolio turnover rate. For Panels (a) and (b), we use a half spread of 37 bps and for Panels (c) and (d), we use a half spread of 9 bps.

Figure 8: Effect of Demand Frequency Adjustment on Delay Costs

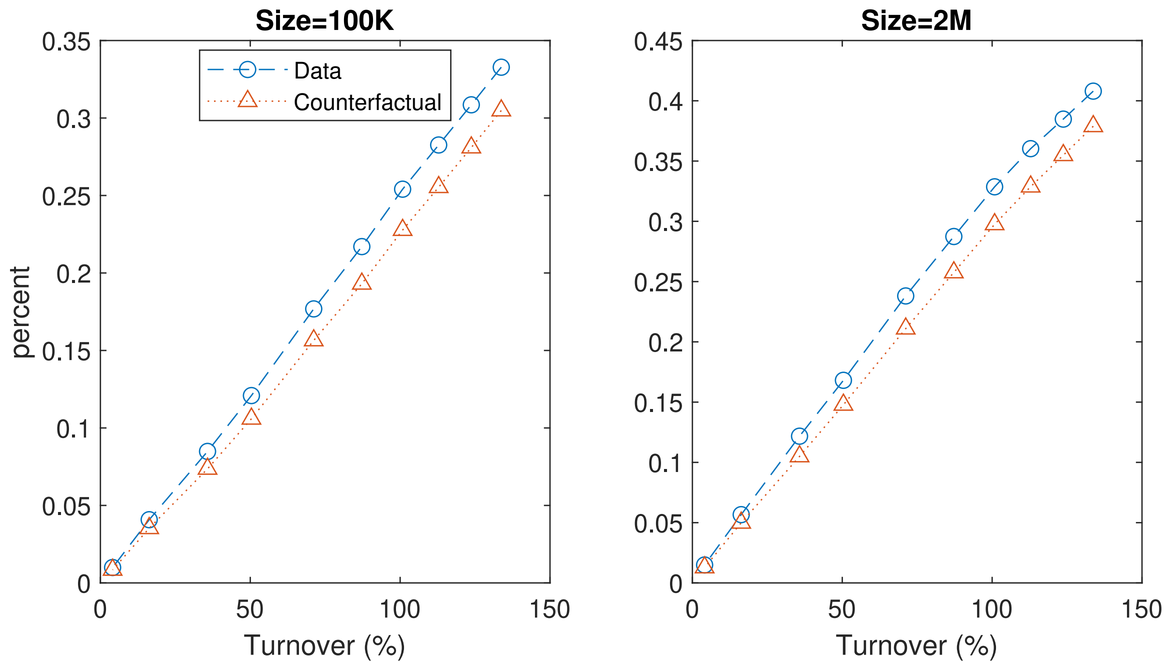

We plot the difference in average returns between the gross factors that have no delays and the net factors that have delays. Delays follow an exponential distribution with parameter λ, which measures the trading intensity on a day. We vary the signal persistence parameter ρ to change the portfolio turnover rate. For this plot, we set $\sigma_v = 0.2\% \times \sqrt{1-\rho^2}$ so that the gross factor return is 0.56% per month. The data are based on the trading intensity parameter λ estimated with the data, and the counterfactual is based on λ for the most active investor with the 99th percentile portfolio turnover rate. The left panel compares delay costs when the target trade size is \$100,000 and the right panel is for a trade size of \$2 million.

Figure 9: Median Delays Over Time

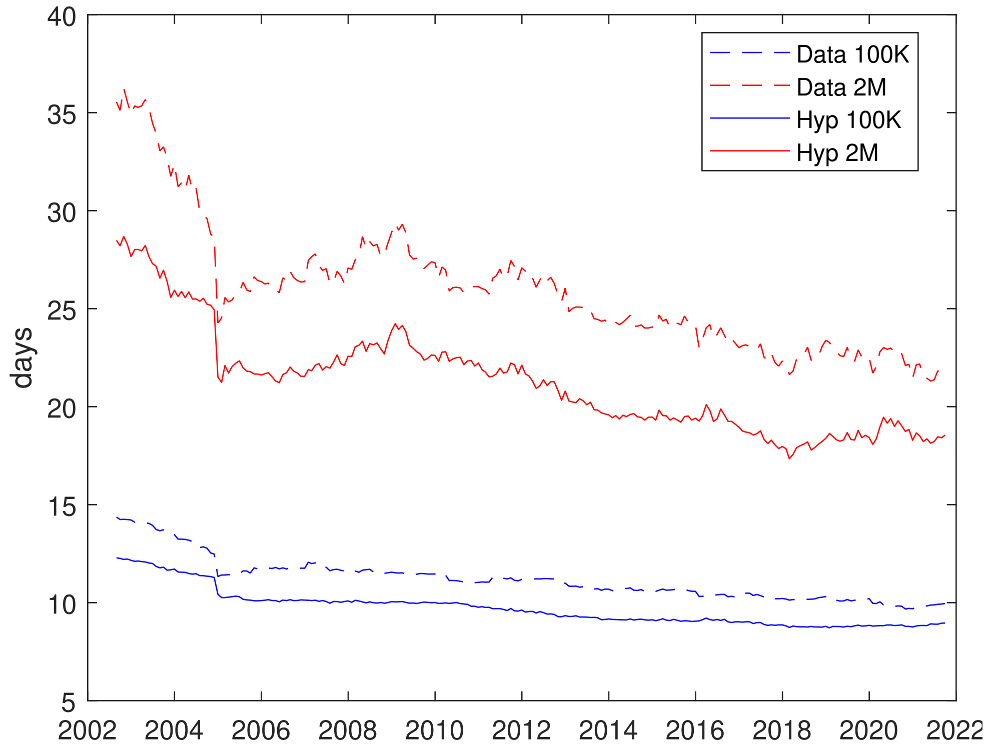

This figure plots the median delay in each month. We use the estimate of the logit regression reported in Table 4 and compute the expected delay ($1/\lambda$) for each bond in each month. The data uses the predictor of trades in the data and Hyp ('hypothetical') uses the 99th percentile of the investor composition measure instead of the values in the data. Each month, we compute the median across bond delays separately for the trades with \$100,000 and \$2 million.

