

Corporate Bond Multipliers: Substitutes Matter *

Manav Chaudhary[†] Zhiyu Fu[‡] Jian Li[§]

August 2023

Abstract

Many economic questions require estimating the price impact of demand shifts in the bond market. In spite of corporate bonds having salient characteristics that distinguish close versus distant substitutes, existing estimates of corporate bond multipliers (the price increase in response to demand shifts) typically assume that all bonds, regardless of their characteristics, are equally good substitutes. In this paper, we show that accounting for the heterogeneous substitutability between bonds is critical for estimating multipliers. By allowing rich heterogeneous substitution patterns among bonds, we demonstrate that security-level multipliers are an order of magnitude smaller than previously estimated and are essentially zero. Nonetheless, aggregated portfolios exhibit substantially larger multipliers, reflecting the reduced availability of near substitutes for more aggregated portfolios. Furthermore, we find that the price impact reverts after a quarter, and that the multiplier is larger for high-yield bonds, longer-maturity bonds, and bonds with greater arbitrage risks.

Keywords: Corporate bonds, inelastic demand, mutual funds, demand-based asset pricing

JEL codes: G10, G12, G23

*We thank Brendan Berthold, Anna Cieslak, Kent Daniel, Olivier Darmouni, Valentin Haddad, Lars Peter Hansen, John Heaton, Jean Helwege, Ralph Koijen, Erik Loualiche, Stefan Nagel, Paul Tetlock and seminar participants at the Economic Dynamics and Financial Markets Working Group at the University of Chicago, AFFECT 2023 mentorship workshop, Columbia Business School, David Backus Memorial Conference on Macro-Finance and NBER Summer Institute for helpful comments. This research was funded in part by the Fama-Miller Center for Research in Finance at the University of Chicago Booth School of Business.

[†]Department of Economics and Booth School of Business, University of Chicago. Email: mchaudhary@chicagobooth.edu

[‡]Department of Economics and Booth School of Business, University of Chicago. Email: zhiyufu@uchicago.edu

[§]Columbia Business School. Email: jl5964@columbia.edu

1 Introduction

Our paper seeks to understand the ability of the corporate bond market to absorb demand-driven price pressures. This is critical for a wide range of policy questions, for example, the impact of quantitative easing or tightening, the equilibrium implications of regulations on insurance companies, and market stability concerns related to mutual fund runs (e.g., [Vissing-Jorgensen and Krishnamurthy 2011](#); [Kojien et al. 2017](#); [Goldstein et al. 2017](#)).¹ Thus far, the literature has used both reduced-form and structural methods to provide quantitative answers to such questions (e.g., [Ellul et al. 2011](#); [Manconi et al. 2012](#); [Bretscher et al. 2022](#)). However, when estimating price impacts at the security level, existing methods often treat all other securities as equally good substitutes and overlook the heterogeneity in substitutability among assets. This is at odds with reality, since Apple’s 10-year bond is presumably a good substitute for Google’s 10-year bond, as both bonds are AA+ rated, whereas it is a poor substitute for bonds issued by Bed Bath & Beyond, which are in default.

Textbook theory suggests that the impact of demand-driven price pressure crucially depends on the availability of close substitutes. On the one hand, assets that have close substitutes should have smaller price impacts because arbitrageurs can easily hedge risks and aggressively trade against idiosyncratic demand shocks. On the other hand, the same mechanism leads the demand-driven price impact to spill over across close substitutes.

In this paper, we focus on the corporate bond market to highlight the importance of accounting for the correct substitutes when estimating price impacts. Perhaps unique to corporate bonds, salient characteristics such as credit rating and maturity make identifying close substitutes relatively straightforward. Bonds within the same rating and maturity category are closer substitutes than bonds that do not share these characteristics. We show that accounting for heterogeneous substitution patterns reduces the price impact estimated at the security level by an order of magnitude. However, shocks at the portfolio level still induce significant price movements, and the price impact increases monotonically with the aggregation level.

Specifically, we measure demand-driven price pressure using multipliers—the percentage increase in the asset’s price if mutual funds exogenously increased their demand for 1% of the

¹The bond market disruption in March 2020 has raised concerns regarding whether the corporate bond market has become more vulnerable to demand shocks. See [Hubbard et al. \(2021\)](#).

asset's amount outstanding (Gabaix and Koijen, 2021). If we assume homogeneous substitution among all bonds, we get a positive significant multiplier around 0.4 at the security level, consistent with the magnitudes found in existing studies (Bretscher et al., 2022; Siani, 2022b; Darmouni et al., 2023). However, once we allow bonds to have a different cross-elasticity for bonds with similar versus different characteristics, our estimate for the multiplier drops to 0.05 and becomes insignificant. Our results show that ignoring heterogeneous substitution patterns leads to a considerable upward bias.

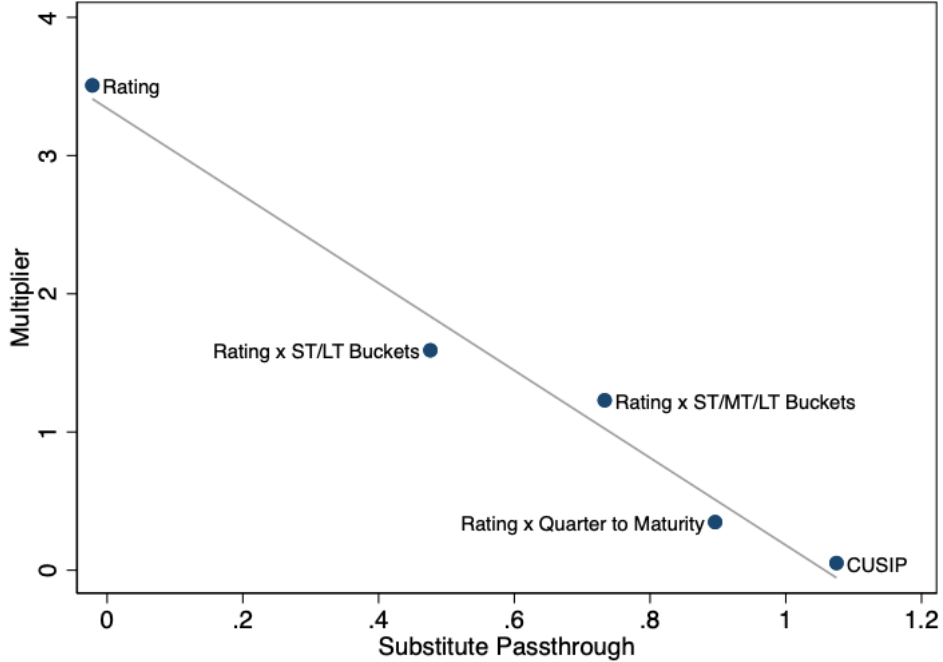
In addition, we estimate the “substitute passthrough” of close substitutes, which is defined as the percentage increase in an asset's price due to its close substitutes' prices rising by 1%. This parameter is important for understanding how shocks are transmitted across assets. At the security level, we find the substitute passthrough is around 1, implying a significant spillover effect among close substitutes.

Our findings suggest that individual bonds are highly substitutable and that the market is quite good at absorbing security-specific demand shocks. With that said, the market is not as good at absorbing demand shocks at more aggregated levels — for example, demand shocks for a specific-rating category. As Figure 1 shows, the more aggregated the demand shock, the higher the multiplier and the lower the substitute passthrough. The considerable difference in the price impact between the security level and the more aggregated portfolio level highlights the double-edged nature of substitutability: while high substitutability reduces the price impact to direct demand shocks, it increases the spillover effect from indirect demand shocks to substitutes. This result also speaks to the optimal implementation of quantitative easing policies in order to achieve the maximum price impact.

Finally, we investigate the persistence of the price impact and its heterogeneity along various portfolio characteristics. We find that the price impact reverts after a quarter and that the price impact is larger for high-yield bonds, longer-maturity bonds, and bonds with greater arbitrage risks.

Our results have important implications for the growing literature estimating price impacts using both reduced-form methods and structural models. Recent work using reduced-form shocks has increasingly added more fixed effects for identification (e.g., Choi et al. 2020; Coppola 2022; Holm-Hadulla and Leombroni 2023). Our framework clarifies that the specifications with saturated characteristic-time fixed effects are fundamentally estimating a different parameter compared to the specifications without saturated fixed effects. The

Figure 1: Demand is more inelastic for more aggregated portfolios



The figure plots the multipliers (inverse of elasticity) and the substitute passthrough coefficients (the link between substitute and test asset returns) for various levels of portfolio aggregation; see Table 4 for descriptions of the aggregation. The figure shows a negative relationship between multipliers and substitute passthroughs. It also shows multipliers are monotonically increasing in aggregation, whereas the substitute passthrough coefficients are monotonically decreasing in aggregation.

former estimates the price impact at the individual security level, controlling for close substitutes' prices, whereas the estimates from the latter specification are potentially averages between portfolio-level and security-level price impacts. A second strand of the literature estimates bond demand elasticities using structural models (e.g., Bretscher et al. 2022; Fang 2022; Darmouni et al. 2023). Our findings suggest that using nested logit models may be particularly important when there is strong heterogeneity in substitutability across assets. Alternatively, one can model the cross-elasticity explicitly as a function of the similarity of two assets.

We construct exogenous demand shocks from flow-induced trading by mutual funds.

Following the literature, we assume that mutual funds invest their flows into their existing portfolios proportionally (Lou, 2012).² We allow for autocorrelations in fund flows and extract the flow innovation terms. Furthermore, we adjust the demand shock to account for the expected cumulative trading from this flow innovation.³ We then strip out common factor variations to remove potential confounding factors. Finally, we construct bond-level shocks by summing up fund-level shocks weighted by the fund's lagged holding share for each bond.

We then use these demand shocks to estimate price multipliers.⁴ We show that as long as lagged mutual fund holding shares are orthogonal to unobserved demand shocks, we can use our (observed) demand shocks to identify the multipliers, even if the flows are endogenous. Our identification strategy is an application of the Bartik instrument, where exogeneity of shares guarantees instrument validity (Goldsmith-Pinkham et al., 2020; Chaudhry, 2022). From this perspective, our identification strategy can be viewed as pooling multiple exogenous exposure research designs. Each bond is exposed to flows into a given mutual fund, but the degree of exposure is governed by how much this mutual fund held the bond in the previous period, which is plausibly exogenous conditional on controls.

To estimate multipliers for individual bonds, we regress bond returns on the constructed demand shocks, controlling for the relevant substitute portfolio's returns. Under the homogeneous cross-elasticity assumption (i.e., any given pair of bonds have the same cross-elasticity), we just need to control for the market return, or equivalently, a time fixed effect. The estimated multiplier is significantly positive and around 0.4, similar to the estimate in the literature.

However, once we relax the homogeneous elasticity assumption, we find a much smaller multiplier. To allow for heterogeneous cross-elasticity, we add another control, which is the returns of an additional substitute portfolio formed by bonds with the same detailed rating category as the testing asset.⁵ To deal with the endogeneity of prices, we instrument

²In Appendix C.1, we relax this one-to-one passthrough assumption by directly looking at mutual fund rebalancing to estimate lower bounds for the demand passthrough and use this to find upper bounds for the multiplier. This exercise suggests that the security-level multiplier is at most 0.1, which is still 3-6 times smaller than what is typically found in studies that assume homogeneous cross-elasticities.

³This is the theoretically relevant shock when there are forward-looking investors (Gabaix and Koijen, 2021).

⁴We focus on the multipliers in this paper, but one can easily translate the estimates into the elasticity space by taking the inverse of the multipliers.

⁵We refer to AAA, AA+, AA – AA + ... as detailed rating categories and AAA, AA, A, BBB... as coarse

the return of the substitute portfolio with the demand shocks to bonds in the substitute portfolio. Under this specification, we find the multiplier is around 0.05 and statistically indistinguishable from zero. In other words, the price impact of these demand shocks from mutual funds is close to zero. On the other hand, the substitute passthrough is close to one (i.e., the spillover effect among close substitutes is almost one to one). We also run a specification where we directly control for substitute prices using substitute portfolio-time fixed effects rather than using instrumental variables, and the multiplier estimates are essentially unchanged.

The reason that ignoring close-substitutes leads to significant over-estimation of the multiplier boils down to the fact that demand shocks for similar bonds are often correlated. When a bond is hit by a positive demand shock, its close substitutes likely also receive similar demand shocks because of common ownership. The increase in the bond price comes from both its own demand shock and also its substitutes' prices being higher. If the latter channel is not accounted for correctly, then the degree of price change due to its own demand shock would be mismeasured. Furthermore, the spillover effect from the asset's substitute price is larger if it is a close substitute. Assuming homogeneous substitution effectively underweighs the impact from close substitutes and attributes too much of the price movement to the asset's own demand shocks.⁶

We then move on to estimate the multipliers for corporate bond portfolios. In contrast to the near-zero multiplier of individual bonds, we find that portfolios have significantly positive multipliers. As the portfolio multiplier is monotonically increasing in the degree of aggregation, the substitute passthrough is monotonically decreasing. For portfolios formed by bonds in the same rating category with the same quarter to maturity, the multiplier is 0.35, significantly larger than the 0.05 estimate at the bond level. Furthermore, the portfolio substitute passthrough is estimated to be slightly lower than that at the bond level. At the most aggregated level, where each portfolio is formed by bonds in the same coarse rating group, the portfolio multiplier is 3.5 and the passthrough coefficient drops to 0.

Consistent with demand being more elastic in the long run, we find that the price response to demand shocks is not permanent and seems to revert fully. For our baseline level of

rating categories.

⁶Accounting for heterogeneous substitution is particularly important for the corporate bond market relative to the equity market. Controlling for close substitutes does reduce the multiplier estimated in the equity market, but only slightly compared to the case for bonds.

aggregation, we find that the price impact takes around two quarters to revert fully. This reversion rate is somewhat faster than the four-quarter reversion Li (2021) found for stock portfolios. Overall, our results suggest that demand is much more elastic in the long run.

Furthermore, we investigate heterogeneity in multipliers to uncover potential structural drivers of inelastic demand. We focus on portfolios formed by bonds in the same detailed rating category with the same quarter to maturity. First, we explore heterogeneity along credit risk and duration risk dimensions and find that portfolios consisting of bonds with lower ratings and longer maturities have significantly larger multipliers. Investment-grade (IG) portfolios have near-zero multipliers, whereas high-yield (HY) portfolios have multipliers around 0.6. In other words, our baseline-level multiplier estimate of 0.35 is predominantly driven by HY bonds. Specifically, the sharp change in multipliers around the IG/HY cutoff indicates that elasticity depends on investor clientele.

Finally, we compute the Sharpe ratio of a strategy that takes advantage of the price deviation between the testing asset and its substitute portfolio for one quarter. This is a near-arbitrage strategy as prices may not converge in the following quarter as a result of non-flow shocks and future flow shocks. We find the (annualized) Sharpe ratios are generally small, suggesting unhedged risk may be a factor impeding arbitrage activities. Indeed, we find that in the cross section of portfolios formed by detailed rating and maturity, portfolios with high arbitrage risks have significantly higher multipliers. A one-standard-deviation increase in arbitrage risk increases the multiplier by 0.45.

In the remainder of this section, we discuss the literature. We explain our demand structure and estimation methodology in Section 2. Section 3 describes our data and construction of demand shocks in detail. Section 4 presents the baseline estimation results, and Section 5 focuses on the implications of the multiplier estimates. Section 6 concludes.

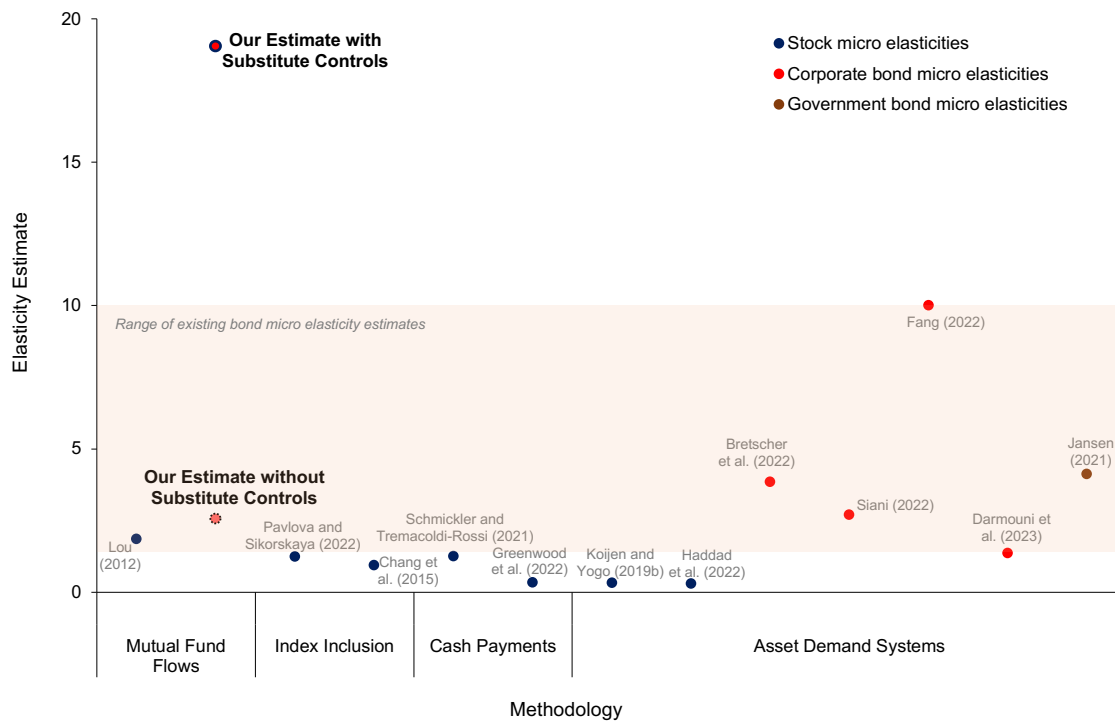
1.1 Literature Review

Estimating multipliers directly contributes to the large literature in asset pricing that estimates the slope of the asset demand curve; multipliers are the inverse of demand elasticities.⁷ The literature has shown that asset demand is considerably less elastic than standard theories suggest. Partially because of data availability, this literature has traditionally focused

⁷Our security-level multiplier estimate of 0.05 implies an elasticity of $1/0.05 = 20$.

on estimating the demand elasticity of individual stocks (micro elasticities). More recently, researchers have begun estimating the micro elasticity of other asset classes, such as bonds, and more aggregate-level elasticities, such as those of the entire equity market (macro elasticities). Our paper contributes to this broad body of work. Figure 2 shows how our estimates compare to a representative selection of micro elasticity estimates from the literature. Our estimates confirm that bond demand is more elastic than stocks, but additionally we find that bonds are a lot more elastic than suggested by prior studies.

Figure 2: Our estimate versus other micro stock and bond elasticity estimates



This graph presents a representative set of stock and bond micro elasticity estimates from the literature. The y -axis is the magnitude of the estimate; if authors provided a range for the estimate, the graph shows the midpoint of the range. The x -axis describes the methodologies used for estimating the elasticities.

Most directly, our estimates speak to the burgeoning literature estimating corporate bond micro elasticities. Several recent papers have adapted the asset demand system developed by

Koijen and Yogo (2019b) for the equity market to study corporate bond demand.⁸ Bretscher et al. (2022) focus on the heterogeneous demand elasticities of different investors and find the average corporate bond elasticity is around 3.8. Darmouni et al. (2023) adopt a two-layer framework to capture the interaction effect between fund flows and asset market inelasticity. Their demand elasticity estimate ranges from 0.8 to around 2, depending on the time period. Siani (2022a) estimates the elasticity of the primary market demand and finds it is between 1.9 to 3.5. Fang (2022) uses a nested logit structure to allow for more flexible substitution within and across IG and HY bonds and finds a higher elasticity of around 10. Compared to these studies, we find corporate bond markets are considerably more elastic, with a bond-level micro elasticity of around 20. Our method has several advantages over the demand system estimates. First, our approach allows for greater flexibility in substitution patterns. At the investor level, a standard logit demand structure imposes homogeneous cross-substitution across assets (in a nested logit system, homogeneity is within the nest). In the context of corporate bonds, this is a strong assumption. For example, it assumes AA- and CCC bonds are equally good substitutes for AA bonds. In line with findings in Fang (2022), allowing for greater flexibility in substitution results in more elastic estimates for the slope of an individual bond's demand curve. Our approach is relatively model agnostic compared to asset demand systems. Furthermore, our approach is not encumbered by not observing the holdings of investors such as households, hedge funds, private pension funds, and so on. Because of data availability issues, corporate bond demand systems are typically only able to estimate the demands for insurance companies, bond mutual funds, and public pension funds. Hence, the bottom-up aggregated market demand elasticity omits potentially important investors. Since our approach starts with the market clearing condition, it is not affected by these data availability concerns. Finally, our top-down approach makes it more suitable for studying bond-specific heterogeneity in micro elasticities, for example, resulting from rating-based segmentation.

In terms of stock micro elasticities, an early strand of the literature shows that the inclusion of stocks in indices results in significant positive abnormal returns. For example, see Shleifer (1986); Harris and Gurel (1986); Beneish and Whaley (1996); Wurgler and Zhu-

⁸Jansen (2021) has also used demand systems to estimate the elasticity of European government bonds and finds estimates of around 4.11 for maturity bucket portfolios, and for other the papers applying the method to equity markets see Haddad et al. (2022)

ravskaya (2002); Chen et al. (2004); Chang et al. (2015); Pavlova and Sikorskaya (2022) for stocks and Calomiris et al. (2022) for bonds. Generally, they find a significant and permanent price impact; the estimated micro elasticities are around three orders of magnitude smaller than those implied by standard models (Petajisto, 2009). Among these index-inclusion papers, Wurgler and Zhuravskaya (2002) have the closest connection with our paper. They show that demand is more elastic for stocks with closer substitutes. This finding is consistent with arbitrageurs being able to better hedge stock-specific idiosyncratic risk and trade more aggressively against demand shocks. We show this is also true for corporate bonds; those with closer substitutes exhibit more elastic demand curves. Overall, our demand shock and setting of corporate bonds have several advantages over Wurgler and Zhuravskaya (2002) for establishing the importance of close substitutes in flattening the demand curve. First, stocks lack a prominent characteristic for identifying substitutes, whereas, in the case of corporate bonds, credit rating and maturity provide natural dimensions for identifying close substitutes. Second, index inclusion (or deletion) only affects a handful of assets in any given year, whereas flow-induced trading shocks have an impact on all assets to some extent. The broader exposure allows us to conduct our analysis at the bond portfolio level. At this higher level of aggregation, arbitrage risks are likely to play a more significant role in limiting arbitrage and hence are easier to detect.⁹

Another strand of the stock elasticity literature uses mechanical portfolio rebalancing as a source of demand shocks, showing that mechanical trading by mutual funds in response to flows has a significant price impact (Lou, 2012; Coval and Stafford, 2007; Jotikasthira et al., 2012; Edmans et al., 2012; Li, 2021). In contrast, Choi et al. (2020) find little evidence for fire-sale price pressure in the corporate bond market. Other variants of mechanical-trading induced demand shocks include using dividend-payment induced trading (Hartzmark and Solomon, 2022; Schmickler and Tremacoldi-Rossi, 2022), reinvestment of stimulus payments by US households (Greenwood et al., 2022), and investing refunds from unsuccessful bids in Chinese IPOs (Li et al., 2021). Overall, these papers find a large price impact as well. Li and Lin (2022) use equity net flows and show stock prices respond more to demands at more aggregated levels. Our identification strategy directly builds on the mutual fund flow-induced trading literature. We make three contributions compared to the standard methodology.

⁹Inclusion in an index may bring better liquidity and better coverage by analysts, which may drive the observed price impact. Trade-induced shocks are less contaminated by such effects.

First, we take into account the close substitute portfolio's return in the estimation; this helps address the omitted variable bias generated by flow-induced demand shocks being correlated across assets. Second, we use the method proposed by [Gabaix and Koijen \(2021\)](#) to control for the omitted variable bias generated by forward-looking agents anticipating predictable fund flows. Third, we formalize the identification by appealing to shares' exogeneity rather than shocks' exogeneity. This formalization of the identification strategy builds a direct parallel with an exogenous share identification of Bartik instruments ([Goldsmith-Pinkham et al., 2020](#); [Chaudhry, 2022](#)).

At more aggregated levels, [Gabaix and Koijen \(2021\)](#) find that the macro elasticity of the overall equity market demand is around 0.2. [Hartzmark and Solomon \(2022\)](#) and [Li et al. \(2021\)](#) also find relatively inelastic macro elasticities. Overall, the equity market seems considerably more inelastic than individual stocks. Other papers have attempted to estimate equity portfolio-level elasticities and found them to be somewhere in between equity micro and macro elasticities ([Peng and Wang, 2019](#); [Li, 2021](#)). We show that more aggregated portfolios of individual bonds are more inelastic, and they are also less sensitive to changes in the prices of their close substitutes.

2 Demand Framework

In this section, we outline our demand framework. To help frame the discussion, we begin by introducing a fully general demand system that, while infeasible to estimate, allows for complete flexibility in the cross-elasticities between bonds. We then outline the homogeneous cross-elasticity restriction typically made by existing methods. This restriction makes estimation feasible; however, it risks introducing positive omitted variable bias to the multiplier estimates. Finally, we outline our demand system, which relaxes the homogeneous cross-elasticity restriction by allowing for *heterogeneous* cross-elasticity between close and distant substitutes, where close and distant substitutes are identified using bond characteristics such as rating and maturity.¹⁰ Our demand system brings us one step closer to the general demand system but remains feasible for estimation.

¹⁰In Appendix F, we show how our approach links to logit demand systems and nested logit demand systems, as in [Koijen and Yogo \(2019b\)](#) and [Koijen and Yogo \(2019a\)](#).

2.1 Fully general demand system

For a fund i in our sample (denote the set as MF), we assume its demand for the N available assets is

$$\underbrace{\mathbf{q}_{i,t}}_{N \times 1} = \underbrace{\Gamma}_{N \times N} \underbrace{\mathbf{p}_t}_{N \times 1} + \underbrace{\mathbf{u}_{i,t}}_{N \times 1} + \underbrace{\nu_{i,t}}_{N \times 1} \quad \text{for } i \in MF, \quad (1)$$

where $\mathbf{q}_{i,t}$ is an $N \times 1$ by vector, where element $q_{i,j,t}$ denotes the log quantity for bond j demanded by fund i . Similarly, \mathbf{p}_t is the $N \times 1$ price vector, where element p_j denotes the log price of bond j . The term Γ is an $N \times N$ matrix that governs demand elasticities and substitution patterns among bonds. To simplify the notation, we assume Γ is homogeneous among investors, but this does not matter for our estimation.¹¹ The slope of the demand curve for asset j is $\frac{\partial q_j}{\partial p_j} = \Gamma_{j,j}$. Furthermore, $\frac{\partial q_j}{\partial p_k} = \Gamma_{j,k}$ is the cross-elasticity of asset j to asset k 's price. The observed demand shock from flow-induced trading is captured by $\mathbf{u}_{i,t}$, where element $u_{i,j,t}$ denotes fund i 's exogenous demand for bond j due to flows. Finally, the vector $\nu_{i,t}$ captures the unobserved demand shocks to fund i at time t . For now, we assume that the observed demand shocks $\mathbf{u}_{i,t}$ are orthogonal to the unobserved shocks $\nu_{i,t}$. We explain how we construct this shock and the identification assumption that guarantees its orthogonality in the next section.

Equation (1) can be viewed as a log-linearization of any generic demand function. For example, it nests the standard demand function derived from maximizing mean-variance utility. However, deriving the demand function from utility maximization would inevitably impose some restrictions on investors' utility. To preserve the maximum degree of flexibility, we model the demand in a reduced-form way.

For a fund that is not in our sample,¹² and thus we cannot observe its demand shocks, we denote the demand of this fund as

$$\mathbf{q}_{i,t} = \Gamma \mathbf{p}_t + \nu_{i,t} \quad \text{for } i \notin MF. \quad (2)$$

¹¹If Γ is heterogeneous among investors, then the final expression for market multipliers should be the average of investor elasticities weighted by their assets under management (AUM) shares.

¹²We use the term "fund" to refer to all investors in the bond market.

With the constant net supply of bonds, the market clearing condition gives

$$\Delta \left(\sum_i \exp(\mathbf{q}_{i,t}) \right) = 0. \quad (3)$$

Log-linearizing the market clearing condition around the last period values, we get for bond j

$$\Delta p_{j,t} = M u_{j,t} + \tilde{M} \Delta p_{j,t}^{sub} + \tilde{\nu}_{j,t}, \quad (4)$$

where

$$\begin{aligned} u_{j,t} &\equiv \sum_i S_{i,j,t-1} u_{i,t} \equiv \sum_i \frac{\exp(q_{i,j,t-1})}{\sum_i \exp(q_{i,j,t-1})} u_{i,t} \\ M &\equiv -\frac{1}{\Gamma_{j,j}} \\ \tilde{M} &\equiv -\frac{\sum_{k \neq j} \Gamma_{j,k}}{\Gamma_{j,j}} \\ \Delta p_{j,t}^{sub} &\equiv \sum_{k \neq j} w_k^j \Delta p_{k,t} \equiv \sum_{k \neq j} \frac{\Gamma_{j,k}}{\sum_{k \neq j} \Gamma_{j,k}} \Delta p_{k,t} \end{aligned}$$

and $\tilde{\nu}_{j,t} \equiv \nu_{j,t}/(-\Gamma_{j,j})$, where $\nu_{j,t}$ is the market-share weighted average of demand from the fund-level unobserved demand shocks $\nu_{i,t}$. With a slight abuse of notation, we keep using $u_{i,t}$ for shocks instead of $\Delta u_{i,t}$.

If investors have downward-sloping demand with respect to the asset's own price, then the diagonal of Γ should be negative (i.e., $\Gamma_{j,j} < 0$), which implies $M > 0$. If assets are broadly speaking substitutes with each other, then the off-diagonal terms are positive, and \tilde{M} should be positive. If assets are complements with each other, then the off-diagonal terms are negative, implying a negative \tilde{M} . We refer to M as the asset's own multiplier and \tilde{M} as the substitute passthrough.

Equation (4) suggests that if we observe asset-specific demand shocks $u_{j,t}$ and asset-specific substitutes $\Delta p_{j,t}^{sub}$, we can run a regression to estimate M and \tilde{M} . However, the construction of $\Delta p_{j,t}^{sub}$ is infeasible without further restrictions on cross-elasticities. Asset j 's $\Delta p_{j,t}^{sub}$ is a weighted average of the returns of other assets, where the weights depend on

unobserved cross-elasticities; the closer an asset is as a substitute, the larger its weights in the substitute portfolio. Hence, we will need to impose some structure on the substitution pattern, which will allow us to construct (or control for) the substitute portfolio without knowing the exact magnitudes of the cross-elasticities Γ_{jk} .

2.2 Demand Structure with Heterogeneous Cross-elasticity

Homogeneous substitution. The existing literature commonly assumes homogeneous substitution patterns. For example, logit-demand systems such as those in [Kojien and Yogo \(2019b\)](#) and [Bretscher et al. \(2022\)](#) imply homogeneous cross-elasticities; that is, $\Gamma_{j,k}$ is proportional to the market share of bond k .¹³ In this special case, $\Delta p_{j,t}^{sub}$ is essentially the market portfolio and is constant across j .¹⁴ Therefore, researchers can use the specification in equation (5) to correctly identify the multiplier M by exploiting the cross-sectional variation. Existing literature that uses flow-induced trading to identify the price impact with time fixed effects implicitly assumes a similar structure ([Lou, 2012](#)):

$$\Delta p_{j,t} = M u_{j,t} + \text{Time fixed effects} + \tilde{\nu}_{j,t}. \quad (5)$$

However, in the corporate bond market, the homogeneous substitution assumption is likely to be violated. If the price of a BBB+ bond increases, investors will probably substitute much more into another BBB+ bond than a AAA bond (all else equal). In addition, there will probably be very little substitution in the high-yield bond class. With heterogeneous substitution patterns, $\Delta p_{j,t}^{sub}$ is no longer constant across j .

When the model is misspecified, we may still correctly identify the multiplier M from equation (4) without controlling for the substitute portfolio *if the observed demand shocks $u_{j,t}$ are orthogonal to $\Delta p_{j,t}^{sub}$* . However, demand shocks among close substitutes are commonly correlated in the cross section. In our case, the demand shocks are constructed from

¹³To see the substitution is homogeneous in a logit-demand system, consider the logit-demand specification $\ln\left(\frac{w_{i,j,t}}{w_{i,0,t}}\right) = \gamma_i p_j + \beta^\top X_j + \epsilon_{i,j,t}$, where X_j is the relevant characteristic vector that affects demand. When asset k 's price increases, $\ln\left(\frac{w_{i,j,t}}{w_{i,0,t}}\right)$ does not change. In other words, $\frac{\partial \ln(w_{i,j,t})}{\partial p_k} = \frac{\partial \ln(w_{i,0,t})}{\partial p_k}$ is constant for all $j \neq k$. For the complete derivation of cross-elasticities, see Appendix F.

¹⁴Technically, the substitute is the market excluding the test asset. When each test asset is small relative to the whole market as in our baseline, the difference is negligible. Under the logit demand system, the substitute portfolio is the *exact* market portfolio. See Appendix F for details.

flows into mutual funds, which hold portfolios of closely substitutable bonds. Hence, the constructed shocks are correlated with each other. Similarly, in the demand system estimation, the instruments for bond prices are commonly constructed from investment mandates, which are mostly stipulated at the market segment level and rarely at the CUSIP level. As we demonstrate later in this section, omitting substitute portfolio return $\Delta p_{j,t}^{sub}$ will bias the multiplier estimates upward when demand shocks are positively correlated among close substitutes.

Two-layer substitution structure. We impose a two-layer structure to capture the heterogeneous substitution patterns. We allow cross-elasticities to have two levels: a stronger cross-elasticity for bonds within the same detailed rating (close substitutes) and a weaker cross-elasticity for bonds outside this category (distant substitutes). Formally, let Γ be

$$\Gamma_{j,k} = \begin{cases} \gamma^o & j = k \\ \gamma^d w_k + \gamma^c w_{k|g} & j \neq k, j, k \text{ in the same detailed rating } g \\ \gamma^d w_k & \text{otherwise,} \end{cases} \quad (6)$$

where w_k is the share of asset k in the whole market and $w_{k|g}$ is the conditional share of k in the detailed rating group g . Intuitively, the cross-elasticity is scaled by market share so that the substitution effect is in proportion to their relative sizes. Appendix F shows that this demand structure also arises from a nested-logit demand system.

Under this demand structure, the substitute portfolios for each asset j can be decomposed into two components: the portfolio of bonds in the same detailed rating group and a market portfolio. We call the former portfolio its close substitute. For example, for a 10-year BBB-rated bond, its close-substitute portfolio includes all the other BBB- bonds. Specifically, equation (4) can be rewritten as

$$\Delta p_{j,t} = M u_{j,t} + \tilde{M} \Delta p_{g(j),t} + \tilde{M}^m \Delta p_t^m + \tilde{v}_{j,t} \quad (7)$$

or alternatively as

$$\Delta p_{j,t} = M u_{j,t} + \tilde{M} \Delta p_{g(j),t} + \text{Time fixed effects} + \tilde{v}_{j,t}, \quad (8)$$

where $\Delta p_{g(j),t} \equiv \sum_{k \in g} w_{k|g,t-1} \Delta p_{k,t}$ is the portfolio return of assets in the same detailed rating category as asset j , weighted by each asset's lagged market value, and Δp_t^m is the market return. We instrument portfolio returns with demand shocks to solve endogeneity issues, as will be explained in Section 3.1. The multiplier M measures the price response to a one-percentage point increase in demand. The passthrough coefficient \tilde{M} captures the comovement of the test asset with its close substitutes, conditional on the market return. A passthrough close to 1 indicates a strong substitutability among assets within groups.