Figure 10: Turnover Rate and Gross Returns of 341 Strategies

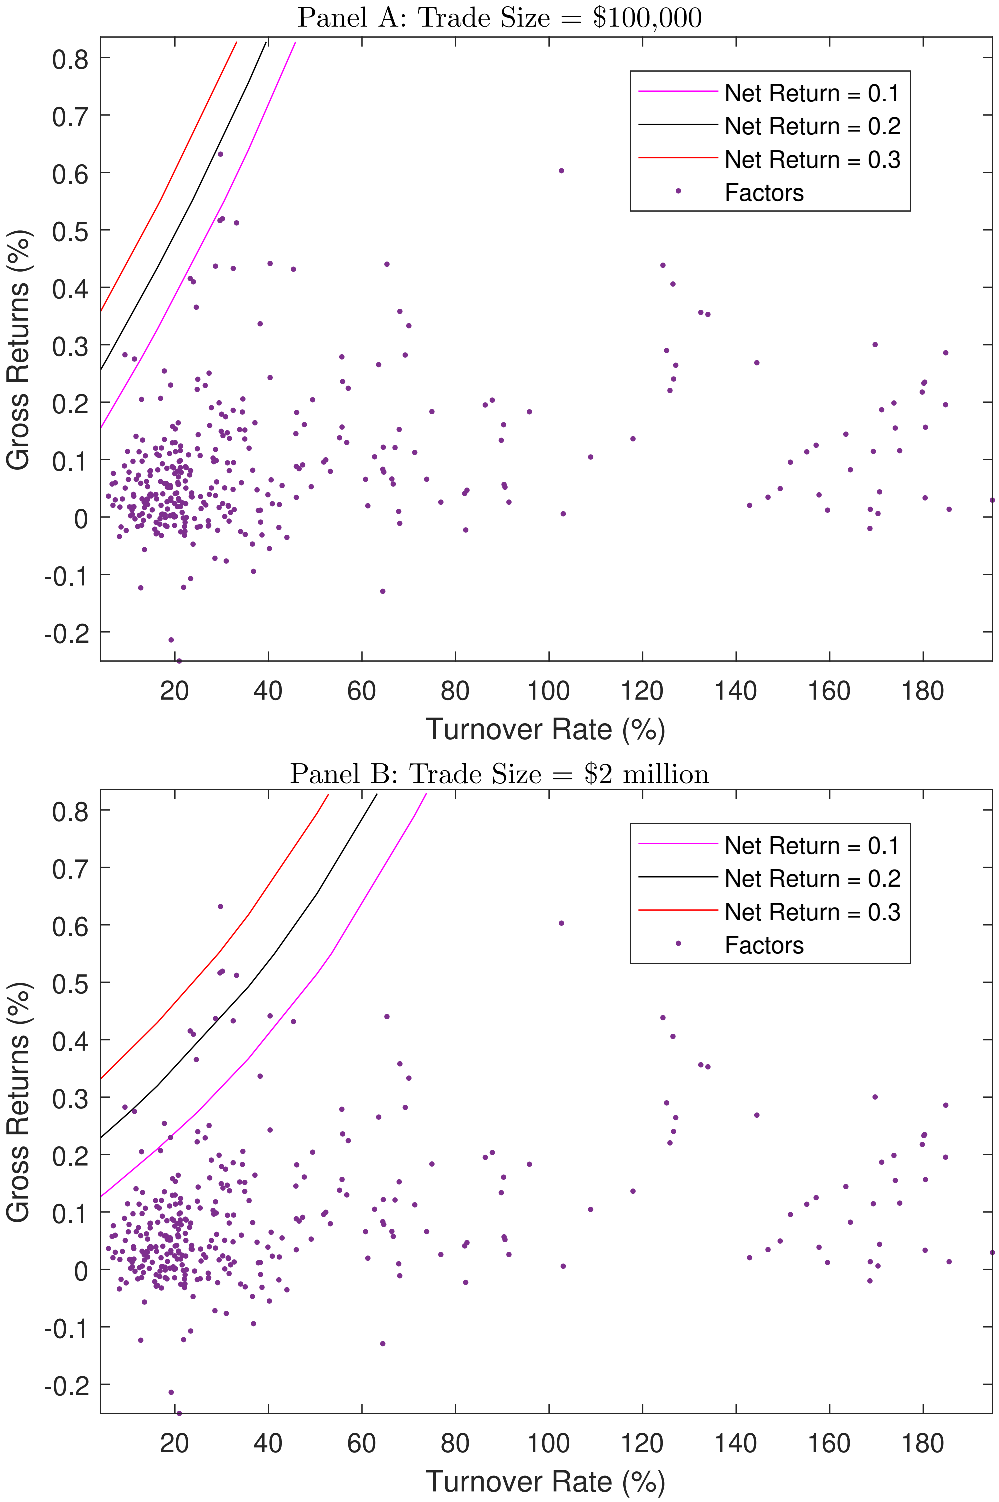

We plot the combination of monthly portfolio turnover rate and gross returns that produces a given level of net returns. Each curve represents the indifference curve with a target level of net returns, ranging from 0.1% per month to 0.3% per month. The plot is based on the estimated counterfactual trade intensity for a given trade size. We vary ρ and σ_v to generate different values of turnover rates and gross returns. A dot shows the gross returns and turnover rate of 341 corporate bond factors.

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

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

Table 2: Simulation Parameters

N	1,000	σ_f	2.00
T	244	σ_u	0.2
σ_v	$0.002\sqrt{1-\rho^2}$	$\mathrm{E}[f]$	1.00

This table reports the set of parameters used in the simulation analysis.

Table 3: Simulated Cost of Delay, Percent Per Month

Panel A. Fix $\sigma_v = 0.002$ and Vary λ

λ	Expected Delay	Turnover Rate (%) 4.1	16.3	35.7	50.4	71.2	87.2	100.8	112.9	123.7	133.7
1.000	1	0.005	0.022	0.048	0.071	0.111	0.143	0.173	0.198	0.223	0.247
0.143	7	0.008	0.035	0.074	0.107	0.157	0.194	0.229	0.256	0.282	0.306
0.048	21	0.014	0.054	0.115	0.160	0.227	0.275	0.316	0.347	0.373	0.396
0.024	42	0.019	0.070	0.153	0.210	0.289	0.341	0.383	0.410	0.432	0.452
0.004	252	0.056	0.168	0.321	0.397	0.461	0.491	0.513	0.522	0.530	0.537

Panel B. Fix $\lambda = 0.048$ and Vary σ_v

σ_v	Gross Returns	Turnover Rate (%) 4.1	16.3	35.7	50.4	71.2	87.2	100.8	112.9	123.7	133.7
0.0005	0.136	0.004	0.014	0.028	0.037	0.055	0.068	0.081	0.089	0.095	0.101
0.0010	0.274	0.007	0.027	0.057	0.078	0.113	0.137	0.159	0.175	0.187	0.199
0.0020	0.552	0.014	0.054	0.115	0.160	0.227	0.275	0.316	0.347	0.373	0.396
0.0030	0.830	0.021	0.080	0.174	0.242	0.342	0.413	0.473	0.520	0.558	0.593

This table reports the simulated cost of delays, which is the difference between gross and net returns. Half spreads are set to zero so that the difference only reflects the cost of delay. λ is the trade intensity and σ_v is the volatility of factors. We use the speed of mean reversion of the signal, ρ, ranging from 0.3 to 0.9995 which corresponds to the monthly turnover rate of 4.1% to 133.7%. Other parameters are $T = 244$, $N = 1,000$, and the number of paths is 1,000.

Table 4: Forecasting Regression of Corporate Bond Transactions

		Customer Sell				Customer Buy			
		\$100K		\$2M		\$100K		\$2M	
Panel A. Regression Estimates									
log invcomp		0.018	(22.74)	0.028	(30.23)	0.017	(21.11)	0.030	(31.42)
inventory28		-0.010	(-37.29)	-0.047	(-63.54)	0.009	(28.94)	0.045	(55.28)
capratio		0.002	(1.67)	0.016	(17.02)	0.000	(-0.22)	0.013	(13.78)
rating		0.054	(29.14)	0.109	(60.60)	0.054	(28.65)	0.111	(61.45)
maturity		-0.017	(-9.26)	-0.008	(-5.71)	-0.020	(-11.60)	-0.002	(-1.68)
log facevalue		0.099	(66.64)	0.206	(151.89)	0.104	(68.30)	0.203	(142.50)
coupon		-0.028	(-16.18)	-0.003	(-1.52)	-0.028	(-15.97)	0.002	(1.34)
age		-0.031	(-17.91)	-0.077	(-44.15)	-0.036	(-20.22)	-0.078	(-45.64)
Intercept		0.877	(623.95)	0.588	(463.61)	0.868	(610.82)	0.576	(450.68)
Adj. R^2		0.142		0.251		0.151		0.248	
N		814,145		814,145		814,145		814,145	
Panel B. Model-Based Estimates of Trade Intensity and Delays									
Prob. Trade	$Prob^{Data}$	0.877		0.588		0.868		0.576	
	$Prob^{LB}$	0.931		0.673		0.920		0.665	
λ	λ^{Data}	0.100		0.042		0.097		0.041	
	λ^{LB}	0.127		0.053		0.120		0.052	
Exp. Delay	data	10.014		23.672		10.362		24.442	
(days)	cntr. factual	7.873		18.779		8.320		19.184	

This table reports the coefficient estimates for the regressions of the bond trading dummy on characteristics in Equation (10). Variable invcomp is the investor composition of Li and Yu (2023), which measures the activeness of bond investors at the end of the previous quarter. Variable inventory28 is the difference between customer buys and customer sells in the preceding 28 days, capratio is the intermediary capital ratio of He et al. (2017) on the previous month. The right-hand-side variables are standardized for ease of interpretation. Values in parentheses are t-statistics clustered at the bond level. The sample is at the month frequency, spanning August 2002 to November 2022.

Table 5: First Stage Regression of Half Spreads on Instruments

	Demand Curve Estimate		Supply Curve Estimate	
capratio	-0.250	(-27.94)		
inventory28	0.015	(2.91)		
log invcomp			0.040	(4.74)
rat	0.524	(23.82)	0.521	(23.42)
tau	0.331	(27.55)	0.334	(27.97)
log facevalue	-0.022	(-2.14)	-0.010	(-0.96)
coupon	0.158	(11.58)	0.152	(11.23)
age	0.138	(11.58)	0.147	(12.24)
Intercept	1.655	(166.25)	1.661	(169.10)
Adj. R^2	0.062		0.055	
N	703,346		703,346	

This table reports the estimates of the panel regressions of half spreads on the set of instruments. Half spreads for each trade is computed as the difference between a dealer-customer transaction price and the most recent interdealer transaction price. Then, the trade-level data is aggregated at the monthly frequency. Variable invcomp is the investor composition of Li and Yu (2023), which measures the activeness of bond investors at the end of the previous quarter. Variable inventory28 is the difference between customer buys and customer sells in the preceding 28 days, capratio is the intermediary capital ratio of He et al. (2017) on the previous month. The right-hand-side variables are standardized for ease of interpretation. Values in parentheses are t-statistics clustered at the bond level. The sample is at the month frequency, spanning August 2002 to November 2022.

Table 6: IV Estimates of the Liquidity Demand and Supply Curves

		Customer Sell				Customer Buy			
		\$100K		\$2M		\$100K		\$2M	
Panel A. Regression Estimates									
Liquidity Demand									
$\hat{h}$		-0.008	(-2.19)	-0.075	(-19.76)	0.004	(1.00)	-0.040	(-10.36)
rating		0.061	(21.66)	0.153	(54.57)	0.055	(19.40)	0.136	(47.96)
maturity		-0.016	(-7.46)	0.015	(7.78)	-0.023	(-11.10)	0.008	(4.12)
log facevalue		0.101	(66.82)	0.207	(146.11)	0.107	(69.18)	0.207	(148.17)
coupon		-0.024	(-13.01)	0.013	(7.04)	-0.026	(-13.68)	0.015	(8.05)
age		-0.034	(-18.73)	-0.072	(-38.80)	-0.040	(-21.90)	-0.079	(-42.98)
Intercept		0.890	(148.69)	0.712	(110.67)	0.862	(142.08)	0.643	(98.44)
Adj. R^2		0.138		0.239		0.148		0.237	
N		814,145		814,145		814,145		814,145	
Liquidity Supply									
$\hat{h}$		0.447	(22.78)	0.710	(30.27)	0.432	(21.14)	0.739	(31.33)
rating		-0.179	(-17.68)	-0.261	(-21.15)	-0.171	(-16.18)	-0.275	(-22.14)
maturity		-0.166	(-25.23)	-0.245	(-31.18)	-0.164	(-23.91)	-0.249	(-31.62)
log facevalue		0.103	(67.92)	0.210	(151.47)	0.109	(70.17)	0.210	(154.40)
coupon		-0.096	(-25.70)	-0.111	(-26.85)	-0.094	(-24.21)	-0.109	(-26.03)
age		-0.097	(-30.43)	-0.182	(-50.99)	-0.100	(-30.52)	-0.187	(-52.05)
Intercept		0.134	(4.07)	-0.591	(-15.12)	0.151	(4.38)	-0.651	(-16.57)
Adj. R^2		0.141		0.241		0.151		0.239	
N		814,145		814,145		814,145		814,145	
Panel B. Counterfactuals									
Prob. Trade	Data	0.877		0.588		0.868		0.576	
	$Prob^{LB}$	0.931		0.673		0.920		0.665	
	$Prob^{CF}$	0.930		0.665		0.920		0.661	
$h^{CF} - h^{Data}$		0.117		0.108		0.121		0.114	

Panel A reports the estimates of the liquidity demand and supply functions in Equations (7) and (6). We regress the bond trading dummies on the half spreads and the set of control variables. $\hat{h}$ is an instrumented half spread with the intermediary capital ratio of He et al. (2017) and 28-day dealer inventory changes for the demand function estimates and with the investor composition measure of Li and Yu (2023) for the supply function. Values in parentheses are t-statistics clustered at the bond level. The sample is at the monthly frequency, spanning August 2002 to November 2022. Panel B reports the average probability of trades in the data and the counterfactuals obtained from the estimated demand function. $Prob^{LB}$ is the counterfactual due to a parallel shift in the demand curve by setting the $InvComp$ at the 99th percentile. $Prob^{CF}$ and h^{CF} are the counterfactual equilibrium probability of transactions and half spreads, respectively, which account for the slope of the demand and supply curves.

Table 7: Forecasting Regression of Corporate Bond Transactions By Credit Rating

		IG Bonds: Customer Sell				HY Bonds: Customer Sell			
		\$100K		\$2M		\$100K		\$2M	
Panel A. Regression Estimates									
log invcomp		0.016	(16.71)	0.034	(31.44)	0.011	(9.54)	0.009	(6.61)
inventory28		-0.001	(-31.39)	-0.005	(-55.85)	-0.001	(-19.40)	-0.004	(-35.00)
capratio		0.006	(0.09)	1.458	(20.09)	0.234	(2.35)	0.534	(4.76)
rating		0.009	(11.92)	0.018	(24.89)	0.007	(6.54)	0.018	(16.84)
maturity		-0.001	(-6.94)	-0.001	(-5.13)	-0.003	(-4.96)	-0.002	(-4.08)
log facevalue		0.167	(58.09)	0.323	(137.92)	0.117	(34.52)	0.286	(76.50)
coupon		-0.017	(-14.84)	0.004	(3.75)	-0.010	(-7.61)	-0.010	(-6.69)
age		0.000	(-14.06)	0.000	(-40.41)	0.000	(-8.48)	0.000	(-18.37)
Intercept		-0.076	(-3.56)	-1.521	(-81.77)	0.262	(9.02)	-1.117	(-34.50)
Adj. R^2		0.149		0.250		0.119		0.222	
N		580,855		580,855		233,290		233,290	
Panel B. Model-Based Estimates of Trade Intensity and Delays									
Prob. Trade	$Prob^{Data}$	0.862		0.545		0.914		0.696	
	$Prob^{LB}$	0.947		0.724		0.966		0.740	
λ	λ^{Data}	0.094		0.037		0.117		0.057	
	λ^{LB}	0.140		0.061		0.161		0.064	
Exp. Delay	data	10.589		26.694		8.559		17.613	
(days)	cntr. factual	7.148		16.310		6.206		15.584	

This table reports the coefficient estimates for the logit regressions of the bond trading dummy on characteristics in Equation (10). The left panel uses IG bonds and the right panel uses HY bonds. Variable invcomp is the investor composition of Li and Yu (2023), which measures the activeness of bond investors at the end of the previous quarter. Variable inventory7 is the difference between customer buys and customer sells in the preceding seven days, capratio is the intermediary capital ratio of He et al. (2017) on the previous month. The right-hand-side variables are standardized for ease of interpretation. Values in parentheses are t-statistics clustered at the bond level. The sample is at the month frequency, spanning August 2002 to November 2022.