For the purpose of estimating the multiplier M alone, we can also simply absorb $\Delta p_{g(j),t}$ and Δp_t^m with group-time fixed effects, as in equation (9):

$$\Delta p_{j,t} = M u_{j,t} + \text{Fixed effects}_{g,t} + \tilde{v}_{j,t}. \quad (9)$$

Although this group-time fixed effect specification is more straightforward, the baseline specification in equation (8) allows us to estimate the substitute passthrough coefficients \tilde{M} , which is informative about the spillover effect among close substitutes. In our analysis, we estimate both specifications (8) and (9). We find that the choice of specification has little impact on the estimate of M , suggesting the IV estimation strategy for equation (8) successfully addresses the endogeneity issues.

Omitted variable bias. Failing to control for close substitutes will lead to overestimation of the multiplier and hence underestimation of the elasticity. The following example showcases the source of the bias. Suppose bond-level shocks are generated as $u_{k,t} = u_{g,t} + \tilde{u}_{k,t}$, where the bond-specific shock $\tilde{u}_{k,t}$ is orthogonal to the group-level shock $u_{g,t}$. Also for simplicity, we assume each bond is minuscule relative to the group so $Cov(p_k^g, \tilde{u}_k) = 0$. If the econometrician controls for time-fixed effects but fails to control for $\Delta p_{j,t}^g$ when estimating the multiplier, the omitted variable bias can be shown as

$$\hat{M} - M = \tilde{M} \frac{Cov(p_k^g, u_g) / Var(u_g)}{1 + \frac{Var(\tilde{u}_k)}{Var(u_g)}} \equiv \tilde{M} \frac{M_g}{1 + \frac{Var(\tilde{u}_k)}{Var(u_g)}} > 0. \quad (10)$$

The bias has three components: the passthrough coefficient \tilde{M} , the group-level multiplier M_g , and the variance ratio of the bond-specific shocks to group shocks $\frac{Var(\tilde{u}_k)}{Var(u_g)}$. Therefore, failure to control for close substitutes will contaminate the bond-level multiplier with the

group-level multiplier. As we show below, multipliers increase as we move to more aggregated portfolio levels, so the bias is positive and can be potentially large.¹⁵

This issue is not specific only to our empirical specification, but it arises whenever shocks across bonds are correlated—a typical scenario for most other demand shocks commonly used in the literature, such as institution mandates. If the substitution is indeed homogeneous as commonly assumed, \tilde{M} is close to zero so the bias is small. However, as we show in our estimation results, at the bond level the substitute passthrough is close to 1, so ignoring close substitutes dramatically overestimates the multiplier.

As we show in Section 4, at the CUSIP level the multiplier M is close to zero once the substitute portfolio is controlled, suggesting the market is fairly elastic. Therefore, for most of the heterogeneity analysis, we conduct the analysis at the portfolio level. In our baseline portfolio specification, we treat each asset as a market-value weighted portfolio formed by bonds with the same detailed rating category and quarter to maturity. Like the individual bond analysis, we assume there are two cross-elasticities, one for portfolios with the same detailed ratings (close substitutes) and another for portfolios outside this category (distant substitutes).

Finally, with enough data, one can easily extend our analysis to allow for multiple layers of substitution: same detailed rating category, the same investment rating (investment-grade versus high-yield), and all the other cases. We verify that the estimation results are quantitatively very similar to the case with two different levels of cross-elasticities. So we will focus on only two levels of heterogeneity for most parts of the paper.

3 Empirical Strategy

In this section, we first describe our identification strategy. We then introduce the data sources and explain the construction of our demand shocks in detail. We present summary statistics regarding the demand shocks at the end of this section.

¹⁵We are implicitly assuming that the cross-elasticity between corporate bonds and other asset classes (such as Treasury securities, mortgage-backed securities) is small. Hence, the bias that is due to not including those asset returns explicitly is negligible.

3.1 Identification

In order to identify the asset's multiplier M and the substitute passthrough \tilde{M} from the specifications in equation (8) and equation (9), we need the unobserved demand shock $\tilde{\nu}_{j,t}$ to be orthogonal to the observed demand shocks $u_{k,t}$ for any k . In addition, in equilibrium, $\Delta p_{j,t}^{sub}$ is endogenous and is determined by the whole vector of demand shocks \mathbf{u}_t and ν_t . Hence, we also need an instrument for $\Delta p_{j,t}^{sub}$ when estimating equation (8).

In Appendix A, we show that under Assumption 1, $\tilde{\nu}_{j,t}$ is orthogonal to the constructed demand shocks $u_{j,t}$. To solve the endogeneity problem of $\Delta p_{j,t}^{sub}$, we instrument it with the portfolio's demand shocks $u_{j,t}^{sub}$, defined as

$$u_{j,t}^{sub} \equiv \sum_{j \neq k, k \in g(j)} w_{k,t}^j u_{k,t}. \quad (11)$$

Under Assumption 1, $u_{j,t}^{sub}$ is also orthogonal to $\nu_{j,t}$ in the cross section.

Assumption 1. *For any mutual fund i in each period t , the mutual fund's lagged portfolio share in bond j is orthogonal to its unobserved demand shocks:*

$$\mathbb{E}_{i,t} [\tilde{\nu}_{j,t} S_{i,j,t-1}] = 0 \quad \forall i, t.$$

Appendix A shows that under Assumption 1, the multiplier and the passthrough coefficient are identified.¹⁶ Our identification strategy uses the same insight as in Goldsmith-Pinkham et al. (2020) that Bartik instruments can be viewed as exogenous share instruments. As a result, the strategy is similar to an exposure research design, where bonds have exogenous exposures to common shocks, and the degree of exposure depends on the holding shares of mutual funds.

To make the pooled exposure design intuition more concrete, suppose households receive a preference shock that increases their demand for bonds (this shock may be correlated with the fundamental characteristics of the bond market). As a result of the household preference

¹⁶One potential concern about Assumption 1 is that if the unobserved shock $\nu_{j,t}$ loads on a factor, such as $\nu_{j,t} = \tilde{\nu}_{j,t} + \beta_j \eta_t$, and there exist factor-focused funds whose portfolio shares are functions of bonds' factor loading (e.g., $Q_{j,t-1} = \psi_{t-1} \beta_j$), then the orthogonality condition in Assumption 1 will be violated. We address this issue in Appendix C. In essence, we show that as long as the fund-specific demand shock $u_{i,t}$ is orthogonal to factors, our identification is still effective. Our results are robust to different factor specifications in flows (see Appendix C).

shock, bond mutual funds will receive inflows. Consider a mutual fund that holds two bonds and before the shock held 10% of bond A and 5% of bond B. Suppose this mutual fund receives an inflow that is 10% of its assets under management (AUM), and mechanically invests these inflows according to its previous period's portfolio weights. In this case, bond A would experience a demand shock of 1%, whereas bond B would experience a demand shock of 0.5%. If the previous period's holding shares of the mutual fund are orthogonal to the shocks, then the increase in the demand for bond A relative to that for bond B will be as good as randomly assigned.¹⁷ Hence, given the exogeneity of funds' lagged holding shares and unobserved shocks, flow-induced trading by a single mutual fund amounts to an exogenous exposure research design. Our identification strategy pools many such instances of flow-induced trading, and hence it can be viewed as a pooled exposure design.

Robustness of Assumption 1. Below we discuss the potential violations for Assumption 1 and how we address them. One potential threat to this assumption is market-timing skills by mutual fund managers. Suppose some mutual fund managers have superior information about bond idiosyncratic returns in the future and allocate their portfolio in advance to front-run the market; in this case, the share exogeneity will be violated. This issue can be addressed using longer lags for shares. In Appendix C.4, we show the results are robust if we aggregate shocks using shares with one-year lags. It is unlikely that mutual fund managers are able to predict bond-level *idiosyncratic* returns in one year and allocate assets accordingly.

Another potential concern is factor loadings driving portfolio shares as well as bond returns. Bond returns can load on common factors, such as credit risks, and there are style-specialized funds that overweigh bonds with higher loadings on one particular factor. For example, suppose the unobserved shocks $\nu_{j,t}$ have the following factor structure:

$$\tilde{\nu}_{j,t} = \beta_j \eta_t + \check{\nu}_{j,t},$$

where η_t is the factor and $\check{\nu}_{j,t}$ is the true idiosyncratic shock. Also suppose one mutual fund

¹⁷If the fund has mandates restricting what types of bonds the fund can hold, then the part of the previous period's holding shares that is determined by mandates is naturally orthogonal to demand shocks.

i specializes in factor η , which means

$$S_{i,j,t-1} = \xi_i \beta_j + \check{S}_{i,j,t-1}.$$

In this case, the share exogeneity condition under $\nu_{j,t}$ will be violated:

$$\mathbb{E}_{i,t} [\nu_{j,t} S_{i,j,t-1}] = \mathbb{E}_{i,t} [\beta_j^2] \xi_i \eta_t \neq 0.$$

The solution to this issue is straightforward. We can explicitly control for common factors in the final regression, allowing for heterogeneous loadings:

$$\Delta p_{j,t} = M u_{j,t} + \tilde{M} \Delta p_{j,t}^{sub} + \beta_j \eta_t + \check{\nu}_{j,t}, \quad (12)$$

so that shares are exogenous to $\check{\nu}_{j,t}$. As is common in the bond literature, we make an additional assumption that β_j depends on the characteristics of bond j , such as maturity or ratings. Then the factor-loading term $\beta_j \eta_t$ can be absorbed by characteristic-time fixed effects.¹⁸ This provides a robust method of controlling for common factors without estimating a factor model.

3.2 Data

Mutual fund data: We obtain detailed data on open-end mutual funds in the US from Morningstar Inc. Morningstar is one of the largest providers of investment research to the asset management industry. Detailed holdings and fund flows are collected from surveys of mutual fund managers and cross-validated by Morningstar against publicly available sources such as regulatory filings to ensure their accuracy. Most funds report at least once per quarter. In order to keep as many funds as possible in our sample, we conduct our analysis at the quarterly frequency. The Morningstar mutual fund coverage is quite extensive: the dataset's total AUM lines up closely with the flow of funds open-end mutual fund sector AUM (see Figure 9 in Appendix B).

Since our identification relies on flow-induced trading in the US corporate bond market, we apply some additional filters to narrow in on domestic bond mutual funds. First, we

¹⁸Our results are also robust to controlling for off-the-shelf factors from the bond literature, such as [Bai et al. \(2019\)](#).

restrict our sample to funds that report their portfolio value in US dollars and also to securities denominated in US dollars. Second, we focus on funds that have at least \$10 million of bonds under management and bonds make up 50% to 120% of their portfolio's AUM. The lower bound is to filter out non-bond funds, and the upper bound is to filter out potential misreporting. Taken together, this database gives us coverage of 1,151 unique funds from 2002Q1 to 2021Q3, totaling around 5.6 million bond-fund-quarter observations.

After applying our data filters, we then define mutual fund i 's dollar flows F_{it} scaled by lagged $AUM_{i,t-1}$ as

$$f_{it} := \frac{F_{it}}{AUM_{i,t-1}}.$$

This is a key building block for constructing our demand shocks. We denote the set of investors for whom we observe flow data as MF .

Bond return data: We use corporate bond returns data from WRDS Bond Returns file, which is constructed using transaction-level data from FINRA's TRACE (Trade Reporting and Compliance Engine) database, and Mergent FISD data for bond issue and issuer characteristics. Specifically, our quarterly return measure is the cumulative end-of-month returns of the months that fall in a given quarter. We further merge this dataset with Morningstar's CUSIP-level bond characteristics.

Additionally, we apply filters to ensure we are capturing corporate bonds with reliable return data. First, we restrict our sample to securities that WRDS classifies as corporate bonds. Second, we restrict our attention to corporate bonds that have at least a CCC- or higher rating by S&P. Overall, we have 377,753 bond-quarter observations.

3.3 Constructing Flow Shocks

Our shock measure builds on the flow-induced trading (FIT) measure proposed by Lou (2012). The measure exploits the fact that mutual funds tend to mechanically invest in-flows/outflows according to their existing portfolio weights (Coval and Stafford, 2007). Ultimately, flow-induced demand shocks are essentially a Bartik instrument, where the key identifying assumption is the exogeneity of lagged mutual funds holding shares (Goldsmith-Pinkham et al., 2020; Chaudhry, 2022). With that said, several omitted variable bias concerns arise when directly using FIT as a demand shock measure; hence, we need to apply

several adjustments to get from FIT to our demand shock (Gabaix and Koijen, 2022). Below, we walk through how we construct our bond-specific demand shock measure.

Step 1: Measuring flow-induced trading by mutual funds. We assume that fund i mechanically reinvests dollar flow F_{it} according to the fund's existing portfolio weight $\theta_{ij,t-1}$. This implies that the dollar amount of the flow-induced trading of bond j will be $\theta_{ij,t-1}F_{it}$. The percentage change in asset j 's holdings due to flows is

$$f_{it} = \frac{\theta_{ij,t-1}F_{it}}{\theta_{ij,t-1}AUM_{i,t-1}}$$

Implicitly, in this step we have assumed that the flows fully pass through to the demand shocks to each bond one to one. If the passthrough is less than one to one, then we risk underestimating the multiplier, since we overestimate the quantity of the shock. We test this assumption in Appendix C.1 by estimating the passthrough of flows to actual trading. Notice that this passthrough is a lower bound of the passthrough to demand, as the actual trading also reflects other equilibrium objects such as prices. Overall, we find that the estimated passthrough is sizeable (around 0.5 to 0.6) at the individual bond level. Therefore, we can use our passthrough estimates to get upper bounds on the multipliers.¹⁹ We find that the upper bounds are still considerably smaller than the micro multipliers typically found for stocks and bonds.²⁰

Step 2: Adjusting for predictability in flows. Because of return chasing by households, mutual fund flows and flow-induced trading tend to be predictable (Lou, 2012). As a result, forward-looking investors will trade today in anticipation of future flows. Hence, the relevant demand shock is not just the flow-induced trading today, but rather, the total amount of flow-induced trading predicted by innovations in flows today (Gabaix and Koijen, 2021). To determine flow innovations, we estimate an AR(3) model with a time trend for each fund

¹⁹At the bond level, the upper bound multiplier estimate is 0.1 (elasticity of 10), and for our baseline level of aggregation, it is around 0.33 (elasticity of 3).

²⁰One may be concerned about funds selectively trading more liquid bonds. We do not find the passthrough to be correlated with measures of liquidity such as bid-ask spreads. Finally, our estimates are quantitatively the same if we exclude crisis periods, where the selling of bonds is more likely to be tilted toward liquid assets.

²¹*i*,

$$f_{it} = \rho_{i,0} + \sum_{k=1}^3 \rho_{i,k} f_{i,t-k} + \delta_i t + \varepsilon_{it}.$$

Assuming relatively little discounting by forward looking agents at the quarterly level, the relevant quantity of the demand shock is the total cumulative trading predicted by the innovation,

$$K_i \varepsilon_{it},$$

where $K_i = \frac{1}{1 - \sum_{k=1}^3 \rho_{i,k}}$.

Step 3: Even though our main identification assumption does not rely on flow exogeneity, as additional robustness checks we also remove potential common factors from the flow shocks. We estimate a model of the form

$$\varepsilon_{i,t} = \alpha_i + \delta_t + \lambda_i^f \eta_t + u_{i,t},$$

where $u_{i,t}$ is the fund-specific idiosyncratic flow innovations. Hence, our measure for fund-specific demand shocks is $K_i u_{it}$. We find that mutual fund flows have a weak factor structure after controlling for time fixed effects. Therefore, in the baseline, we do not control additional factors η_t to avoid introducing additional estimation errors. Controlling for additional factors makes little difference to the estimates (see Appendix C).

Step 4: Aggregate fund-level shocks to bond-level shocks. To obtain bond-level demand shocks, we aggregate the fund-level shocks $K_i u_{it}$ to the bond level using funds' lagged holding share weights. With a slight abuse of notation, we denote bond j 's demand shock in quarter t as u_{jt} . Specifically, u_{jt} is defined as

$$u_{jt} = \sum_{i \in MF} S_{ij,t-1} K_i u_{it},$$

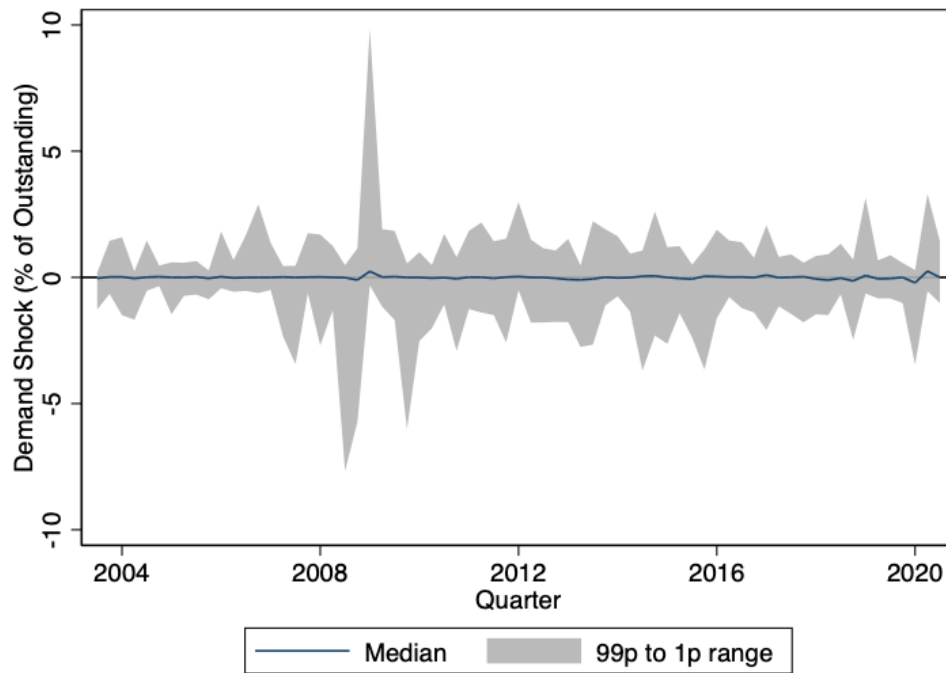
where $S_{ij,t-1}$ is the market share of bond j that is held by investor i in period $t - 1$.