Table 8: CAPM α of 341 Bond Risk Factors After Costs of Delays

	Distribution				# of Factors w. CAPM Alpha			# of Significant Factors
	Mean	10%	50%	90%	[0%, 0.1%)	[0.1%, 0.2%)	[0.2%, ∞)	
Panel A. Performance Before Transaction Costs								
Gross α	0.093	-0.017	0.060	0.241	167	69	50	56
Panel B. Lower Bound Estimates of Delay								
Panel B1. Trade Size of $100K								
Net α	-0.411	-1.301	-0.237	-0.039	14	3	0	0
DelayCost	0.026	0.001	0.008	0.077				
BACost	0.478	0.092	0.300	1.392				
Panel B2. Trade Size of $2M								
Net α	-0.141	-0.500	-0.078	0.027	45	6	5	0
DelayCost	0.034	0.002	0.012	0.101				
BACost	0.199	0.041	0.130	0.564				
Panel C. Counterfactual Estimates of Delay								
Panel C1. Trade Size of $100K								
Net α	-0.498	-1.571	-0.286	-0.055	11	1	0	0
DelayCost	0.026	0.001	0.008	0.077				
BACost	0.565	0.108	0.354	1.645				
Panel C2. Trade Size of $2M								
Net α	-0.204	-0.677	-0.115	0.005	29	8	1	0
DelayCost	0.035	0.002	0.012	0.102				
BACost	0.263	0.054	0.172	0.743				

This table reports the distribution of the bond CAPM alphas across the 341 risk factors we generate. Every month, we sort bonds into deciles based on the latest value of the signals and form value-weighted portfolios. A factor is the return difference between the tenth and first deciles. The factors are signed to guarantee that their gross average returns are non-negative (i.e., the factors are sign-corrected to be increasing in expected returns). Panel A reports the distribution of the CAPM α before transaction costs. Panel B reports the distribution using the lower bound estimates for transaction costs and Panel C reports that using the counterfactual equilibrium estimates. The sample is from August 2002 to November 2022. t-statistics are adjusted by Newey-West 12 lags.

Table 9: Profitability of Bond Factors, Top 15 Gross α with Transaction Costs at the Lower Bound

			Trade Size: $100K			Trade Size: $2M		
Signal	Turnover	Gross α	Net α	DelayCost	BACost	Net α	DelayCost	BACost
ret10	175	0.691	-1.314	0.360	1.644	-0.399	0.441	0.649
		(3.26)	(-6.19)			(-1.88)		
bondprice	30	0.632	0.159	0.100	0.373	0.331	0.139	0.162
		(2.27)	(0.57)			(1.19)		
mom6mspread	103	0.603	-0.979	0.309	1.273	-0.316	0.396	0.522
		(2.71)	(-4.40)			(-1.42)		
yldtoworst	30	0.519	0.040	0.101	0.378	0.215	0.141	0.164
		(2.11)	(0.16)			(0.87)		
yield	30	0.516	0.045	0.099	0.372	0.217	0.138	0.161
		(2.10)	(0.18)			(0.88)		
spread	33	0.512	-0.015	0.111	0.416	0.178	0.154	0.180
		(2.06)	(-0.06)			(0.71)		
sprtod2d	40	0.441	-0.155	0.091	0.506	0.097	0.127	0.218
		(3.14)	(-1.10)			(0.69)		
seas11na	65	0.440	-0.459	0.083	0.816	-0.017	0.111	0.346
		(3.16)	(-3.29)			(-0.12)		
TrendFactor	124	0.438	-1.290	0.195	1.533	-0.416	0.242	0.612
		(2.60)	(-7.64)			(-2.47)		
dts	29	0.437	-0.019	0.096	0.360	0.147	0.134	0.156
		(1.93)	(-0.08)			(0.65)		
swapspread	32	0.433	-0.076	0.101	0.407	0.115	0.141	0.176
		(1.91)	(-0.33)			(0.51)		
strucvalue	45	0.432	-0.243	0.106	0.568	0.040	0.147	0.244
		(2.83)	(-1.59)			(0.26)		
spreadvol	23	0.415	0.075	0.049	0.292	0.221	0.068	0.127
		(2.07)	(0.37)			(1.10)		
ratingxspread	24	0.410	0.042	0.067	0.301	0.185	0.094	0.131
		(1.56)	(0.16)			(0.71)		
mom3mspread	126	0.406	-1.405	0.253	1.559	-0.528	0.313	0.621
		(2.56)	(-8.86)			(-3.33)		

This table reports the CAPM α of the 15 best signals and the transaction cost estimates based on the simulation. The simulated costs are assigned to each strategy using the strategy's gross returns and turnover rates. Net α is the difference between gross α and the simulated total costs. In computing simulated costs, we use the lower bound estimates of trade intensity λ^{LB} reported in Table 4.

Table 10: CAPM α of 341 Bond Factors: 3-Month Holding Period

	Distribution				# of Factors w. CAPM Alpha			# of Significant Factors
	Mean	10%	50%	90%	[0%, 0.1%)	[0.1%, 0.2%)	[0.2%, ∞)	
Panel A. Performance Before Transaction Costs								
Gross α	0.082	-0.009	0.061	0.213	190	72	35	36
Panel B. Lower Bound Estimates of Delay								
Panel B1. Trade Size of $100K								
Net α	-0.218	-0.517	-0.172	-0.024	19	4	1	0
DelayCost	0.012	0.001	0.006	0.029				
BACost	0.289	0.077	0.241	0.622				
Panel B2. Trade Size of $2M								
Net α	-0.060	-0.207	-0.046	0.049	60	9	7	0
DelayCost	0.017	0.001	0.009	0.040				
BACost	0.125	0.034	0.105	0.267				
Panel C. Counterfactual Estimates of Delay								
Panel C1. Trade Size of $100K								
Net α	-0.271	-0.625	-0.220	-0.034	12	3	1	0
DelayCost	0.012	0.001	0.006	0.029				
BACost	0.341	0.091	0.285	0.735				
Panel C2. Trade Size of $2M								
Net α	-0.100	-0.285	-0.080	0.025	42	9	4	0
DelayCost	0.017	0.001	0.009	0.041				
BACost	0.165	0.045	0.139	0.352				

This table reports the distribution of the bond CAPM alphas across the 341 factors we generate. Every month, we sort bonds into deciles based on the latest value of signals and form value-weighted portfolios. We hold the bonds over a staggered three-month period, such that only 33.33% of the portfolio is rebalanced each month (Jegadeesh and Titman, 1993). A portfolio return in month t is the average of the three strategies run in the month. A factor is the return difference between the tenth and first deciles. The factors are signed to guarantee that their gross average returns are non-negative (i.e., the factors are sign-corrected to be increasing in expected returns). Panel A reports the distribution of the CAPM α before transaction costs. Panel B reports the distribution using the lower bound estimates for transaction costs and Panel C reports that using the counterfactual equilibrium estimates. The sample is from August 2002 to November 2022. t-statistics are adjusted by Newey-West 12 lags.

Table 11: Bond Factors, Top 15 Gross α: 3-Month Holding Period

			Trade Size: \$100K			Trade Size: \$2M		
Signal	Turnover	Gross α	Net α	DelayCost	BACost	Net α	DelayCost	BACost
bondprice	18	0.547	0.272	0.052	0.223	0.377	0.072	0.097
		(2.08)	(1.03)			(1.44)		
mom6mspread	53	0.430	-0.345	0.115	0.660	-0.011	0.159	0.282
		(2.10)	(-1.69)			(-0.05)		
spread	20	0.409	0.099	0.058	0.252	0.218	0.081	0.110
		(1.61)	(0.39)			(0.86)		
yldtoworst	21	0.402	0.080	0.065	0.258	0.200	0.090	0.112
		(1.63)	(0.32)			(0.81)		
yield	20	0.402	0.082	0.064	0.255	0.202	0.089	0.111
		(1.63)	(0.33)			(0.82)		
ret10	60	0.370	-0.455	0.079	0.747	-0.055	0.107	0.318
		(3.38)	(-4.16)			(-0.50)		
ratingxspread	16	0.370	0.134	0.041	0.196	0.228	0.056	0.086
		(1.37)	(0.49)			(0.84)		
dts	21	0.354	0.020	0.067	0.267	0.145	0.093	0.116
		(1.70)	(0.10)			(0.69)		
swapspread	20	0.352	0.048	0.054	0.249	0.168	0.075	0.109
		(1.51)	(0.21)			(0.72)		
spreadvol	19	0.339	0.067	0.035	0.238	0.188	0.048	0.104
		(1.51)	(0.30)			(0.84)		
rmax521d	48	0.322	-0.351	0.075	0.598	-0.039	0.104	0.257
		(1.76)	(-1.91)			(-0.21)		
seas11na	38	0.316	-0.186	0.027	0.475	0.072	0.040	0.205
		(2.55)	(-1.49)			(0.58)		
rating	9	0.314	0.178	0.017	0.119	0.238	0.024	0.053
		(1.15)	(0.65)			(0.87)		
sizeb	10	0.310	0.166	0.014	0.130	0.233	0.019	0.057
		(1.85)	(0.99)			(1.39)		
strucvalue	26	0.304	-0.067	0.047	0.324	0.098	0.066	0.141
		(2.02)	(-0.44)			(0.65)		