Figure 3 plots the 1st to 99th percentile range of the shocks for each quarter, along with its median. The magnitudes of shocks are generally quite large, ranging from around 1% to 3% of the total outstanding amount. The range is relatively stable in magnitude apart from

²¹Our results are robust to changing the specification of the AR process (see Appendix C.5).

two crisis episodes in 2008 and 2020. For robustness, we explore including and dropping the crisis periods from our estimation sample and find that our results are robust to this decision.

Figure 3: Demand shocks ($N = 339,847$)



This graph plots the median value of our demand shock measure in a given quarter, and the shaded grey area denotes the range of demand shock values that fall between the 1st and 99th percentile of shocks in any given quarter. The shocks are represented as a percentage of the total outstanding amount.

4 Estimation Results

We present our main findings in this section. We first show that the multiplier at the CUSIP level drops from 0.4 to 0.05 once we account for the heterogeneity in substitution patterns. Second, when we form bonds into our baseline portfolios, the portfolio multiplier becomes significantly larger, indicating demand is more inelastic at the portfolio level. Furthermore, the multiplier monotonically increases as we form increasingly more aggregated portfolios,

whereas the substitute passthrough monotonically decreases. Furthermore, we investigate the dynamics of the price impact at the portfolio level and find the effect almost fully reverts after two quarters, suggesting demand is more elastic in the long run. Finally, we show that the multiplier is larger for bonds with higher default risks and longer time to maturity.

4.1 Baseline Estimates

CUSIP Level Table 1 presents our baseline estimates treating each CUSIP as one asset. The first column corresponds to assuming homogeneous cross-substitution. Under this assumption, every asset has the same substitute portfolio, which is simply the market portfolio. This is equivalent to including time fixed effects in the regression. Hence, we can simply run the regression specified by equation (5). Under this assumption, we get a significant CUSIP multiplier of 0.39, indicating that a 1% increase in demand leads to a 0.39% increase in the bond's price. This implies that the average elasticity for a bond is around 2.6. The magnitude of the average elasticity is in line with the estimate in the literature using a standard logit demand system (Bretschler et al., 2022; Siani, 2022a).

Once we relax the homogeneous cross-substitution assumption and allow certain groups of bonds to be closer substitutes compared to others, the CUSIP multiplier drops to nearly zero, implying very elastic demand at the individual bond level. Columns (2)-(4) in Table 1 present results using the specification in equation (9). We include a group time fixed effect to control for the returns of the substitute portfolios, where groups are defined as bonds within the same detailed rating group. Furthermore columns (5)-(6) show the results from running the regression in equation (8), where we include the returns of a substitute portfolio, $\Delta p_{j,t}^{sub}$. The substitute portfolio is defined as the market-value weighted portfolio formed by all bonds within the same detailed rating category (excluding the asset itself). We instrument the substitute portfolio's ($\Delta p_{j,t}^{sub}$) return with the market-value weighted demand shocks for all bonds in the substitute portfolio $u_{j,t}^{sub}$, defined in equation (11). We also include a time fixed effect in all the regressions, which means we are controlling for the market portfolio return (the distant substitute) as well. Effectively, we are allowing the cross-elasticity to be different for bonds within the same rating category relative to all the other bonds.

Our result implies the demand elasticity at the CUSIP level is around 20,²² which is

²²Under the unit-elasticity null, we estimate the 95% confidence interval for the elasticity is (19.9, 20.1).

much larger than the current estimates in the literature. In Appendix C.6, we show that the magnitude of the estimate does not change when only looking at bonds with large shocks (above median in absolute size). Among the elasticity estimates in the literature, Fang (2022)'s estimate of 10 is the closest to ours. Compared to other demand systems used for estimating corporate bond elasticities, Fang (2022) uses a somewhat more flexible nested logit structure, allowing for differential substitution within and across IG and HY categories. Overall, our result emphasizes the importance of explicitly modeling the rich heterogeneous substitution patterns in the corporate bond market.

As mentioned in Section 3.3, in the case when the passthrough from fund flows to security-level flows is less than one to one, we can use the estimated passthrough coefficient to provide an upper bound for the multiplier. The upper bound is approximately 0.1, implying an elasticity of 10. Furthermore, the bias introduced by imperfect passthrough should be the same in all the specifications; hence, it cannot explain the difference in the estimates between column (1) (assuming homogeneous cross-substitution) and columns (2)-(4) (heterogeneous cross-substitution). The difference between the two cases highlights the importance of allowing heterogeneous cross-substitution patterns. Finally, the one-to-one passthrough assumption is more likely to hold true when shocks are larger. Appendix C.6 shows that our estimates are stable when we only use large shocks.

In addition to a near-zero CUSIP multiplier, the substitute passthrough (\tilde{M}) is highly significant and quantitatively close to 1. This implies that within the detailed rating groups, bonds are effectively perfect substitutes with each other. Finally, our point estimates are stable whether or not we exclude crisis periods or include additional fixed effects.

Our results imply that ignoring the substitute portfolio's return in the regression leads to a significant overestimation of the multiplier. Intuitively, when we observe a demand shock to bond j , it is likely that other similar bonds are also experiencing demand shocks. The price of bond j is higher not only because of its own demand shock but also because all of its substitutes now have higher prices. Hence, ignoring the substitute portfolio will lead to overestimation of the price impact from its own demand shock. Furthermore, in the case when certain bonds are closer substitutes than others, we need to overweight the close substitutes in constructing the substitute portfolio, or alternatively, allow the close substitutes to have a different coefficient from the other bonds.

In Appendix E, we estimate the equity market multiplier at the stock level, following

identical procedures. We define close substitutes as stocks with similar factor loadings or stocks in the same industry group. We find that allowing for heterogeneous cross-substitution reduces the multiplier estimated, as in the corporate bond case. However, quantitatively, the difference is not as big as that in the bond market, suggesting that heterogeneous substitution patterns are a particularly special feature of the bond market. Perhaps because rating and maturity play a significant role in investors' demand, accounting for heterogeneous cross-substitution is much more important in the bond market than in the equity market.

Portfolio Level Next, we form portfolios of bonds based on detailed rating and quarter to maturity, and estimate the multiplier at the portfolio level. In addition to the market portfolio (which is taken care of by the time effect), we include the return of a substitute portfolio that is formed by all bonds in the same detailed rating category as before.

The results are presented in Table 2. As before, column (1) corresponds to the result using the specification in equation (5). Columns (2)-(4) report results using the specification in equation (9) with different definitions of close substitutes. Furthermore, columns (5)-(6) relate to the IV strategy with the specification in equation (8). In the rest of the paper, unless specified otherwise, equation (8) is our baseline specification, where we instrument $\Delta p_{j,t}^{sub}$ with $u_{j,t}^{sub}$, since it allows us to estimate and interpret the coefficient in front of the substitute portfolio returns.

Similar to the CUSIP case, we find that only including the market portfolio as the substitute significantly overestimates the portfolio multiplier. Once we include additional substitute portfolios, the multiplier estimated is considerably lower.²³

A priori, we expect the portfolio-level multiplier to be larger than the CUSIP multiplier. Intuitively, the substitutability among portfolios with different maturities should be smaller compared to the substitutability among CUSIPs with the same maturity. Indeed, Table 2 shows the multiplier on portfolios is significantly larger than the multiplier for individual CUSIPs: a 1% increase in demand for a given portfolio raises the price for that portfolio by 35 basis points. This estimate is stable across various specifications. Additionally, the substitute passthrough at the portfolio level is also smaller than that at the CUSIP level, consistent with the intuition that more aggregated portfolios are less substitutable.

²³In Appendix C.6, we show that the estimate does not change much when only looking at portfolios with large shocks (above median in absolute size).

To verify whether our demand specification sufficiently accounts for important cross-substitution heterogeneity, we further allow heterogeneity in cross-elasticities for bonds in the same investment-grade (IG) or high-yield (HY) group versus bonds that are not. We implement such specification by including an additional substitution portfolio formed by all the IG or HY bonds (excluding the own asset and close substitute's returns). The result is shown in Table 3 column (2). The own multiplier does not change quantitatively. In addition, the substitution passthrough is much larger for bonds within the same rating category compared with bonds not in the same rating category but in the same IG/HY group. Columns (3) and (4) of Table 3 further confirm that our result is not sensitive to the choice of the specific substitute portfolio. Column (3) uses all bonds in the same coarse rating group to form the substitute portfolio (i.e., using a broader definition of the set of close substitutes).²⁴ We can also group bonds into two maturity groups based on the partition: $\{[0, 10), [10, \infty)\}$. Column (4)'s substitute portfolio includes all bonds in the same maturity group and detailed rating category (i.e., a more narrow definition of the set of close substitutes). Across all the different alternative specifications, the estimated own multiplier is similar to that in the baseline case.

²⁴For example, for 11-year BBB-rated bonds, the substitute portfolio includes all the BBB-rated bonds.

Table 1: CUSIP baseline

	Homo. OLS	OLS			First-stage	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)
Shock	0.39*** (0.06)	0.02 (0.05)	0.05 (0.05)	0.07 (0.04)		0.05 (0.05)
Substitute return						1.07*** (0.06)
Group Shock					2.61*** (0.30)	
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Group x Quarter FE	No	Yes	Yes	Yes	No	No
ST/LT x Quarter FE	No	No	No	Yes	No	No
Drop Crisis	No	No	Yes	Yes	Yes	Yes
N	333,537	333,537	314,534	314,534	314,534	314,534
R^2	0.21	0.39	0.36	0.40	0.62	
First-stage F-statistic						72.49

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table presents the CUSIP-level multiplier estimates. Column 1 shows the OLS estimates from regressing bond returns on the demand shock, controlling for time fixed effects, as in equation (5). This specification corresponds to a model in which we assume homogeneous cross-elasticity with all other bonds. Columns (2)-(4) show results from running the regression in equation (9) with different definitions of close substitutes. Columns (2) and (3) directly control for close-substitute prices using detailed rating x time fixed effects. Column (4) additionally controls for maturity (long-term/short-term) x quarter fixed effects to control for a potentially omitted time-varying maturity factor in holding shares. Columns (5) and (6) relate to the IV specification in equation (8). We regress bond returns on the demand shock and substitute returns, while controlling for time fixed effects (and additional controls depending on the specification). We instrument for substitute returns using demand shocks to the substitute portfolio. The IV specification corresponds to the model that allows for heterogeneous cross-elasticities between close and distant substitutes. Standard errors are clustered at the substitute group x time level.

Table 2: Quarter to Maturity x Detailed Rating baseline

	Homo. OLS	OLS		First-stage	2SLS	
	(1)	(2)	(3)	(4)	(5)	(6)
Shock	0.95*** (0.12)	0.32** (0.10)	0.33*** (0.10)	0.35*** (0.10)		0.35*** (0.10)
Substitute return						0.90*** (0.05)
Group shock					2.91*** (0.34)	
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Group x Quarter FE	No	Yes	Yes	Yes	No	No
ST/LT x Quarter FE	No	No	No	Yes	No	No
Drop Crisis	No	No	Yes	Yes	Yes	Yes
N	81,866	81,866	77,387	77,387	76,348	76,348
R^2	0.23	0.47	0.41	0.44	0.47	
First-stage F-statistic						69.85

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table presents the quarter to maturity \times detailed rating level multiplier estimates. Column (1) shows the estimates from regressing portfolio returns on the demand shock, controlling for time fixed effects, as in equation (9). This specification corresponds to a model in which we assume homogeneous cross-elasticity with all other bonds. Columns (2)-(4) show results from running the regression in equation (9) with different definitions of close-substitutes. Columns (2) and (3) directly control for close-substitute prices using detailed rating \times time fixed effects. Column (4) additionally controls for maturity (long-term/short-term) \times quarter fixed effects to control for a potentially omitted time-varying maturity factor in holding shares. Columns (5) and (6) relate to the IV specification in equation (8), in which we regress portfolio returns on the demand shock and substitute returns, while controlling for time fixed effects (and additional controls depending on the specification). We instrument for substitute returns using demand shocks to the substitute portfolio. The IV specification corresponds to the model that allows for heterogeneous cross-elasticities between close and distant substitutes. Standard errors are clustered at the substitute group \times time level.

Table 3: Alternative Substitutes

	(1)	(2)	(3)	(4)
Shock	0.35*** (0.10)	0.35*** (0.10)	0.31** (0.11)	0.33** (0.11)
Detailed rating substitute return	0.90*** (0.05)	0.77*** (0.17)		
IG substitute return		0.18 (0.22)		
Coarse rating substitute return			0.79*** (0.08)	
Det rating x ST/LT substitute return				0.99*** (0.08)
Quarter FE	Yes	Yes	Yes	Yes
Drop Crisis	Yes	Yes	Yes	Yes
<i>N</i>	76,348	76,348	76,296	76,348
First-stage F-statistic	69.85	3.37	18.19	107.94

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table shows that the quarter to maturity \times detailed rating level multiplier estimates are robust to the exact definition of the close substitute. Column (1) is our baseline specification in which the close substitute is defined as all other bonds in the same detailed rating category. In column (2), the substitute includes the detailed rating substitute and all bonds with the same investment rating as an additional substitute. Column (3) includes all other bonds in the same coarse rating category as substitutes. Column (4) includes all bonds in the same detailed-rating group and similar maturity as substitutes. Standard errors are clustered at the substitute group \times time level.

4.2 Aggregate Portfolios

Motivated by the difference in the CUSIP and the portfolio multiplier, we repeat our analysis for portfolios with various aggregation levels, using the specification in equation (8). We find that the more aggregated the portfolio is, the larger the portfolio multiplier and the smaller the substitute passthrough.

Specifically, we estimate the multiplier and substitute passthrough for all the portfolios specified in Table 4. The second column specifies at which level we form the portfolios, and the third column specifies the substitute portfolio included in the regression in addition to the market portfolio. Intuitively, as we move to more aggregated portfolios, they become less substitutable with each other. As a result, we expect demand shocks to have a larger price impact, and the relationship between the portfolio's price and its substitute's price becomes weaker.

Table 4: Aggregate Portfolios

	Asset	Substitute Portfolio
CUSIP	Individual bonds	Other bonds in the same detailed rating category
Rating \times Quarter to Maturity	Portfolios formed by detailed rating and quarter to maturity	Other bonds in the same detailed rating category
Rating \times ST/MT/LT Buckets	Portfolios formed by coarse rating and three maturity groups ($\{[0, 4), [4, 10), [10, \infty)\}$)	Other bonds in the same coarse rating category
Rating \times ST/LT Buckets	Portfolios formed by coarse rating and two maturity groups ($\{[0, 10), [10, \infty)\}$)	Other bonds in the same coarse rating category
Rating	Portfolios formed by coarse rating categories	Other bonds in the same investment grade category

This table presents all levels of portfolios for which we estimate multipliers and substitute passthroughs. The left column specifies at which level the portfolio is formed. All bonds are weighted by their market-value inside the portfolio. The right column specifies the additional substitute portfolio included in the regression in addition to the time fixed effect.

This is indeed what we find. Figure 1 visualizes our findings, and Table 5 presents

the exact estimates. To ensure our results are not driven by tail aggregate shocks, we exclude the crisis periods for all regressions onward. From Figure 1, we see that as we move from the CUSIP-level estimate to the most aggregated portfolios formed by ratings, the asset's own multiplier increases from 0.05 to 3.5, implying the demand elasticity drops from 20 to less than 0.3.²⁵ Intuitively, it is easy to find substitutes for an individual bond, whereas it is much more difficult to find substitutes for the entire portfolio of BBB bonds. Hence, demand is much more elastic for individual CUSIPs compared with more aggregated portfolios. Furthermore, the substitute passthrough decreases from 1.07 to essentially 0. At the rating portfolio level, assets are much weaker substitutes with each other compared to those at the CUSIP level.

Our aggregation results highlight the connection between micro and macro multipliers. Demand elasticities naturally depend on what we define as an asset. Micro multipliers typically refer to the multiplier on an individual security, whereas macro multipliers refer to the multiplier when we treat a whole asset class as one.²⁶ They are at the two ends of a spectrum, and our results show that depending on the portfolio definition, the multiplier estimated can lie anywhere in between. In fact, demand inelasticity at the portfolio level is what allows us to estimate the multipliers at more micro levels: the significance in the first stage of all the IV regressions relies on portfolio prices responding to demand shocks.

Finally, the difference in elasticities and spillover effects across portfolios with different aggregation levels also has interesting implications for unconventional monetary policy implementation. For a given size of an asset purchase program, if the central bank concentrates all the asset purchases on a narrow portfolio, the demand shock (measured as a fraction of amount outstanding) will be large, but the price impact per unit of demand shock will be small. If, however, the central bank purchases more aggregated/broad portfolios, the demand shock is smaller, but the per unit price impact is large. The spillover effects to assets outside of the purchase program will also be different under the two scenarios. Our estimates of M and \tilde{M} serve as important inputs to this analysis. We leave a proper quantitative investigation to future research.

²⁵We have also formed portfolios by issuers. The multiplier estimated is 0.11, and the substitute passthrough is 0.98. Both fall between the CUSIP level estimates and baseline portfolio level estimates.

²⁶Our identification strategy does not allow us to uncover the macro elasticity for the corporate bond market as an asset class.

Table 5: Aggregation

	CUSIP	Rat x Q to Mat	Rat x 3 Mat	Rat x 2 Mat	Rat
	(1)	(2)	(3)	(4)	(5)
Shock	0.05 (0.05)	0.35*** (0.10)	1.23* (0.48)	1.59** (0.58)	3.51*** (0.87)
Substitute return	1.07*** (0.06)	0.90*** (0.05)	0.73*** (0.13)	0.48** (0.15)	-0.02 (0.31)
Quarter FE	Yes	Yes	Yes	Yes	Yes
Drop Crisis	Yes	Yes	Yes	Yes	Yes
N	314,534	76,348	1,407	938	469
First-stage F-statistic	72.49	69.85	33.83	16.15	17.78

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table shows the multiplier and substitute passthrough estimates for various levels of portfolio aggregation (see Table 4 for descriptions of the aggregation). All the specifications are estimated using IV. Specifically, we instrument the substitute's return using demand shocks to the assets in the substitute portfolio. Standard errors are clustered at the substitute group x time level.

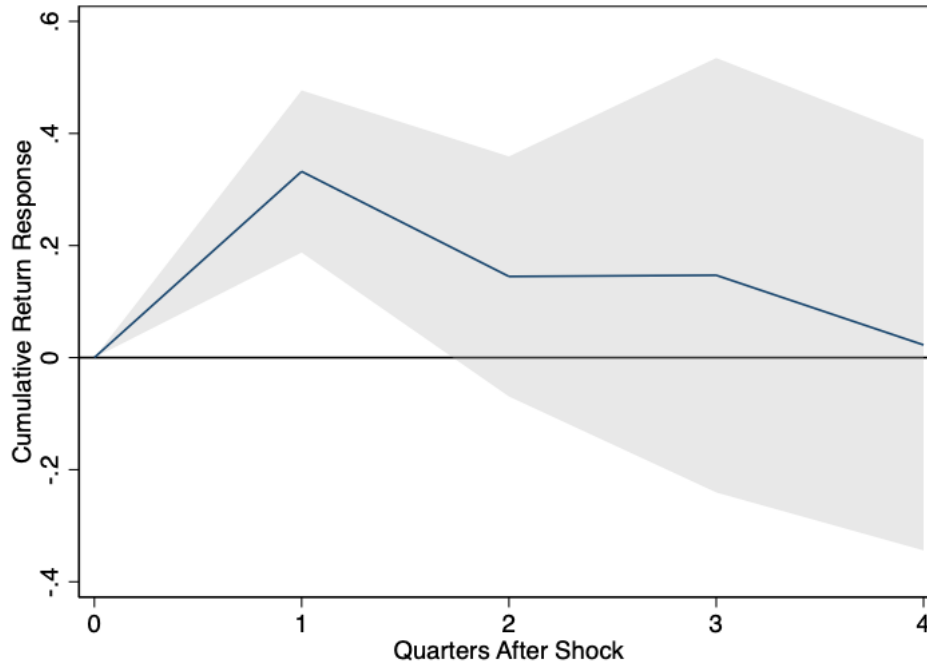
4.3 Dynamic Effects

We study the dynamic effects of price response to demand shocks in this section. We focus on the baseline portfolio specification in Section 4.1 (i.e., portfolios are formed by bonds with the same detailed rating and quarter to maturity). To see the dynamic effects, we regress portfolios' cumulative returns on lagged demand shocks. Specifically, we run the following regression:

$$\Delta p_{j,t:t+h} = M_h u_{j,t} + \text{Fixed effects}_{g,t+h} + \varepsilon_{j,t+h} \quad \text{for } h = 0, 1, \dots \quad (13)$$

Figure 4 plots the response of cumulative returns to demand shocks. Consistent with our previous result, there is a significant on-impact price increase upon positive demand shocks. The effect almost fully reverts in the following quarter, and the price impact is statistically indistinguishable from zero. In other words, the cumulative effect drops to 0 after $T + 1$. The reversal in return after six months is consistent with the hypothesis that long-term demand

Figure 4: Price Impact Dynamics



The graph plots the dynamic price impact estimates of the demand shock. The estimates are for the detailed rating \times quarter to maturity level (our baseline level of aggregation). Specifically, we regress various leads of the return of the portfolio in excess of its close substitute return on the demand shock. The shaded region denotes the 95% confidence interval.

is more elastic than short-term demand. Table 6 shows the exact estimates.

Table 6: Dynamic Effects

	T	T+1	T+2	T+3
	(1)	(2)	(3)	(4)
Shock	0.33***	0.14	0.15	0.02
	(0.07)	(0.11)	(0.20)	(0.19)
Group x Quarter FE	Yes	Yes	Yes	Yes
N	77,387	55,466	45,770	39,584
R^2	0.41	0.57	0.60	0.61

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table presents the dynamic cumulative price impact estimates of the demand shock. The estimates are for the detailed rating \times quarter to maturity level (our baseline level of aggregation). Specifically we regress the cumulative return on the demand shock while controlling for substitute group \times time fixed effects. Standard errors are clustered at the substitute group \times time level.

4.4 Rating and Maturity Heterogeneity

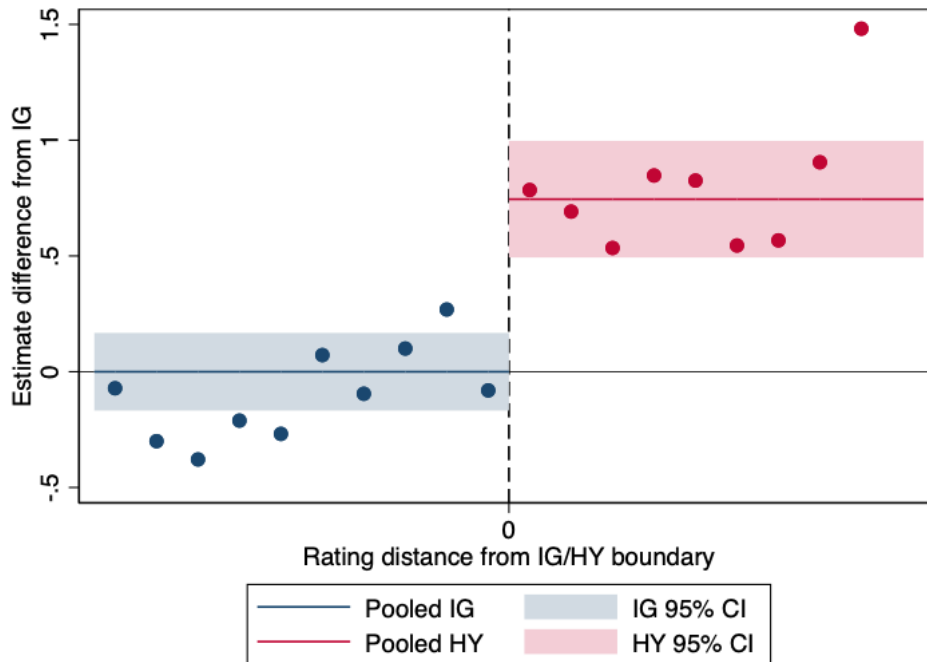
Finally, we look at how multipliers vary across different types of bonds. We focus on the heterogeneity in multipliers for bond portfolios with different maturities and ratings. As before, we use our baseline portfolio specification.²⁷ As shown in Section 4.1, the exact substitute definition does not matter for the estimation results.

We first estimate portfolio multipliers for IG versus HY bonds. We pool all the IG portfolios together and find the portfolio multiplier is close to 0. For the pooled HY portfolios, the estimated multiplier is around 0.6, much higher than that for IG portfolios. The overall estimate of 0.35 masks significant heterogeneity among portfolios with different ratings. We also estimate the portfolio multipliers for each rating category. The results are shown as scattered dots in Figure 5. On average, riskier portfolios have higher portfolio multipliers. HY bonds tend to be riskier with larger idiosyncratic risks. As a result, it may be more difficult to substitute to other assets, which leads to higher portfolio multipliers.

The discrete change in portfolio multipliers around the IG/HY cutoff suggests that investor segmentation may be a source of demand inelasticity. Figure 6 zooms in closer to

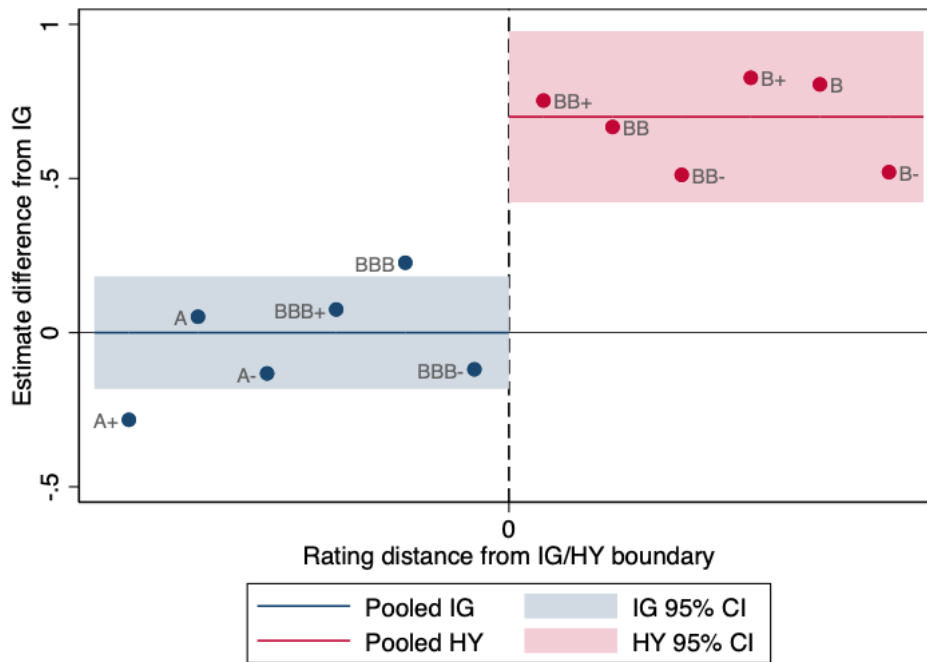
²⁷Each asset is a portfolio formed by bonds with the same detailed rating category and quarter to maturity, and the substitute portfolio includes all the other bonds in the same detailed rating category.

Figure 5: IG/HY Segmentation Effect ($N = 28,192$)



The graph shows how the detailed rating \times quarter to maturity multiplier estimates differ across ratings. Specifically, the estimates are plotted relative to the pooled investment-grade multiplier estimate. The blue and red dots are the heterogeneous estimates for IG and HY detailed rating categories, respectively. The red line represents the relative estimate for HY bonds. The blue and red shaded areas represent the 95% confidence interval for the IG and HY pooled estimate, respectively.

Figure 6: Zooming in on IG/HY Segmentation Effect ($N = 20,372$)

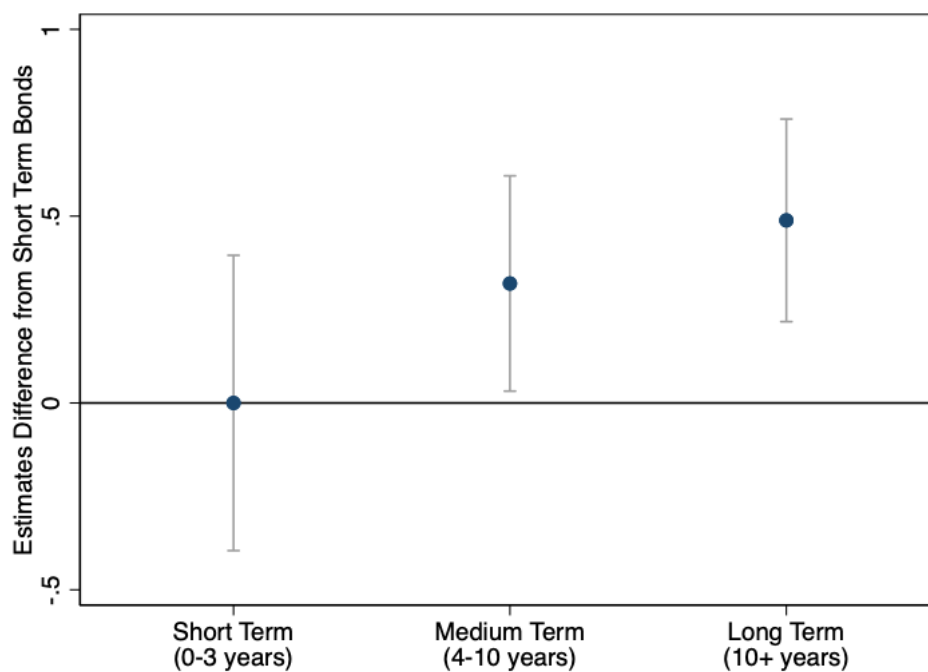


The graph shows how the detailed rating \times quarter to maturity multiplier estimates differ across the six detailed ratings above and below the IG/HY cutoff. Specifically, the estimates are plotted relative to the multiplier estimate of the six IG categories just above the cutoff. The blue and red dots are the heterogeneous estimates for IG and HY detailed rating categories, respectively. The red line represents the relative pooled estimate for six detailed rating categories just below the cutoff. The blue and red shaded areas represent the 95% confidence interval of the zoomed-in pooled IG and HY estimates, respectively.

the IG/HY cutoff. On the left-hand side of the cutoff, the point estimates for BBB-rated portfolios and A-rated portfolios are all close to 0. Once we cross the cutoff, the estimated portfolio multipliers jump discretely upward to around 0.6. Investor segmentation is particularly strong at the IG/HY cutoff. Many investors in the corporate bond market face strong incentives to stay within either the investment-grade universe or the high-yield universe. For example, insurance companies face much higher capital charges when holding HY bonds: 90% of insurance companies' portfolios are in IG bonds. Mutual funds are set up as either investment-grade funds or high-yield funds. The fund mandates limit the managers from investing in the other category of bonds. As a result of these frictions, we see strong investor segmentation around the IG/HY cutoff. While mutual funds hold only about 15% of BBB-rated bonds, they hold about 22% for BB+ rated bonds. The difference in investor base likely contributes to the discrete changes in the portfolio multipliers estimated.

In addition to ratings, we also investigate heterogeneity along the maturity dimension. Specifically, we group portfolios into short, medium, and long maturity groups: 0-3 years, 4-10 years, and 10 years and above. Figure 7 shows longer maturity portfolios have slightly higher portfolio multipliers, although the difference is not statistically significant. Short-term portfolios have a multiplier close to 0; medium-term and long-term portfolios have a higher multiplier around 0.3-0.4. The heterogeneous multipliers for portfolios with different maturities are likely linked to the clientele effects and segmentation along the yield curve (Vayanos and Vila, 2021; Kekre et al., 2022).

Figure 7: Maturity Heterogeneity ($N = 76,348$)



The graph shows how the detailed rating \times quarter to maturity multiplier estimates differ across maturity. The estimates are presented relative to the short-term maturity elasticity. The bars reflect 95% confidence intervals. As the estimates suggest we can reject the null of no heterogeneity at the 95% level.

5 Arbitrage Risk

To better understand the quantitative implications of the multipliers we find in the previous section, we look at the Sharpe ratio of a strategy that takes advantage of the price deviation between the testing asset and its substitute portfolio at the one-quarter horizon. For example, if portfolio A's price is too high relative to its substitute portfolio, arbitrageurs could make money by selling portfolio A and purchasing its substitute. However, if substitute returns do not fully replicate portfolio A's returns, arbitrageurs face tracking error risk, which could result in losses. In other words, this is a near-arbitrage opportunity as the trade contains some risk. Specifically, we consider a strategy that sells the more expensive one between asset j and its close-substitute asset and buys the cheaper one. We unwind this position in the following period, given that we find the price gap disappears by the end of the next period on average.²⁸

The expected return of this strategy is $M \times \text{Mean}(|u_{j,t}|)$, where M is the multiplier estimated in Section 4 and $|u_{j,t}|$ is the size of demand shocks in absolute terms. The volatility of this return is the arbitrage risk, defined as

$$\text{ArbRisk}_j \equiv \text{std}(\tilde{v}_{j,t}^S + Mu_{j,t}). \quad (14)$$

The arbitrage risk comes from non-flow shocks as well as the next period (observable) demand shocks ($u_{j,t}$). Since we have removed all the predictable components in constructing u , the future period u is unknown to the arbitrageur when carrying out the strategy and should be considered as part of the arbitraging risk.

The average Sharpe ratio of this arbitrage activity is then given by

$$SR = \frac{M \times \text{Mean}(|u_{j,t}|)}{\text{Mean}(\text{ArbRisk}_j)}. \quad (15)$$

We find the annualized Sharpe ratios across different aggregation levels are generally small, ranging from 0.006 at the CUSIP level to 0.28 at the most aggregated rating level.²⁹

²⁸The Sharpe ratios are calculated under the implicit assumption that different bonds and portfolios are segmented, and the arbitrageur buys and sells one bond/portfolio at a time. In other words, we are not considering the diversification benefit if the arbitrageur operates across multiple submarkets.

²⁹If we subset to large shocks only, the Sharpe ratios for portfolios at different aggregation levels also fall within this range. The main reason is that the large shocks are accompanied by large noises as well.

The Sharpe ratio increases with the aggregation level, mainly because the multipliers are larger for more aggregated portfolios. However, even the maximum Sharpe ratio is small, and it is before considering potential shorting costs.³⁰ Our results reveal that the risks of engaging in these arbitrages are high relative to the average gain, which potentially explains why we observe these price deviations exist in the data.

Table 7: Arbitrage Risk and Portfolio Multipliers

	M	Arb. risk	Sharpe ratio
CUSIP	.052	.043	.006
Det. rating \times Quarter to maturity	.348	.055	.026
Rating \times ST/MT/LT	1.229	.037	.134
Rating \times ST/LT	1.591	.048	.142
Rating	3.507	.043	.280

The table presents the estimated multiplier (M), the average arbitrage risk (as defined in Equation 14) and the implied annualized Sharpe ratio for portfolios at different aggregation levels.

5.1 Multipliers and Arbitrage Risk

Motivated by the low Sharpe ratios, we explore to what extent arbitrage risk limits investors to substitute across assets. Our arbitrage risk is the part of the risk that cannot be hedged away by constructing the substitute portfolio, and hence it acts as a limit to arbitrage (Wurgler and Zhuravskaya, 2002). Intuitively, the higher the arbitrage risk, the less arbitrageurs are willing to substitute to alternative portfolios and hence higher multipliers. In Appendix G, we provide a simple two-asset case illustrating this intuition.

We use our baseline portfolio (formed by bonds with the same detailed rating group and quarter to maturity) for analysis in this section. We first sort our portfolios into four groups based on the size of their arbitrage risk defined in equation (14): group 1 includes portfolios with the lowest arbitrage risk, and group 4 includes portfolios with the highest arbitrage risk. We run the baseline specification in Section 4.1 including interaction terms of group indicators and demand shocks ($u_{j,t}$). In other words, we allow the portfolio multiplier M to

³⁰Other papers have found strategies that predict bond returns with much bigger Sharpe ratios (see, e.g., Bartram et al. 2020).