This table reports the CAPM α of the 15 best signals and the transaction cost estimates based on the simulation. The simulated costs are assigned to each strategy using the strategy's gross returns and turnover rates. Net α is the difference between gross α and the simulated total costs. In computing simulated costs, we use the lower bound estimates of trade intensity λ^{LB} reported in Table 4.

Table 12: CAPM α of 341 Bond Factors: Banding

	Distribution				# of Factors w. CAPM Alpha			# of Significant Factors
	Mean	10%	50%	90%	[0%, 0.1%)	[0.1%, 0.2%)	[0.2%, ∞)	
Panel A. Performance Before Transaction Costs								
Gross α	0.078	-0.010	0.058	0.210	194	61	37	45
Panel B. Lower Bound Estimates of Delay								
Panel B1. Trade Size of $100K								
Net α	-0.244	-0.695	-0.107	0.016	29	11	1	0
DelayCost	0.015	0.001	0.005	0.036				
BACost	0.308	0.061	0.158	0.826				
Panel B2. Trade Size of $2M								
Net α	-0.071	-0.273	-0.023	0.050	81	15	8	1
DelayCost	0.019	0.001	0.006	0.046				
BACost	0.129	0.027	0.069	0.350				
Panel C. Counterfactual Estimates of Delay								
Panel C1. Trade Size of $100K								
Net α	-0.300	-0.835	-0.140	-0.001	25	8	1	0
DelayCost	0.015	0.001	0.005	0.036				
BACost	0.363	0.072	0.186	0.976				
Panel C2. Trade Size of $2M								
Net α	-0.112	-0.366	-0.045	0.036	57	11	8	1
DelayCost	0.020	0.001	0.007	0.046				
BACost	0.170	0.036	0.092	0.462				

This table reports the distribution of the bond CAPM alphas of the 341 factors we create. Every month, we sort bonds into deciles based on the latest value of signals and form value-weighted portfolios. Once a bond joins portfolio 1 (10), we hold the bonds in the portfolio as long as these bonds remain in the first (last) three deciles (Novy-Marx and Velikov, 2015). A factor is the return difference between portfolios 10 and 1. The factors are signed to guarantee that their gross average returns are non-negative (i.e., the factors are sign-corrected to be increasing in expected returns). Panel A reports the distribution of the CAPM α before transaction costs. Panel B reports the distribution using the lower bound estimates for transaction costs and Panel C reports that using the counterfactual equilibrium estimates. The sample is from August 2002 to November 2022. t-statistics are adjusted by Newey-West 12 lags.

Table 13: Bond Factors, Top 15 Gross α: Banding

			Trade Size: \$100K			Trade Size: \$2M		
Signal	Turnover	Gross α	Net α	DelayCost	BACost	Net α	DelayCost	BACost
ret10	142	0.681	-1.318	0.355	1.644	-0.403	0.435	0.649
		(3.48)	(-6.73)			(-2.06)		
mom6mspread	60	0.486	-0.404	0.144	0.746	-0.026	0.194	0.318
		(2.57)	(-2.14)			(-0.14)		
bondprice	11	0.459	0.290	0.027	0.142	0.359	0.038	0.062
		(2.39)	(1.51)			(1.87)		
mom3mspread	87	0.361	-0.880	0.156	1.085	-0.296	0.205	0.452
		(2.42)	(-5.88)			(-1.98)		
IndRetBig	67	0.352	-0.594	0.105	0.841	-0.144	0.140	0.356
		(2.60)	(-4.39)			(-1.07)		
TrendFactor	78	0.347	-0.718	0.096	0.969	-0.187	0.127	0.407
		(2.57)	(-5.33)			(-1.39)		
yldtoworst	12	0.331	0.152	0.032	0.147	0.222	0.045	0.065
		(1.98)	(0.91)			(1.32)		
yield	11	0.331	0.156	0.032	0.144	0.224	0.044	0.063
		(1.97)	(0.93)			(1.33)		
spread	10	0.328	0.176	0.025	0.128	0.237	0.034	0.056
		(1.74)	(0.93)			(1.26)		
kurt	12	0.328	0.166	0.015	0.146	0.242	0.021	0.064
		(2.91)	(1.48)			(2.15)		
spreadvol	10	0.315	0.168	0.017	0.130	0.234	0.024	0.057
		(1.68)	(0.90)			(1.25)		
ret31	96	0.312	-0.978	0.094	1.197	-0.302	0.121	0.494
		(2.89)	(-9.06)			(-2.80)		
dVolCall	164	0.297	-1.515	0.167	1.644	-0.557	0.205	0.649
		(3.60)	(-18.38)			(-6.76)		
ratingxspread	8	0.282	0.166	0.016	0.100	0.215	0.022	0.044
		(1.37)	(0.81)			(1.05)		
rmax521d	67	0.278	-0.662	0.100	0.840	-0.211	0.133	0.356
		(1.66)	(-3.96)			(-1.26)		

This table reports the bond CAPM α of the 15 best signals and the transaction cost estimates based on simulation. The simulated costs are assigned to each strategy using the strategy's gross returns and turnover rates. Net α is the difference between gross α and the simulated total costs. In computing simulated costs, we use the lower bound estimates of trade intensity λ^{LB} reported in Table 4.

Table 14: Profitability of ML Strategies Based On Simulated Cost of Transactions

			Trade Size: \$100K			Trade Size: \$2M		
Signal	Turnover	Gross α	Net α	DelayCost	BACost	Net α	DelayCost	BACost
enet	124	0.731	-1.231	0.433	1.529	-0.419	0.539	0.611
		(3.34)	(-5.63)			(-1.92)		
ens	124	0.808	-1.158	0.438	1.528	-0.347	0.545	0.611
		(3.26)	(-4.67)			(-1.40)		
lasso	140	0.698	-1.396	0.450	1.644	-0.501	0.551	0.649
		(3.20)	(-6.40)			(-2.30)		
lens	125	0.795	-1.186	0.441	1.540	-0.367	0.548	0.615
		(3.49)	(-5.21)			(-1.61)		
nens	114	0.778	-1.040	0.404	1.415	-0.305	0.511	0.572
		(3.13)	(-4.19)			(-1.23)		
nn	136	0.659	-1.421	0.436	1.644	-0.523	0.533	0.649
		(3.09)	(-6.67)			(-2.46)		
rf	103	0.469	-1.050	0.244	1.276	-0.366	0.313	0.523
		(2.41)	(-5.39)			(-1.88)		
ridge	116	0.747	-1.076	0.391	1.433	-0.324	0.492	0.578
		(3.61)	(-5.20)			(-1.56)		
xt	99	0.810	-0.771	0.350	1.232	-0.148	0.451	0.507
		(3.40)	(-3.24)			(-0.62)		

This table reports the bond CAPM α of the ML algorithms and the transaction cost estimates based on the simulation. The simulated costs are assigned to each ML factor using the factor's gross returns and turnover rates. Net α is the difference between gross α and the simulated total costs. In computing simulated costs, we set the counterfactual trade intensity parameter $\lambda = 0.145$ and $\lambda = 0.058$ to replicate the trade frequency of trade size \$100,000 (left panel) and \$2 million (right panel). The linear models with penalization comprise an elastic net (enet), Lasso (lasso), Ridge (ridge) and the Linear Ensemble (lens). The nonlinear models comprise a feedforward neural network (nn), tree-based aggregation methods comprising an extremely randomized set of trees (xt) and a random forest (rf). nens is the nonlinear ensemble and ens is the ensemble across all models.