Table 8: Arbitrage Risk and Portfolio Multipliers

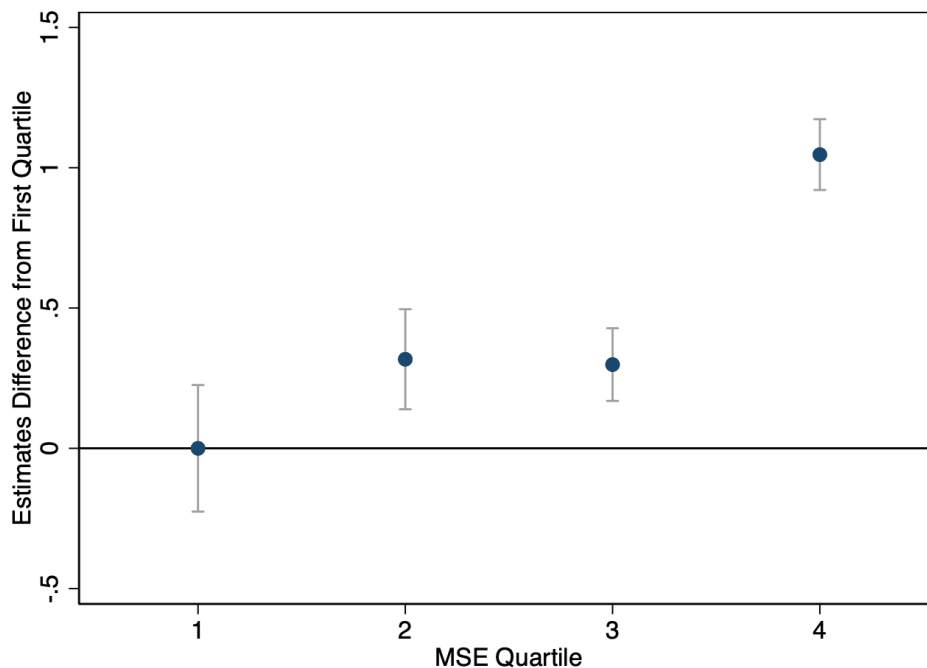
	(1)	(2)
Shock	0.35*** (0.04)	0.03 (0.04)
ArbRisk x Shock		0.45*** (0.03)
Substitute return	0.90*** (0.04)	0.91*** (0.04)
Quarter FE	Yes	Yes
Drop Crisis	Yes	Yes
N	76,348	76,348
First-stage F-statistic	2544.04	2562.86

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table presents estimates for how the detailed rating \times quarter to maturity multiplier estimates depend on the arbitrage risk of the portfolio. We standardized the arbitrage risk measurement by subtracting the mean and dividing it by the cross-sectional standard deviation of the arbitrage risk. Column (1) estimates our baseline specification for the post-winsorized sample. Column (2) includes an interaction term between the demand shocks and arbitrage risk. Robust standard errors are clustered at the substitute group \times time level.

Figure 8: Arbitrage Risk and Portfolio Multipliers ($N = 76,348$)



The graph shows the detailed rating \times quarter to maturity multiplier estimates in different arbitrage risk quartiles. The estimates are presented relative to the lowest arbitrage risk quartile. The bars reflect 95% confidence intervals. As the estimates suggest, we can reject the null of no heterogeneity at the 95% level.

differ by group. We plot the relative magnitudes of the estimates for each group in Figure 8.

We indeed find that portfolios with higher arbitrage risk have higher portfolio multipliers and more inelastic demand. We obtain an almost monotone pattern of multipliers across groups. Groups with higher arbitrage risk have significantly higher portfolio multipliers. For group 1, the estimate for the multiplier is close to 0, suggesting almost frictionless arbitrage between bonds in this group and their substitutes. However, for group 4, the estimate for the multiplier is much larger. [Wurgler and Zhuravskaya \(2002\)](#) find similar patterns in the stock market using index inclusion and exclusion events. We verify that such pattern exists more broadly outside of the index-related events.

The relationship between arbitrage risks and portfolio multipliers is also verified in a regression specification. We add an interaction term between $ArbRisk_j$ and demand shock

u in the baseline regression specification,

$$\Delta p_{j,t} = M_0 u_{j,t} + M_1 u_{j,t} \times ArbRisk_j + \tilde{M} \Delta p_{j,t}^{sub} + \epsilon_{j,t} \quad (16)$$

Table 8 presents the results. The coefficient in front of the interaction term is significantly positive. A one-standard-deviation increase in arbitrage risk increases the portfolio multiplier by 0.45. The effect is particularly strong among the high-yield bonds, which have higher overall risks and also higher idiosyncratic risks.

6 Conclusion

In this paper, we study a wide range of multipliers in the corporate bond market utilizing exogenous demand shocks from mutual fund flows. Studying the corporate bond market is interesting not only because it serves as an important channel for the real economy but also because it is different from the equity market in terms of market structure, investor composition, and correlation of asset payoffs. Comparing the two markets can shed light on the key drivers of market inelasticity.

We emphasize that it is important to account for the correct substitute portfolios when estimating the multipliers. Different from existing methods in the literature, we allow for close and distant substitutes in our estimation. Ignoring the heterogeneous cross-elasticity leads to underestimation of demand elasticities. At the CUSIP-level, the estimated multiplier drops from 0.3 to essentially 0 (both economically and statistically) once we allow for certain bonds to be closer substitutes than others. Relative to equities, we find that individual bonds are much more elastic.

While individual bonds are quite elastic, portfolios of bonds face more inelastic demand. The portfolios formed by rating categories and maturity have a multiplier around 0.35. Furthermore, as the portfolio becomes more aggregated, the portfolio-level multiplier increases and the substitute passthrough decreases, indicating portfolios are less substitutable at more aggregated levels. Finally, we find that riskier portfolios and long-term portfolios have higher multipliers; as do portfolios with larger arbitrage risks. Our results imply that both investor segmentation and arbitrage risk contribute to inelastic demand at the portfolio level. We leave the quantification of different channels to future research.

Finally, we find that the price impact at the portfolio level almost disappears in the next quarter, suggesting that long-run demand is more elastic than short-run demand. Investigating the dynamics of the price impact is interesting and can be informative about the underlying arbitrage frictions. For future work, it would be interesting to link the demand elasticities at different horizons to different arbitrage theories.

References

- Bai, Jennie, Turan G. Bali, and Quan Wen**, “Common risk factors in the cross-section of corporate bond returns,” *Journal of Financial Economics*, 3 2019, 131 (3), 619–642.
- Bartram, Söhnke M., Mark Grinblatt, and Yoshio Nozawa**, “Book-to-Market, Mispricing, and the Cross-Section of Corporate Bond Returns,” *Working paper*, 12 2020.
- Beneish, Messod D. and Robert E. Whaley**, “An Anatomy of the ”S & P Game”: The Effects of Changing the Rules,” *The Journal of Finance*, 12 1996, 51 (5), 1909.
- Bretscher, Lorenzo, Lukas Schmid, Ishita Sen, and Varun Sharma**, “Institutional Corporate Bond Pricing,” *SSRN Electronic Journal*, 1 2022.
- Calomiris, Charles W., Mauricio Larrain, Sergio L. Schmukler, and Tomas Williams**, “Large international corporate bonds: Investor behavior and firm responses,” *Journal of International Economics*, 7 2022, 137, 103624.
- Chang, Yen-Cheng, Harrison Hong, Inessa Liskovich, Yen-Cheng Chang, Harrison Hong, and Inessa Liskovich**, “Regression Discontinuity and the Price Effects of Stock Market Indexing,” *Review of Financial Studies*, 1 2015, 28 (1), 212–246.
- Chaudhry, Aditya**, “Do Subjective Growth Expectations Matter for Asset Prices?,” *SSRN Electronic Journal*, 9 2022.
- Chen, Honghui, Gregory Noronha, and Vijay Singal**, “The Price Response to S&P 500 Index Additions and Deletions: Evidence of Asymmetry and a New Explanation,” 2004.
- Choi, Jaewon, Saeid Hoseinzade, Sean Seunghun Shin, and Hassan Tehranian**, “Corporate bond mutual funds and asset fire sales,” *Journal of Financial Economics*, 11 2020, 138 (2), 432–457.
- Coppola, Antonio**, “In Safe Hands: The Financial and Real Impact of Investor Composition Over the Credit Cycle,” 2022.
- Coval, Joshua and Erik Stafford**, “Asset fire sales (and purchases) in equity markets,” *Journal of Financial Economics*, 11 2007, 86 (2), 479–512.

- Darmouni, Olivier, Kerry Y Siani, and Kairong Xiao**, “Nonbank Fragility in Credit Markets: Evidence from a Two-Layer Asset Demand System *,” 2023.
- Davis, Steven J. and John Haltiwanger**, “Gross job creation, gross job destruction, and employment reallocation,” *Quarterly Journal of Economics*, 1992, *107* (3), 819–863.
- Edmans, Alex, Itay Goldstein, and Wei Jiang**, “The Real Effects of Financial Markets: The Impact of Prices on Takeovers,” *The Journal of Finance*, 6 2012, *67* (3), 933–971.
- Ellul, Andrew, Chotibhak Jotikasthira, and Christian T Lundblad**, “Regulatory pressure and fire sales in the corporate bond market \$,” 2011.
- Fang, Chuck**, “Monetary Policy Amplification through Bond Fund Flows,” *Working paper*, 2022.
- Gabaix, Xavier and Ralph S J Koijen**, “In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis,” Technical Report 2021.
- and –, “Granular Instrumental Variables,” 2022.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik Instruments: What, When, Why, and How,” *American Economic Review*, 8 2020, *110* (8), 2586–2624.
- Goldstein, Itay, Hao Jiang, and David T. Ng**, “Investor flows and fragility in corporate bond funds,” *Journal of Financial Economics*, 12 2017, *126* (3), 592–613.
- Greenwood, Robin, Nber Toomas Laarits, and Jeffrey Wurgler**, “Stock Market Stimulus,” 2022.
- Haddad, Valentin, Paul Huebner, and Erik Loualiche**, “How Competitive is the Stock Market? Theory, Evidence from Portfolios, and Implications for the Rise of Passive Investing,” 2022.
- Harris, Lawrence and Eitan Gurel**, “Price and Volume Effects Associated with Changes in the S&P 500 List: New Evidence for the Existence of Price Pressures,” *The Journal of Finance*, 9 1986, *41* (4), 815.
- Hartzmark, Samuel M and David H Solomon**, “Predictable Price Pressure,” 2022.

- Holm-Hadulla, Frédéric and Matteo Leombroni**, “Heterogeneous Intermediaries and Bond Characteristics in the Transmission of Monetary Policy *,” 2023.
- Hubbard, Glenn, Donald Kohn, Laurie Goodman, Kathryn Judge, Anil Kashyap, Ralph Koijen, Blythe Masters, Sandie O’connor, and Kara Stein**, “Task Force on Financial Stability,” Technical Report 6 2021.
- Jansen, Kristy A.E.**, “Long-term Investors, Demand Shifts, and Yields,” *SSRN Electronic Journal*, 8 2021.
- Jotikasthira, Chotibhak, Christian Lundblad, Tarun Ramadorai, John Campbell, Cam Harvey, Donghui Li, Ludovic Phalippou, Michael Schill, Ajay Shah, and Dimitri Vayanos**, “Asset Fire Sales and Purchases and the International Transmission of Funding Shocks,” *The Journal of Finance*, 12 2012, 67 (6), 2015–2050.
- Kekre, Rohan, Moritz Lenel, and Federico Mainardi**, “Monetary Policy, Segmentation, and the Term Structure *,” 2022.
- Koijen, Ralph and Motohiro Yogo**, “Exchange Rates and Asset Prices in a Global Demand System Exchange Rates and Asset Prices in a Global Demand System *,” 2019.
- Koijen, Ralph S.J. and Motohiro Yogo**, “A Demand System Approach to Asset Pricing,” <https://doi.org/10.1086/701683>, 8 2019, 127 (4), 1475–1515.
- , **François Koulischer, Benoît Nguyen, and Motohiro Yogo**, “Euro-Area Quantitative Easing and Portfolio Rebalancing,” *American Economic Review*, 5 2017, 107 (5), 621–27.
- Li, Jennifer, Neil D Pearson, and Qi Zhang**, “Impact of Demand Shocks on the Stock Market: Evidence from Chinese IPOS ,” 2021.
- Li, Jiacui**, “What Drives the Size and Value Factors?,” *SSRN Electronic Journal*, 6 2021.
- **and Zihan Lin**, “Prices Are Less Elastic For Less Diversifiable Demand,” *SSRN Electronic Journal*, 6 2022.
- Lou, Dong**, “A Flow-Based Explanation for Return Predictability,” *The Review of Financial Studies*, 12 2012, 25 (12), 3457–3489.

- Manconi, Alberto, Massimo Massa, and Ayako Yasuda**, “The role of institutional investors in propagating the crisis of 2007–2008,” *Journal of Financial Economics*, 6 2012, 104 (3), 491–518.
- Pavlova, Anna and Taisiya Sikorskaya**, “Benchmarking Intensity,” *The Review of Financial Studies*, 8 2022.
- Peng, Cameron and Chen Wang**, “Factor Demand and Factor Returns *,” 2019.
- Petajisto, Antti**, “Why Do Demand Curves for Stocks Slope Down?,” *Source: The Journal of Financial and Quantitative Analysis*, 2009, 44 (5), 1013–1044.
- Schmickler, Simon and Pedro Tremacoldi-Rossi**, “Spillover Effects of Payouts on Asset Prices and Real Investment,” *SSRN Electronic Journal*, 10 2022.
- Shleifer, Andrei**, “Do Demand Curves for Stocks Slope Down?,” *The Journal of Finance*, 7 1986, 41 (3), 579.
- Siani, Kerry**, “Raising Bond Capital in Segmented Markets,” *SSRN Electronic Journal*, 10 2022.
- Siani, Kerry Y**, “Raising Bond Capital in Segmented Markets,” 2022.
- Vayanos, Dimitri and Jean-Luc Vila**, “A Preferred-Habitat Model of the Term Structure of Interest Rates,” *Econometrica*, 1 2021, 89 (1), 77–112.
- Vissing-Jorgensen, Annette and Arvind Krishnamurthy**, “The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy,” *Brookings working paper*, 2011.
- Wurgler, Jeffrey and Ekaterina Zhuravskaya**, “Does Arbitrage Flatten Demand Curves for Stocks?,” *Journal of Business*, 10 2002, 75 (4), 583–608.

A Identification

A.1 Identification of Multipliers

We first consider the specification in Equation (9), restated here:

$$\Delta p_{j,t} = M u_{j,t} + \text{Fixed effects}_{g,t} + \tilde{\nu}_{j,t}.$$

To identify the coefficients correctly, we need the following moment conditions to be satisfied:

$$\mathbb{E}[u_{j,t} \nu_{j,t}] = 0. \tag{17}$$

We now prove that Assumption 1, i.e. $\nu_{jt} \perp S_{ij,t-1}$ for all fund i and time t , is a sufficient condition for satisfying these moment conditions, where $S_{ij,t-1}$ is the lagged market share of mutual fund i . For clarity, we use the notation $\mathbb{E}_i^{(j)}$ to denote the expectation over dimension j conditional on dimension i .

Proof.

$$\begin{aligned} \mathbb{E}^{(j,t)} [u_{j,t} \nu_{j,t}] &= \mathbb{E}^{(j,t)} \left[\mathbb{E}_{j,t}^{(i)} [S_{ij,t-1} u_{i,t}] \nu_{j,t} \right] && \text{(Definition of } u_{j,t}\text{)} \\ &= \mathbb{E}^{(i,j,t)} [S_{ij,t-1} u_{i,t} \nu_{j,t}] && \text{(L.I.E.)} \\ &= \mathbb{E}^{(i,t)} \left[u_{i,t} \mathbb{E}_{i,t}^{(j)} [S_{ij,t-1} \nu_{j,t}] \right] && \text{(conditional Exp.)} \\ &= \mathbb{E}^{(i)} [u_{i,t} 0]. && \text{(I.D. assum.)} \end{aligned}$$

□

The first equality is the definition of bond-level shocks as the aggregation across fund-level shocks. The second line directly follows through the law of iterated expectations. The third equality applies the law of iterated expectations again, and this time we take the expectations across bonds *within each fund*. As shocks $u_{i,t}$ are at the fund level, it can be pulled out of the expectations. This is the key property exploited in the exogenous-share design. Notice that here we do not assume orthogonal flows. The last equality uses the identification assumption.

A.2 Identification of Substitute Passthrough

We estimate the passthrough according to the following equation:

$$\Delta p_{j,t} = M u_{j,t} + \tilde{M} \Delta p_{j,t}^{sub} + \tilde{\nu}_{j,t},$$

where we instrument $\Delta p_{j,t}^{sub}$ with $u_{j,t}^{sub}$, constructed as the value-weighted group-level shocks:

$$u_{j,t}^{sub} \equiv \sum_k w_{k|g(j),t-1} u_k,$$

where $g(j)$ is the group of close substitutes for bond j , and $w_{k|g(j),t-1} = \frac{\mathbb{I}_{k \in g(j)} w_{k,t-1}}{w_{g(j),t-1}}$ is the (lagged) market share of bond k conditional on group $g(j)$. The required moment condition is:

$$\mathbb{E}[u_{j,t}^{sub} \nu_{j,t}] = 0.$$

The proof is slightly more complicated due to the weighting in u^{sub} , but essentially follows a similar strategy. We suppress the time subscript in the proof below for ease of notation.

Proof. Plug in the definitions and use the law of iterated expectations as before, we have:

$$\mathbb{E}^{(j)} [u_j^{sub} \nu_j] = \mathbb{E}^{(j)} [w_{k|g(j)} u_k \nu_j] \tag{18}$$

$$= \mathbb{E}^{(j)} [w_{k|g(j)} \mathbb{E}_k^{(i)} [u_i S_{ik}] \nu_j] \tag{19}$$

$$= \mathbb{E}^{(i)} [\mathbb{E}_i^{(k)} [\mathbb{E}_{ik}^{(j)} [w_{k|g(j)} \nu_j] S_{ik}] u_i]. \tag{20}$$

Notice that the inner expectation is virtually the group-level idiosyncratic shocks:

$$\mathbb{E}_{ik}^{(j)} [w_{k|g(j)} \nu_j] = \mathbb{E}_{ik}^{(j)} \left[\frac{\mathbb{I}_{k \in g(j)} w_k}{w_{g(j)}} \nu_j \right] = w_k \underbrace{\mathbb{E}_{ik,j \in g(k)}^{(j)} \nu_j}_{\nu_g}.$$

Plug it back into Equation 20, we have:

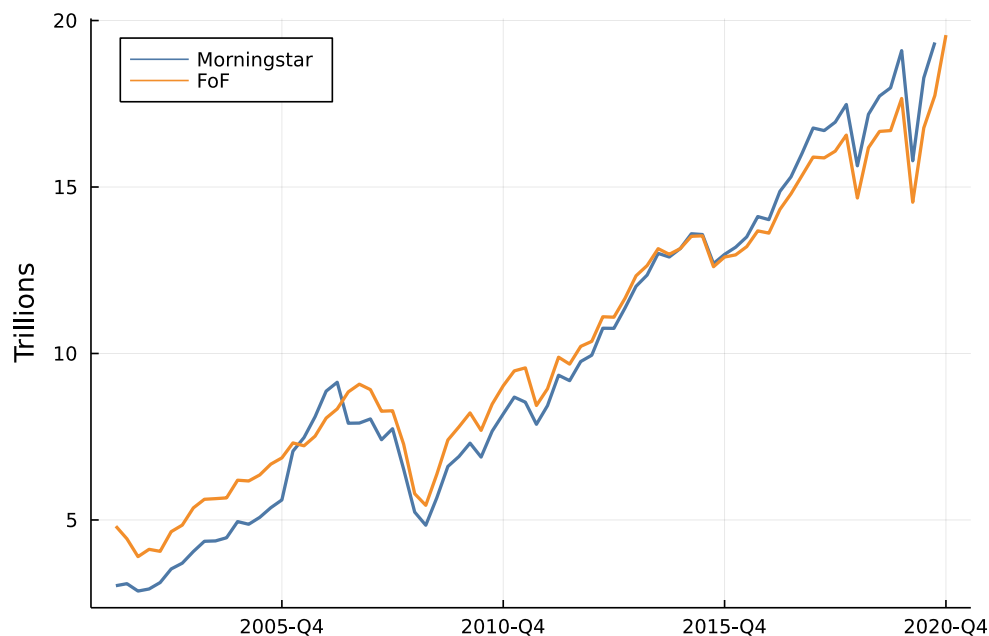
$$\begin{aligned}
 \mathbb{E}_i^{(k)} \left[\mathbb{E}_{ik}^{(j)} [w_{k|g(j)} \nu_j] S_{ik} \right] &= \mathbb{E}_i^{(k)} [w_k S_{ik} \nu_g] \\
 &= \mathbb{E}_i^{(g)} \left[\underbrace{\mathbb{E}_g^{(k)} [w_k S_{ik}]}_{S_{i,g}} \nu_g \right] \\
 &= 0.
 \end{aligned}$$

Here we rely on the exogenous share assumption at the group level, i.e., $S_{i,g} \perp \nu_g$. □

B Morningstar Data Coverage

We find that Morningstar mutual fund dataset coverage is quite extensive and closely lines up with the coverage of the Flow of Funds estimate of mutual fund holdings.

Figure 9: Total Assets of All Mutual Funds in Flow of Funds vs. Morningstar



This figure plots the net total assets of all mutual funds reported to Morningstar (blue) against total mutual fund shares outstanding in the US, retrieved from Table L.244 from Financial Accounts (Flow of Funds).

C Robustness of Shock Construction

In this section, we explore the sensitivity of our results to various deviations from our baseline assumptions used in the shock construction process.

C.1 Assessing passthrough of fund flows into trading

In constructing our shocks, we assumed that mutual funds mechanically reinvest inflows according to their existing portfolio weights. This assumption implies a one-to-one passthrough from inflows into trades. In this section, we kick the tires on this assumption and find that while the passthrough is not precisely one-to-one, it is still relatively high, around 50% at the individual bond level and 60% at our baseline level of aggregation. Overall, the less than one-to-one passthrough suggests that our multiplier estimates are potentially attenuated.

However, we can find upper bounds on the multipliers using our passthrough estimates. The upper bound estimates are still considerably smaller than those typically found for stocks. At the bond level, the upper bound multiplier estimate is 0.08, and for our baseline level of aggregation, it is around 0.33.

Instead, we run the following specification ignoring the price response,

$$\Delta q_{ijt} = \alpha + \beta f_{it} + \varepsilon_{ijt}$$

where $\Delta q_{ijt} = 2 \frac{Q_{ijt} - Q_{ij,t-1}}{Q_{ijt} + Q_{ij,t-1}}$ is the Davis and Haltiwanger (1992) change in mutual fund holdings, and f_{it} is the funds flows as a percentage of lagged AUM.³¹ Naturally, our β estimates are biased due to omitting the prices. However, since mutual fund demand curves are presumably downward sloping, we can sign the bias as negative. As a result, our β estimate is a lower bound of the true passthrough.

Table 9 shows that passthrough is not equal to one but is still sizeable and highly significant. At the individual bond (CUSIP) level, the passthrough is around 0.51. This estimate is smaller than what Lou (2012) found for individual stocks—0.62 for inflows and 0.97 for outflows. The lower passthrough for individual bonds is, however, somewhat expected. It is presumably easier to find another GM stock, than it is to find another 34-quarter GM bond maturing in 12 months. Consistent with the difficulty in finding exact bond CUSIPs, the passthrough increases at more aggregate levels—it is easier to find another BBB bond than a specific CUSIP. At the detail rating-quarter to maturity level, the passthrough rises to around 0.58. At the rating letter level, the passthrough is closer to 0.95.

Table 10 separately assesses the passthrough for inflows and outflows. Two things stand out. Firstly, we see higher passthrough at more aggregate levels. Secondly, the passthrough is higher for outflows than inflows, echoing Lou (2012) findings for equities.

Overall, these passthrough estimates suggest that our multiplier estimates are potentially attenuated. By assuming a passthrough of 1, we are overestimating the size of the demand shock and hence underestimating the multiplier. Since our passthrough estimates are lower bounds for the actual passthrough, we can use them to estimate the upper bounds on the true multiplier. If our lower bound passthrough estimate is $\hat{\beta} < 1$, then the true multiplier

³¹Davis and Haltiwanger (1992) change is a second order approximation of the log difference, and has several desirable properties over the standard percentage change measure, for example, it is symmetric, and it does not generate outliers due to small base effects.

Table 9: Estimating passthrough from flows into holdings

	CUSIP		Det Rating x Q to Mat		Rating	
	(1)	(2)	(3)	(4)	(5)	(6)
Flows	0.51*** (0.02)	0.49*** (0.01)	0.58*** (0.01)	0.57*** (0.01)	0.94*** (0.04)	0.91*** (0.05)
Time + Fund FE	No	Yes	No	Yes	No	Yes
<i>N</i>	4,861,780	4,861,779	1,193,197	1,193,177	163,515	163,502

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table summarizes the estimates of the passthrough of funds flows into trading for different levels of portfolio aggregation. In terms of levels of aggregation the columns (1) and (2) contain estimates for security-level passthrough, columns (3) and (4) contain passthroughs at the detailed-rating x quarter-to-maturity level of aggregation, and columns (5) and (6) contain rating level portfolio passthroughs.

Table 10: Estimating passthrough from flows into holdings

	CUSIP		Det Rating x Q to Mat		Rating	
	(1)	(2)	(3)	(4)	(5)	(6)
Flows	0.45*** (0.05)	0.66*** (0.05)	0.55*** (0.02)	0.79*** (0.04)	0.94*** (0.03)	1.08*** (0.13)
Time + Fund FE	Yes	Yes	Yes	Yes	Yes	Yes
Flow Type	Inflow	Outflow	Inflow	Outflow	Inflow	Outflow
<i>N</i>	2,893,530	1,968,240	690,300	502,855	84,099	79,391

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table summarizes the estimates of the passthrough of funds flows into trading for different levels of portfolio aggregation, separating the samples into inflows and outflows. In terms of levels of aggregation the columns (1) and (2) contain estimates for security-level passthrough, columns (3) and (4) contain passthroughs at the detailed-rating x quarter-to-maturity level of aggregation, and columns (5) and (6) contain rating level portfolio passthroughs.

must be $M \leq \frac{1}{\beta} \hat{M}$, where \hat{M} is the multiplier we estimated in section 4.1. Accordingly, this calculation suggests that the true CUSIP multiplier could be as high as 0.08, and our detailed rating-quarter to maturity could be as high as 0.33. While these upper bounds are higher than our estimates, they are still considerably smaller than 1, the magnitude typically found for equities.

C.2 Non-parametric estimation of flow factor structure

In this section, we implement the factor models for flow shocks. We first show that flow shocks have weak factor structures in this section; in the next section we explicitly remove the common factors and show our results are robust.

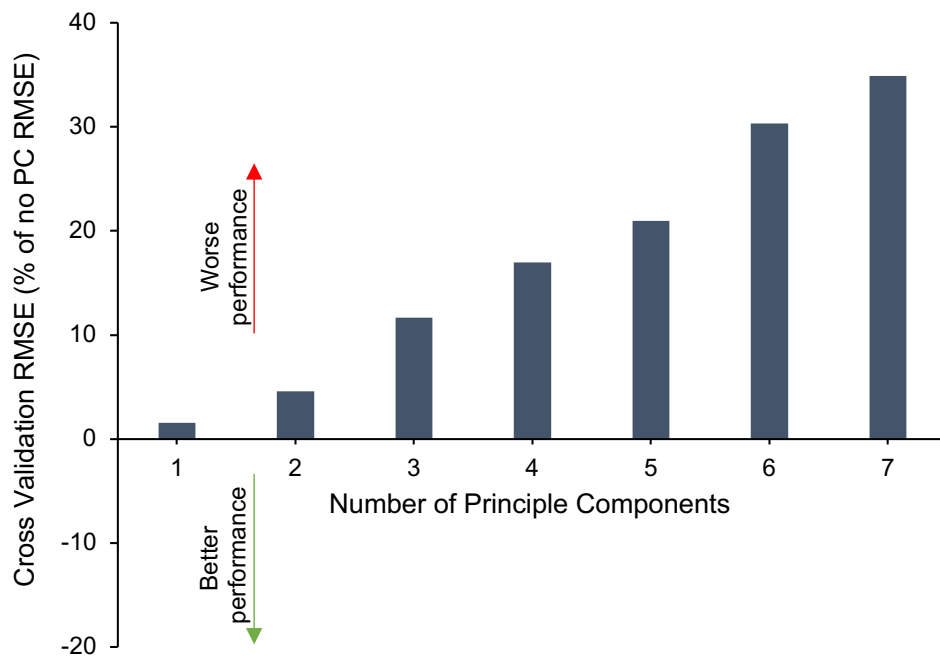
We first implement the non-parametric approaches to estimate common factors. Due to missing values in our data, we estimate the factor model using alternating least squares (ALS). We find the data exhibits a relatively weak factor structure, with the first factor only explaining around 8% of the variation in mutual fund flows.

With that said, non-parametric factor estimation methods are known to have poor finite sample performance when there are a lot of missing value. In our context, due to mutual fund entry and exit, we have a highly unbalanced panel of flow data—relative to a fully balanced panel we are “missing” 49.8% of observations. Hence, there is a risk that the factor model is simply fitting noise rather than identifying actual factors.

To assess this overfitting risk, we run a five-fold cross-validation exercise. The procedure splits the sample into five subsamples, estimates the factor model on one of them, and then assess the models performance on the four subsamples left out. It then repeats this five times and calculates the average performance of the method.

Figure 10 shows the average root mean-squared errors from the five-fold cross-validation exercise. The results suggests the factors are likely being fit on noise. In fact, adding more factors seems to make the out-of-sample performance even worse. While this finding does not suggest that there is a strong factor structure in flows, it does suggest that non-parametric factor estimation may not be able to extract it in our setting with a highly unbalanced panel. Due to this concern, in the next subsection we explore a parametric approach to estimating factors—the additional structure should help reduce the risk of overfitting.

Figure 10: Cross-validation of non-parametric factor estimation



This figure shows the average out-of-sample root mean-squared errors from the five-fold cross-validation exercise, relative to the average out-of-sample root-square error for a model with no principle components. The exercise is repeated for models with a larger number of principle components.

C.3 Parametric estimation of flow factor structure

Due to the poor performance of the non-parametric method in estimating factors, in this section we follow a parametric approach below. Specifically, let the data-generating process of innovation be:

$$\epsilon_{i,t} = \delta_t + \Lambda_{i,t}\eta_t + u_{i,t}, \quad (21)$$

where δ_t is the time fixed effects, $\Lambda_{i,t}\eta_t$ is the contribution from common factors, and $u_{i,t}$ is the desired idiosyncratic demand shocks.

Following the common approach in asset pricing, we assume the factor loading are a linear function of characteristics, $\Lambda_{i,t} = C'_{i,t}\lambda$, where $C_{i,t}$ is a vector of observable characteristics of fund i , including (lagged) log AUM of the firm, the share in high-yield bonds, and the average duration in the portfolio, and λ is a constant vector. The data generating process of $\epsilon_{i,t}$ under this parameterization is then:

$$\epsilon_{i,t} = \delta_t + C'_{i,t}(\lambda\eta_t) + u_{i,t}. \quad (22)$$

Notice (22) can be estimated by running a panel regression of $\epsilon_{i,t}$ on $C_{i,t}$ with time-varying coefficients and time fixed effects. In the baseline where no additional factors are controlled, we simply regress $\epsilon_{i,t}$ on time fixed effects. The estimated residual is recovered as:

$$\hat{u}_{i,t} = \epsilon_{i,t} - \hat{\delta}_t - C'_{i,t}(\widehat{\lambda\eta}_t) \quad (23)$$

To minimize noise due to extreme outliers and volatile funds, we winsorize innovations $\epsilon_{i,t}$ at the 5% level before estimating (22). Importantly, the winsorized innovations are only used in the estimation of coefficients. In (23) we use the original $\epsilon_{i,t}$ to recover the demand shock $\hat{u}_{i,t}$, as outliers are also valid idiosyncratic shocks to demand.

C.4 Lagging market shares for shock construction

Another potential concern to the exogenous share condition is that some mutual fund managers have superior information and therefore can front-run the market, leading to a positive correlation between market share and bond returns. This issue can be addressed with fur-

ther lagging the market share used in aggregating fund flow shocks: it is highly unlikely for a mutual fund manager to predict *idiosyncratic* returns at the bond level in one year and allocate their portfolio in one year to benefit from it.³²

Table 11 and 12 report our baseline results using shocks aggregated with one-year-lag market shares. The results remain largely unchanged.

C.5 Sensitivity of multiplier estimation to shock construction specifications

Below we explore how sensitive our multiplier estimates are to changing various choices we made in the shock construction process. Overall, our estimates do not seem to change very much if we (i) winsorize extreme values and use inverse variance weighting in calculating fixed effects, (ii) control for additional factors as described in C.3, or (iii) use different specifications for our AR regressions. Furthermore, we also find that our substitute specification does a quite good job at control for substitutes, as the results are unchanged if we instead include detailed rating \times quarter fixed effects.

C.6 Multipliers and shock sizes

In this section, we estimate the magnitudes of the bond-level and portfolio-level multipliers for bonds that experienced large shocks. There are two reasons why this might be interesting. First, the effect of the demand shocks may be non-linear, and one might expect large shocks generate larger price responses than small shocks. Second, the bonds with large demand shocks are likely to belong to funds that experienced large fund flows, in which case the passthrough assumption is more likely to hold. The CUSIP-level results are presented in Table 14, and the portfolio-level results are presented in Table 15. The estimates are quantitatively similar to what we have in the main text.

³²It might be more plausible if they are able to predict the dynamics of some systematic factors in one year and trade accordingly—the common factor issue is addressed in Section by removing the common factors from flows.

Table 11: CUSIP specification lagging holding shares by one-year

	Homo. OLS	OLS		First-stage	2SLS
	(1)	(2)	(3)	(4)	(5)
Shock	0.39*** (0.07)	-0.03 (0.06)	0.01 (0.05)		0.02 (0.06)
Substitute return					1.08*** (0.04)
Group Shock				3.93*** (0.47)	
Quarter FE	Yes	Yes	Yes	Yes	Yes
Group x Quarter FE	No	Yes	Yes	No	No
Drop Crisis	No	No	Yes	Yes	Yes
N	277,336	277,336	261,144	261,144	261,144
R^2	0.21	0.40	0.37	0.62	
First-stage F-statistic					94.96

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table shows the multiplier estimates from repeating the baseline CUSIP level specification for shocks constructed by lagging holdings shares by one-year. Column 1 is the OLS estimates from regressing bond returns on the demand shock, controlling for time fixed effects, as in equation (5)—this specification corresponds to a model in which we assume homogeneous cross-elasticity with all other bonds. Columns 2-4 show results from running the regression in equation (9) with different definitions of close-substitutes. Columns 2 and 3 directly control for close-substitute prices using detailed rating x time fixed effects. Column 4 additional controls for maturity (long-term/short-term) x quarter fixed effects to control for potentially omitted time-varying maturity factor in holding shares. Column 5 and 6 relate to the IV specification specified in equation (8). We regress bond returns on the demand shock and substitute returns, while controlling for time fixed effects (and additional controls depending on the specification). We instrument for substitute returns using demand shocks to substitute assets. The IV specification corresponds to the model that allows for heterogeneous cross-elasticities between close and distant substitutes. The parentheses contain standard errors clustered at the substitute group x time level.