Table 15: Average and Median Half Spreads (Percent) of Insurance Firms' Bond Trade

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

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

Table 16: Trading Costs and Trade Size in the Corporate Bond Market.

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

This table regresses trading costs (the bond half-spread) on the log of trade size, various fixed effects and bond-level control variables. The bond level control variables include bond rating, coupon, maturity and the log of the issue size (amount outstanding). t-statistics in parentheses are based on two-way clustered standard errors at the day and client level.

Internet Appendix to

Factor Investing with 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`.

Bank of America Merrill Lynch (BAML) database

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

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

The following ICE-specific filters are then applied:

1. Only include corporate bonds, `Ind_Lvl_1 == 'corporate'`
2. Only include bonds issued by U.S. firms, `Country == 'US'`
3. Only include corporate bonds denominated in U.S. Dollars, `Currency == 'USD'`

BAML/ICE Bond Returns. Total bond returns are computed in a standard manner in ICE, and no assumptions about the timing of the last trading day of the month are made because the data is quote based, i.e., there is always a valid quote at month-end to compute a bond return. This means that each bond return is computed using a price quote at exactly the end of the month, each and every month. This introduces homogeneity into the bond returns because prices are sampled at exactly the same time each month. ICE only provides bid-side pricing, meaning bid-ask bias is inherently not present in the monthly sampled prices, returns and credit spreads. The monthly

ICE return variable is (as denoted in the original database), is `trr_mtd_loc`, which is the month-to-date return on the last business day of month t. We use this return specification (in excess of the one-month risk free rate of return) and the bond returns in excess of a portfolio duration matched U.S. Treasury bond returns denoted as `ex_rtn_mtd` in the ICE dataset as the dependent variables to train the machine learning models.

Enhanced TRACE Database

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

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

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

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

A.2 Bond and stock characteristics

We describe our 53 'custom-made' bond and equity characteristics in Table A.4. Panel A describes our 37 bond-based characteristics which span the vast majority of those used in the literature which forms corporate bond factors. Panel B describes additional equity-based characteristics that are not included in publicly available equity repositories, but have all been used in research that attempts to predict bond returns or form bond factors. All rank demeaned characteristics are made publicly available on `openbondassetpricing.com/data`.

For the publicly available equity databases, we follow Chen and Velikov (2023) and drop characteristics that are discrete (i.e., exchange indicators) or dominated by missing values at the stock level. When there are overlaps between the Chen and Zimmermann (2022), CZ, and Jensen et al. (2023), JKP characteristics we drop the CZ version of the characteristic. In total, after dropping characteristics based on the above, we are left with 137 CZ characteristics and 151 JKP characteristics.

B Net of Fees Corporate Bond Market Factor

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

$$R_{HYG,t} = \beta_0 + \beta_{BB} \cdot R_{BB,t} + \varepsilon_t,$$

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

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` in Table A.5.

C Machine Learning Model Estimation and Cross-Validation

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

We report the respective sets of hyperparameters which we cross-validate over in Table A.6.

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

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

Feed Forward Neural Network Ensemble Panel C reports the hyperparameters for the feed forward neural network (NN). We estimate a shallow network with a single layer and 32 neurons. Since our sample starts off with a relatively smaller sample size than that of Gu et al. (2020) and other work which utilizes equity data only, we set the batch size to 1024 (with batch normalization) and the number of epochs to 100. We cross-validate over the learning rate which is $\in [0.001, 0.01]$

and an ℓ_1 penalty $\in [0.001, 0.01]$. We also implement early stopping with the 'patience' parameter set to 5. The prediction variance of each individually estimated neural network is high. In order to reduce prediction variance across estimated neural network models, at each training date we estimate 10 models with different randomly assigned initial weights. In doing so, we select the best performing 5 models based on the smallest mean squared error estimated in the validation sample at that training date. This means, that at each date $t + 1$, we produce five predictions from the five best performing models estimated at the training date. The overall $t + 1$ prediction is the average over these five best performing models. At each training date, we then repeat this process ten times, yielding ten ensembled predictions. The final NN prediction for each month $t + 1$ is the average over these ten ensembled predictions, i.e., an ensemble over the ensemble.

D Additional Results

This section reports the additional tables referred to in the paper.

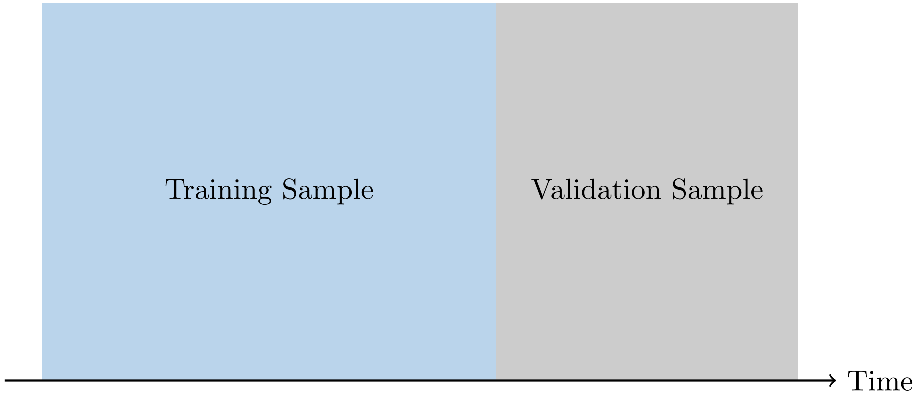

Figure A.1: Sample Splitting for Cross-Validation of Model Hyperparameters.

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

Table A.1: Logit Regression of Corporate Bond Transactions

		Customer Sell				Customer Buy			
		$100K		$2M		$100K		$2M	
Panel A. Regression Estimates									
log invcomp		0.270	(23.13)	0.158	(28.44)	0.247	(22.12)	0.163	(29.72)
inventory28		-0.260	(-45.46)	-0.382	(-80.93)	0.331	(60.68)	0.398	(91.88)
capratio		0.111	(12.27)	0.101	(20.10)	0.092	(10.55)	0.087	(17.08)
rating		0.679	(34.65)	0.611	(60.17)	0.657	(34.90)	0.620	(61.66)
maturity		-0.136	(-9.98)	-0.025	(-3.39)	-0.159	(-12.25)	0.006	(0.74)
log facevalue		1.587	(78.30)	1.239	(146.26)	1.598	(81.97)	1.222	(145.81)
coupon		-0.245	(-12.68)	0.023	(2.40)	-0.231	(-12.32)	0.049	(5.27)
age		-0.188	(-14.85)	-0.414	(-40.35)	-0.216	(-17.65)	-0.420	(-41.85)
Intercept		2.817	(150.49)	0.533	(76.25)	2.732	(152.55)	0.472	(67.33)
Pseudo R^2		0.230		0.221		0.236		0.219	
N		814,145		814,145		814,145		814,145	
Panel B. Model-Based Estimates of Trade Intensity and Delays									
Prob. Trade	data	0.877		0.588		0.868		0.577	
	cntr. factual	0.937		0.672		0.927		0.664	
λ	data	0.100		0.042		0.096		0.041	
	cntr. factual	0.132		0.053		0.125		0.052	
Exp. Delay	data	10.013		23.670		10.363		24.441	
(days)	cntr. factual	7.587		18.833		8.015		19.249	

This table reports the coefficient estimates for the logit regressions of the bond trading dummy on characteristics. Variable invcomp is the investor composition of Li and Yu (2023), which measures the activeness of bond investors at the end of the previous quarter. Variable inventory28 is the difference between customer buys and customer sells in the preceding 28 days, capratio is the intermediary capital ratio of He et al. (2017) on the previous month. The right-hand-side variables are standardized for ease of interpretation. The sample is at the month frequency, spanning August 2002 to November 2022.