Table 12: Quarter to Maturity x Detailed Rating specification lagging holding shares by one-year

	Homo. OLS	OLS		First-stage	2SLS
	(1)	(2)	(3)	(4)	(5)
Shock	1.04*** (0.15)	0.31** (0.12)	0.33** (0.12)		0.33** (0.12)
Substitute return					0.96*** (0.04)
Group Shock				4.36*** (0.51)	
Quarter FE	Yes	Yes	Yes	Yes	Yes
Group x Quarter FE	No	Yes	Yes	No	No
Drop Crisis	No	No	Yes	Yes	Yes
N	81,866	81,866	77,387	76,348	76,348
R^2	0.22	0.47	0.41	0.48	
First-stage F-statistic					69.48

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table presents the quarter-to-maturity \times detailed rating level multiplier estimates for demand shocks constructed by lagging holdings shares by one-year. Column 1 is the estimates from regressing portfolio returns on the demand shock, controlling for time-fixed effects, as in equation (9). This specification corresponds to a model in which we assume homogeneous cross-elasticity with all other bonds. Columns 2-4 show results from running the regression in equation (9) with different definitions of close-substitutes. Columns 2 and 3 directly control for close-substitute prices using detailed rating \times time fixed effects. Column 4 additional controls for maturity (long-term/short-term) \times quarter fixed effects to control for potentially omitted time-varying maturity factor in holding shares. Columns 5 and 6 relate to the IV specification in equation (8), in which we regress portfolio returns on the demand shock and substitute returns, while controlling for time fixed effects (and additional controls depending on the specification). We instrument for substitute returns using demand shocks to substitute assets. The IV specification corresponds to the model that allows for heterogeneous cross-elasticities between close and distant substitutes. The parentheses contain standard errors clustered at the substitute group \times time level.

Table 13: Shock robustness

	Bench.	Robustness						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Shock	0.33*** (0.10)	0.29** (0.10)	0.32** (0.11)	0.37*** (0.11)	0.26** (0.08)	0.28*** (0.08)	0.21*** (0.06)	0.18 (0.13)
AR lags	3	3	1	2	1	2	3	0
Time Trend	Yes	Yes	Yes	Yes	No	No	No	No
Factors	No	Yes	No	No	No	No	No	No
Quarter \times Sub FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Drop Crisis	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	77,387	77,387	77,387	77,387	77,387	77,387	77,387	77,387
R^2	0.41	0.41	0.41	0.41	0.41	0.41	0.41	0.41

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table summarizes the estimates for the portfolio multipliers when we use different specifications to extract demand shocks. All estimates are for portfolios formed by detailed rating and quarter-to-maturity as in Section 4.1. “Factors” refer to the parametric estimation of flow factor structure explained in Appendix C.3.

Table 14: CUSIP Multipliers from Large Shocks

	Homo. OLS	OLS			First-stage	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)
Shock	0.35*** (0.05)	0.06 (0.04)	0.05 (0.04)	0.04 (0.04)		0.05 (0.05)
Substitute return						1.18*** (0.06)
Group Shock					2.54*** (0.28)	
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Group x Quarter FE	No	Yes	Yes	Yes	No	No
ST/LT x Quarter FE	No	No	No	Yes	No	No
Drop Crisis	No	No	Yes	Yes	Yes	Yes
N	166,749	166,747	157,247	157,247	157,248	157,248
R^2	0.20	0.40	0.37	0.39	0.61	0.17
First-stage F-statistic						79.08

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table summarizes the estimates for the CUSIP multipliers when we only consider the subset of bonds with shocks that are above-median shock in size. Column 1 is the OLS estimates from regressing bond returns on the demand shock, controlling for time fixed effects—this specification corresponds to a model in which we assume homogeneous cross-elasticity with all other bonds. Columns 2 and 3 directly control for close-substitute prices using detailed rating x time fixed effects. Column 4 additional controls for maturity (long-term/short-term) x quarter fixed effects to control for potentially omitted time-varying maturity factor in holding shares. Column 5 and 6 relate to the IV specification in which we regress bond returns on the demand shock and substitute returns, while controlling for time fixed effects (and additional controls depending on the specification). We instrument for substitute returns using demand shocks to substitute assets. The IV specification corresponds to the model that allows for heterogeneous cross-elasticities between close and distant substitutes. The parentheses contain standard errors clustered at the substitute group x time level.

Table 15: Portfolio Multipliers from Large Shocks

	Homo. OLS	OLS			First-stage	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)
Shock	0.80*** (0.11)	0.33*** (0.10)	0.28** (0.10)	0.23* (0.09)		0.27** (0.10)
Substitute return						0.97*** (0.07)
Group Shock					2.76*** (0.32)	
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Group x Quarter FE	No	Yes	Yes	Yes	No	No
ST/LT x Quarter FE	No	No	No	Yes	No	No
Drop Crisis	No	No	Yes	Yes	Yes	Yes
N	39,324	39,323	37,086	37,086	37,086	37,086
R^2	0.25	0.51	0.47	0.49	0.51	0.24
First-stage F-statistic						71.62

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table presents the quarter to maturity \times detailed rating level multiplier estimates using the subset of portfolios that experienced shocks that are above-median in absolute size. Column 1 is the OLS estimates from regressing bond returns on the demand shock, controlling for time fixed effects—this specification corresponds to a model in which we assume homogeneous cross-elasticity with all other bonds. Columns 2 and 3 directly control for close-substitute prices using detailed rating \times time fixed effects. Column 4 additional controls for maturity (long-term/short-term) \times quarter fixed effects to control for potentially omitted time-varying maturity factor in holding shares. Column 5 and 6 relate to the IV specification in which we regress bond returns on the demand shock and substitute returns, while controlling for time fixed effects (and additional controls depending on the specification). We instrument for substitute returns using demand shocks to substitute assets. The IV specification correspond to the model that allows for heterogeneous cross-elasticities between close and distant substitutes. The parentheses contain standard errors clustered at the substitute group \times time level.

D Alternative Specifications

In this section we explore other specification to estimate equation 7 i.e.,

$$\Delta p_{j,t} = M u_{j,t} + \tilde{M} \Delta p_{g(j),t} + \tilde{M}^m \Delta p_t^m + \tilde{v}_{j,t}$$

Our objective was to estimate for M , hence our baseline specification was a fixed effect regression, which directly controls for $\Delta p_{g(j),t}$ and Δp_t^m by including group \times quarter fixed effect. While this specification is most robust, we lose economic content by not being able to estimate \tilde{M} . Hence, we additionally explored an IV specification, where we include $\Delta p_{g(j),t}$ (instrumenting it with $u_{g(j),t}$) and controlling for Δp_t^m using time-fixed effects. In the interest of brevity we left a third specification where we also estimate \tilde{M}^m for the appendix. In this specification we can directly include $\Delta p_{g(j),t}$ and Δp_t^m , which we instrument with $u_{g(j),t}$ and $u_{m,t}$ respectively.

Table 16 presents the estimates for all three specification. The own multiplier estimates are unchanged across the three specifications. The (near) substitute passthrough is close to one, where as the distant substitute passthrough is close to zero. These estimates are in line with what we would economically expect—a securities price is more closely tilted to near substitutes than distant substitutes.

Table 17 presents similar estimates for the baseline level of aggregation. The own multiplier estimates are extremely similar across the three specifications. The (near) substitute passthrough is close to one, where as the distant substitute passthrough is significant but smaller than the near substitute. Overall, \tilde{M}^m being larger for this more aggregated portfolio, compared to the CUSIP \tilde{M}^m is also in line with what we would economically expect.

Table 16: CUSIP-level multipliers: alternative specifications

	OLS	2SLS	
	(1)	(2)	(3)
Shock	0.05 (0.05)	0.05 (0.05)	0.05 (0.05)
Substitute return		1.07*** (0.06)	1.07*** (0.07)
Mkt return (far sub.)			0.03 (0.10)
Quarter FE	Yes	Yes	No
Group x Quarter FE	Yes	No	No
Drop Crisis	Yes	Yes	Yes
N	314,534	314,534	314,534
R^2	0.36		
First-stage F-statistic		72.49	33.54

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table presents CUSIP level multiplier estimates from an alternative instrumental variable specification. As in the baseline column 1 corresponds to directly controlling for close-substitute prices using detailed rating x time fixed effects. Similarly, as before, Column 2 corresponds to the IV specification in which we regress bond returns on the demand shock and substitute returns, while controlling for market returns using time fixed effects. We instrument for substitute returns using demand shocks to substitute assets. Column 3, is the alternative specification, that directly includes market returns, which we instrument for using demand shocks to the bond market as a whole. The parentheses contain standard errors clustered at the substitute group x time level.

Table 17: Baseline aggregation level multipliers: alternative specifications

	OLS	2SLS	
	(1)	(2)	(3)
Shock	0.33*** (0.10)	0.35*** (0.10)	0.36*** (0.11)
Substitute return		0.90*** (0.05)	0.83*** (0.06)
Mkt return (far sub.)			0.45*** (0.11)
Quarter FE	Yes	Yes	No
Group x Quarter FE	Yes	No	No
Drop Crisis	Yes	Yes	Yes
N	77,387	76,348	76,348
R^2	0.41	0.21	0.36
First-stage F-statistic		69.85	24.14

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table presents the quarter to maturity \times detailed rating level multiplier estimates from an alternative instrumental variable specification. As in the baseline column 1 corresponds to directly controlling for close-substitute prices using detailed rating \times time fixed effects. Similarly, as before, Column 2 corresponds to the IV specification in which we regress bond returns on the demand shock and substitute returns, while controlling for market returns using time fixed effects. We instrument for substitute returns using demand shocks to substitute assets. Column 3, is the alternative specification, that directly includes market returns, which we instrument for using demand shocks to the bond market as a whole. The parentheses contain standard errors clustered at the substitute group \times time level.

E Substitutions in the Equity Market

In the main text of our paper, we focus on the corporate bond market. A priori, there are clear close substitutes in the bond market, which helps mitigate the effect of demand shocks. Ignoring close substitutes leads to biased multiplier estimates. Our estimates further validate this prior. Stocks, however, do not have clear close substitutes, and therefore it is harder to accommodate demand shocks without large price impacts. Assuming homogeneous substitution when estimating the multiplier is also not far off.

In this section, we confirm this prior on the equity market. Using the same Morningstar fund data, we follow identical procedures in shock construction and identification as in our baseline estimation. We estimate the stock multiplier with and without controlling for its close substitutes—defined in terms of loading on Fama-French three factors, or its industries. We find that controlling for close substitutes does reduce the estimated multiplier, but only very slightly, suggesting that substitutions are way less influential in the equity market than in the bond market.

Table 18: Multiplier estimates for stock markets

	Stock Return		
	(1)	(2)	(3)
Shock	0.380*** (0.086)	0.254*** (0.040)	0.252*** (0.072)
Group x Quarter FE	None	FF3	Industry
Quarter FE	Yes	Yes	Yes
N	144,768	136,270	135,201
R^2	0.188	0.332	0.263

This table reports estimates of the multipliers for the stock market at the security level, using the specification in (9). Column (1) reports the multiplier estimate *without* controlling for heterogeneous substitution patterns. Columns (2)-(3) control for group-time fixed effects. In Column (2), the close substitute group is defined as $3 \times 3 \times 3$ bins sorted by loadings on Fama-French three factors. In Column (3), the close substitute group is defined as industry groups following the Global Industry Classification Standard (GICS), with 27 groups in total. The parentheses contain standard errors clustered at the substitute group x time level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 18 reports the estimates of the equity multipliers at the stock level. The shocks are constructed in the same way as in the baseline specification for corporate bonds, and the sample covers the CRSP universe of the U.S. listed stocks from 2003Q1 to 2020Q4. In the first column, we regress stock returns on stock level shocks, controlling for the time-fixed effect only. This specification assumes homogeneous substitution patterns across stocks. The point estimate is 0.38, indicating a one-percent demand shock to a single stock leads to a 38 basis points increase in the stock price.

In Column (2), we add group-time fixed effects to control for the close substitutes. The group here is defined as the stocks with similar factor loading as the test stock. Specifically, we compute the loading of each stock on Fama-French three factors, and then triple-sort them into $3 \times 3 \times 3$ groups. The total group numbers (27) are comparable with the detailed rating groups in the bond market (21). Controlling for the group-time fixed effects does reduce the point estimate, suggesting that the equity market also exhibits heterogeneous substitution patterns. Nevertheless, the reduction in the multiplier is much less dramatic than the bond market counterpart as reported in 1. After controlling for close substitutes, the multiplier is still around 0.25, only 34% smaller than the univariate regression. To put it into perspective, in the bond market we see an almost 90% drop in the multiplier once controlled for the close substitutes. The multiplier is also much larger than the point estimate in corporate bonds, indicating much inelastic demand in the equity market. This is also consistent with the prior that it is harder to substitute away in the stock market than the bond market. In Column (3) we define the close substitutes in terms of the Global Industry Classification Standard (GICS) industry groups (27 groups in total). The estimate is very close.

In conclusion, we find that the force of heterogeneous substitution is at work in the equity market as well, but it is much weaker than the bond market, consistent with our prior.

F Interpretations Based on Logit Demand Systems

In this section, we show how our reduced form approach is linked to a nested logit demand system. Specifically, we derive the elasticity matrix Γ from a nested logit demand system and show that it yields the reduced-form specification as we used in the main text.

Following [Kojen and Yogo \(2019a\)](#), consider the demand structure for a representative investor as follows:

$$\begin{aligned}
w(j | g) &= \frac{\exp(\delta(j, g))}{\sum_{j \in g} \exp(\delta(j, g))} \\
w(g) &= \frac{(\sum_{k \in g} \exp(\delta(k, g)))^\lambda}{1 + \sum_{g'} (\sum_{k \in g'} \exp(\delta(k, g'))^\lambda)} \\
\delta(j, g) &= \beta_g p_j + \beta X + u_j,
\end{aligned} \tag{24}$$

where $w(j | g)$ is the conditional share of asset j in group g , $w(g)$ is the share of group g in the whole portfolio. The unconditional share of asset j is therefore $w(j) = w(j | g)w(g)$. The parameter λ controls the substitution at the group level. When $\lambda = 1$, this system is reduced to the simple logit system as

$$w(j) = w(j | g)w(g) = \frac{\exp(\delta(j, g))}{1 + \sum_{g'} (\sum_{k \in g'} \exp(\delta(k, g'))^\lambda)}.$$

Taking the derivative of $w(j)$ with respect to p_k , we have:

$$\frac{\partial w(j)}{\partial p_k} = \begin{cases} \beta_g w(j) (1 - \lambda w(j) - (1 - \lambda)w(j | g)) & j = k \\ -\beta_g w(j) (\lambda w(k) + (1 - \lambda)w(k | g)) & j \neq k, j \& k \in g \\ -\beta_{g'} \lambda w(j) w(k) & j \in g, k \in g' \neq g \end{cases} \tag{25}$$

Log linearize the system (24) with the partial derivatives above, reorganize, we have:

$$\log \Delta w(j) = \beta_g \Delta p_j - \beta_g (1 - \lambda) \sum_{k \in g} w(k | g) \Delta p_k - \lambda \sum_{g'} \sum_{k \in g'} \beta_{g'} w(k) \Delta p_k + u_j$$

Applying market clearing condition $\log \Delta w(j) = \Delta p(j)$, we can derive the equation for estimation:

$$\Delta p_j = \underbrace{\frac{1}{(1 - \beta_g)} u_j}_M + \underbrace{\frac{\beta_g}{(\beta_g - 1)} (1 - \lambda) \sum_{k \in g} w(k | g) \Delta p_k}_{\tilde{M}} + \underbrace{\lambda \frac{1}{(\beta_g - 1)} \sum_{g'} \sum_{k \in g'} \beta_{g'} w(k) \Delta p_k}_{\text{Time FE}} + \tilde{v}_j \tag{26}$$

This equation maps directly to our empirical specification. When $\lambda = 1$, the nested system is reduced to standard logit demand, and $\tilde{M} = 0$. In this case, we do not need

to control for price changes for close substitutes. Our estimates of \tilde{M} strongly reject this assumption. Therefore, when specifying demand for corporate bonds, a nested system is preferred over the standard logit demand.

G Arbitrage Risk

Consider a two-period economy populated with measure 1 of homogeneous investors. Investors have CARA utility, with absolute risk aversion coefficient γ . There are two risky assets with normally distributed payoff in period 2

$$\begin{pmatrix} D_A \\ D_B \end{pmatrix} \sim N(\mu, \Sigma) \quad (27)$$

$$\Sigma = \begin{pmatrix} \sigma_A^2 & \rho\sigma_A\sigma_B \\ \rho\sigma_A\sigma_B & \sigma_B^2 \end{pmatrix} \quad (28)$$

In addition, there is a risk-free asset with total return normalised to 1. Investors choose portfolio x to maximize expected utility,

$$x = \frac{\Sigma^{-1}}{\gamma}(\mu - p) \quad (29)$$

$$x_A = \frac{\gamma\sigma_B^2(\mu_A - p_A) - \rho\gamma\sigma_A\sigma_B(\mu_B - p_B)}{|\gamma\Sigma|} \quad (30)$$

where $|\cdot|$ denotes matrix determinant. Apply market clearing, we get

$$x + u + \nu = \bar{x} \quad (31)$$

where u is the observed exogenous demand and ν is the unobserved demand shocks. Focusing on asset A , we can express its price as

$$p_A = \underbrace{\gamma\sigma_A^2(1 - \rho^2)}_{=M} u_A + \frac{\rho\sigma_A}{\sigma_B} p_B + \text{constant} + \underbrace{\gamma(\sigma_A^2\nu_A + \rho\sigma_A\sigma_B\nu_B)}_{\text{error term}} \quad (32)$$

The micro multiplier maps to the term $\gamma\sigma_A^2(1 - \rho^2)$ in this case, which is determined by both the risk aversion coefficient as well as asset A's residualized risk. To understand the term

$\sigma_A^2(1 - \rho^2)$, regress asset A's price on its substitute, asset B's price,

$$p_A = \beta_0 + \beta_1 p_B + \epsilon \tag{33}$$

the variance of ϵ captures the risk in asset A that cannot be hedged by holding asset B, which is equal to $\sigma_A^2(1 - \rho^2)$.

Hence we have shown that the multiplier M is increasing in both the risk aversion coefficient and the residualized risk that cannot be hedged by the substitute portfolio.