Table A.2: Forecasting Regression of Corporate Bond Transactions: Nonlinear Model

		Customer Sell				Customer Buy			
		\$100K		\$2M		\$100K		\$2M	
Panel A. Regression Estimates									
invcomp		0.051	(11.58)	0.016	(2.95)	0.053	(11.80)	0.008	(1.45)
1/invcomp		-0.008	(-22.79)	-0.009	(-30.48)	-0.008	(-23.24)	-0.009	(-29.80)
invcomp^2		-0.003	(-9.25)	-0.001	(-2.63)	-0.003	(-9.51)	0.000	(-1.02)
log invcomp		-0.055	(-15.59)	-0.034	(-9.38)	-0.058	(-16.15)	-0.030	(-8.14)
inventory28		-0.001	(-37.70)	-0.005	(-63.79)	0.001	(29.28)	0.004	(55.46)
capratio		0.297	(4.77)	1.426	(23.00)	0.191	(3.05)	1.240	(19.74)
rating		0.014	(28.30)	0.028	(60.13)	0.014	(27.81)	0.028	(60.87)
maturity		-0.001	(-7.03)	0.000	(-1.64)	-0.002	(-9.29)	0.000	(2.27)
log facevalue		0.145	(66.09)	0.302	(153.29)	0.153	(67.79)	0.299	(142.95)
coupon		-0.013	(-16.91)	-0.002	(-2.27)	-0.014	(-16.70)	0.001	(0.69)
age		0.000	(-15.20)	0.000	(-41.07)	0.000	(-17.47)	0.000	(-42.37)
Intercept		-0.066	(-3.69)	-1.4888	(-89.20)	-0.11421	(-6.23)	-1.48129	(-84.95)
Pseudo R^2		0.148		0.257		0.157		0.255	
N		814,145		814,145		814,145		814,145	
Panel B. Model-Based Estimates of Trade Intensity and Delays									
Prob. Trade	data	0.877		0.588		0.868		0.576	
	cntr. factual	0.912		0.587		0.896		0.585	
λ	data	0.100		0.042		0.097		0.041	
	cntr. factual	0.116		0.042		0.108		0.042	
Exp. Delay	data	10.014		23.672		10.362		24.442	
(days)	cntr. factual	8.636		23.732		9.272		23.882	

This table reports the estimates of the logit regression of monthly bond transactions. The estimates follow that of Table A.1, except we add the nonlinear transformation of the investor composition measure of Li and Yu (2023).

Table A.3: Forecasting Regression of Corporate Bond Transactions: Interactions between Investor Compositions and Bond Characteristics

		Customer Sell				Customer Buy			
		\$100K		\$2M		\$100K		\$2M	
Panel A. Regression Estimates									
log invcomp		0.259	(28.51)	0.172	(17.96)	0.266	(27.91)	0.179	(16.91)
log invcomp rat		-0.004	(-12.16)	-0.002	(-6.50)	-0.003	(-11.58)	-0.002	(-5.16)
log invcomp tau		0.000	(1.59)	0.000	(2.39)	0.000	(2.80)	0.000	(1.93)
log invcomp fv		-0.037	(-29.88)	-0.016	(-12.19)	-0.038	(-29.29)	-0.016	(-11.38)
log invcomp cpn		0.002	(4.63)	-0.005	(-9.08)	0.002	(3.65)	-0.006	(-10.99)
log invcomp age		0.000	(3.33)	0.000	(1.69)	0.000	(2.27)	0.000	(2.11)
inventory28		-0.001	(-38.36)	-0.005	(-63.47)	0.001	(32.22)	0.004	(56.44)
capratio		0.064	(1.07)	1.170	(18.69)	-0.046	(-0.77)	0.987	(15.54)
rating		0.007	(11.26)	0.025	(33.60)	0.007	(11.04)	0.026	(34.89)
maturity		-0.001	(-4.99)	0.000	(-0.47)	-0.001	(-5.24)	0.000	(1.00)
log facevalue		0.074	(28.97)	0.276	(81.88)	0.080	(29.47)	0.273	(72.46)
coupon		-0.008	(-7.54)	-0.011	(-8.38)	-0.009	(-8.18)	-0.011	(-8.32)
age		0.000	(-8.26)	0.000	(-23.36)	0.000	(-10.37)	0.000	(-24.74)
Intercept		0.484	(25.79)	-1.195	(-47.03)	0.454	(22.61)	-1.185	(-41.62)
Pseudo R^2		0.148		0.252		0.157		0.250	
N		814,145		814,145		814,145		814,145	
Panel B. Model-Based Estimates of Trade Intensity and Delays									
Prob. Trade	data	0.877		0.588		0.868		0.576	
	cntr. factual	0.959		0.731		0.948		0.725	
λ	data	0.100		0.042		0.097		0.041	
	cntr. factual	0.152		0.062		0.141		0.062	
Exp. Delay	data	10.014		23.672		10.362		24.442	
(days)	cntr. factual	6.583		16.014		7.106		16.247	

This table reports the estimates of the logit regression of monthly bond transactions. The estimates follow that of Table A.1, except we add the interaction terms between the investor composition measure of Li and Yu (2023) and bond characteristics such as credit rating (rat), time to maturity (tau), log face value (fv), coupon rate (cpn), and time since the issuance (age).

Table A.4: List of the Corporate Bond and Stock Characteristics.

Num.	ID	Characteristic Name and Description	Reference	Source
		Panel A: Bond Characteristics Available on `openbondassetpricing.com/data`		
1	age	Bond age. The number of years the bond has been in issuance scaled by the tenor of the bond.	Israel et al. (2018)	BAML/ICE
2	ave12mspread	Rolling 12-month moving average of bond option adjusted credit spreads skipping the prior month	Elkamhi et al. (2021)	BAML/ICE
3	bond_mom_ipr	6-month bond credit momentum skipping the prior month. Demeaned with duration-times-spread.	Israel et al. (2018)	BAML/ICE
4	bond_price	Bond price. Clean price of the bond.	Bartram et al. (2023)	BAML/ICE
5	convexity	Bond convexity. Convexity of the bond.	–	BAML/ICE
6	coupon	Bond coupon. The annualised bond coupon payment in percent (%)	Chung et al. (2019)	BAML/ICE
7	dspread	First difference in bond option adjusted credit spread.	–	BAML/ICE
8	dts	Duration-times-spread. Annualized bond duration multiplied by the bond option adjusted credit spread.	Dor et al. (2007)	BAML/ICE
9	duration	Bond duration. The derivative of the bond value to the credit spread divided by the bond value, and is calculated by ICE	Israel et al. (2018)	BAML/ICE
10	emp_value	Empirical bond value. Defined as the percentage difference between the actual credit spread and the fitted ("fair") credit spread for each bond. The fitted spread is derived from cross-sectional regressions of the log of bond option adjusted credit spreads onto the log of duration, bond ratings and bond credit return volatility. Demeaned with duration-times-spread.	Israel et al. (2018)	BAML/ICE
11	faceval	Face value. The bond amount outstanding in units	Israel et al. (2018)	BAML/ICE
12	impliedspread	The fitted spread used to estimate the 33. value characteristic.	Houweling and Van Zundert (2017)	BAML/ICE
13	kurtosis	Bond kurtosis. Rolling bond excess kurtosis computed with a minimum amount of rolling observations equalling 12 which then expands up to 60-months	–	BAML/ICE
14	ltrev30_6	Bond long-term reversal (medium-term). Computed as the rolling sum of the prior 30-months of bond returns (minimum 12 monthly values) minus the rolling sum of the most recent 6-month returns (minimum 6 monthly values).	Subrahmanyam (2023)	BAML/ICE
15	ltrev48_12	Bond long-term reversal (long-term). Computed as the rolling sum of the prior 48-months of bond returns (minimum 12 monthly values) minus the rolling sum of the most recent 12-month returns (minimum 6 monthly values).	Bali et al. (2021)	BAML/ICE
16	mom3mspread	Mom. 3m log(Spread). The log of the spread 3 months earlier minus current log spread	–	BAML/ICE
17	mom6	Corporate bond momentum. The sum of the last 6-months of bond returns minus the prior month	Gebhardt et al. (2005)	BAML/ICE
18	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
19	mom6mspread	Mom. 6m log(Spread). The log of the spread 6 months earlier minus current log spread	–	BAML/ICE
20	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
21	rating	Bond S&P rating. Bond numerical rating. AAA = 1, ... , D = 22	Kelly et al. (2021)	BAML/ICE
22	ratingxspread	Corporate bond ratings multiplied by credit spread.	–	BAML/ICE
23	sizeb	Bond market capitalization. Computed as bond units amount outstanding multiplied by the clean price of the bond	Houweling and Van Zundert (2017)	BAML/ICE
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)	BAML/ICE

25	sprtod2d	Spread-to-Distance-to-Default. Spread-to-D2D is the option-adjusted spread, divided by one minus the CDF of the distance-to-default	Kelly et al. (2021)	CRSP/COMP
26	spread	Bond option adjusted credit spread. The option adjusted spread of the bond provided by ICE	Kelly et al. (2021)	BAML/ICE
27	spreadvol	Volatility of the first difference of the bond option adjusted credit spread. Rolling period of 24-months with a minimum required observations of 12	–	BAML/ICE
28	strevb	Short-term bond reversal. Defined as the previous months bond return	–	BAML/ICE
29	struc_value	Structural bond value. Defined as the percentage difference between the actual credit spread and the fitted ("fair") credit spread for each bond. The fitted spread is derived from cross-sectional regressions of the log of bond option adjusted credit spreads onto the log of the probability of default computed with firm-level distance-to-default. Demeaned with duration-times-spread.	Israel et al. (2018)	BAML/ICE
30	swap_spread	Bond swap spread. The swap spread of the bond provided by ICE	Kelly et al. (2021)	BAML/ICE
31	tmt	Bond time to maturity	–	BAML/ICE
32	value	Bond value. Defined as the percentage difference between the actual credit spread and the fitted ("fair") credit spread for each bond. The fitted spread is derived from cross-sectional regressions of the log of bond option adjusted credit spreads onto the 3-month change in spreads, maturity and credit ratings.	Houweling and Van Zundert (2017)	BAML/ICE
33	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
34	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
35	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
36	yield	Bond yield. The yield-to-maturity of the bond provided by ICE	Gebhardt et al. (2005)	BAML/ICE
37	yld_to_worst	Bond yield-to-worst. The yield-to-worst of the bond provided by ICE	–	BAML/ICE
Panel B: Equity Characteristics Available on `openbondassetpricing.com/data`				
38	book_lev	Book leverage. Shareholder's equity and long-/short-term debt (`DLTTQ` + `DLCQ`) and minority interest (`MIBTQ`) minus cash and inventories (`CHEQ`), divided by shareholder's equity minus preferred stock	Kelly et al. (2021)	COMP
39	book_prc	Book-to-price. Firm Book-to-price is the sum of shareholder's equity and preferred stock divided by equity market capitalization for the issuing firm	Kelly et al. (2021)	CRSP/COMP
40	chg_gp_at	Profitability change. The 5-year change in gross profitability	Asness et al. (2019)	COMP
41	d2d	Distance-to-default. Computed as in Bharath and Shumway (2008)	Bharath and Shumway (2008)	CRSP/COMP
42	debt_ebitda	Debt-to-EBITDA. Total debt (`DLTTQ` + `DLCQ`) divided by EBITDA (`SALEQ` − `COGSQ` − `XSGAQ`)	Kelly et al. (2021)	CRSP/COMP
43	eqtyvol	Equity volatility defined as the month-end value from a 180-day rolling-period	Campbell and Taksler (2003)	CRSP
44	gp_at	Profitability. Sales (`REVTQ`) minus cost-of-goods-sold (`COGSQ`), divided by assets (`ATQ`)	Choi and Kim (2018)	COMP
45	gross_prof_ipr	Demeaned profitability. Duration-times-spread demeaned gp_at	Israel et al. (2018)	COMP/ICE
46	me	Equity market capitalization	Choi and Kim (2018)	CRSP
47	mkt_lev	Market leverage. Market capitalization and long-/short-term debt (`DLTTQ` + `DLCQ`) and minority interest (`MIBTQ`) and preferred stock minus cash and inventories (`CHEQ`), divided by market capitalization	Kelly et al. (2021)	CRSP/COMP

48	mkt_lev_ipr	Demeaned Market leverage. Duration-times-spread demeaned mkt_lev	Israel et al. (2018)	COMP/ICE
49	ni_me	Earnings-to-price. Net income (`NIQ`) divided by market equity	Correia et al. (2012)	CRSP/COMP
50	operlvg	Operating leverage. Sales (`SALEQ`) minus EBITDA (`SALEQ` − `COGSQ` − `XSGAQ`), divided by EBITDA	Gamba and Saretto (2013)	COMP
51	stock_mom_ipr	Demeaned equity momentum. Duration-times-spread demeaned momentum (sum of the past 6 months, skipping the most recent month).	Israel et al. (2018)	COMP/ICE
52	totaldebt	Total firm debt (`DLTTQ` + `DLCQ`)	Kelly et al. (2021)	COMP
53	turnvol	Turnover volatility. Turnover volatility is the quarterly standard deviation of sales (`SALEQ`) divided by assets (`ATQ`). The volatility is computed over 80 quarters, with a minimum required period of 10 quarters. Thereafter, the volatility is averaged (smoothed) over the preceding 4-quarters in a rolling fashion	Kelly et al. (2021)	CRSP/COMP

This table presents information on the 53 characteristics that are generated by the authors and which are made publicly available on `openbondassetpricing.com/data`. The remaining characteristics are sourced from the publicly available equity data repositories build by Chen and Zimmermann (2022) and Jensen et al. (2023). Panel A reports the 37 characteristics which relate to *bond*-only characteristics which are constructed using the BAML/ICE corporate bond dataset. Panel B reports the 16 characteristics which relate to *equity-and-bond* characteristics which are constructed using CRSP and COMPUSTAT (COMP). Additional resources and description notes for the equity `openassetpricing.com`-based data can be downloaded here. Documentation for the Jensen et al. (2023) data can be found here.

Table A.5: 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.

Table A.6: Hyperparameters Across the Machine Learning Models.

Panel A: Linear Models with Penalties: LASSO, RIDGE & ENET

Parameter	`sklearn` mnemonic	Value
Intercept	`fit_intercept=True`	True
ℓ_1 penalty	`alphas`	Variable
ℓ_2 penalty	`alphas`	$\in [0.0001, \ldots, 1]$
Num. Penalties	`n_alphas`	100
ℓ_1 ratio	`l1_ratio`	$\in$ [0.001, 0.01,0.99,0.999]

Panel B: Tree-Based Ensembles: RF and XT

Parameter	`sklearn` mnemonic	Value
Num. Trees	`n_estimators`	100
Max depth	`max_depth`	$\in$ [2,4,6]
Split features	`max_features`	$\in$ [5,10,20]
Min leaf samples	`min_samples_leaf`	$\in$ [1,10,50]

Panel C: Feed Forward Neural Network: NN

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

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