The Effect of Ownership Concentration on Government Bond Yields

Timothy R. McDade
tim.mcdade@duke.edu

Department of Political Science
Duke University

April 9, 2022

Abstract

Bonds are one of the most important sources of funding for modern governments, but we understand little about the implications of who owns government bonds. Other, more aggregated cross-country analyses have found a link between size of bond debt and fiscal balances, but I take a granular approach. I introduce a novel data set of the holders of individual securities to examine the pricing effects of variations in ownership structure. How does the concentration of a government bond's ownership structure affect its yield? I propose that a government bond whose ownership is more concentrated will have more volatile prices because it is more susceptible to movements by large investors. This volatility incurs a premium and corresponds to higher secondary market yields. I expect, and find empirical support, that the issuer's general riskiness does not create endogeneity by affecting both yield and ownership concentration. Although a risk-based match between investor and security might affect investors selecting into the market, position size is driven by other factors. Sourcing enterprise-quality data on individual securities from multiple financial data providers, I find support for my hypotheses using a variety of time-series econometric methods. This paper contributes to the literature security-level clarity on the relationship between ownership structure and yield and a clear depiction of the trade-off for reliance on a small group of investors. Because governments are responsible for debt service payments that are affected by secondary market yields, my findings have implications for fiscal policy and the tension between governmental accountability and economic credibility.

Contents

List of Figures

1 Introduction

Countries fund themselves using money from a variety of sources: taxation, loans, state-owned assets, and debt markets. Since the Brady debt restructurings in the early- to mid-1990s, countries have increasingly relied on issuing bonds. I seek to understand a dynamic of this market that can be opaque and unnoticed: the number of entities that own the government's debt. This paper will answer the question, "how does a concentrated ownership structure of government debt securities affect a government fiscally?"

The connection between governments and investors is increasingly important because of today's international financial markets. The government's ability to manage the possibly conflicting interests of bondholders and constituents is relatively understudied and stands apart from the literature on financial markets, collective action, bureaucracy, and regulation. Government capture by bond markets could lead to dependency, and the distribution of property among bondholders is something we do not yet understand well. That creditors had enough bargaining power over the Argentine government to (temporarily) impound an Argentinian naval frigate shows that this topic matters to governments, investors, and financial intermediaries.

Forecasting financial market movements is important because it helps explain how investors capitalize political phenomena through specific changes in value of securities. Although equity markets can provide insight into tariff policy and others that affects publicly listed companies, government bond markets reflect the market's assessment of a government's ability and willingness to repay its debt. While aggregated cross-country analyses have found that countries with more bond debt have improved fiscal balances, I focus instead on the effects of changes in the ownership of specific securities. Because of this granular focus, my approach is useful for forecasting financial market movements.

The concentration of a security's ownership structure changes based on how many investors own what share of the security. For example, one bond whose outstanding debt is held exclusively by two investors is more concentrated that that of a bond that is held broadly by a large investor base. Notwithstanding the concentration of a security when issued, relative movements in the concentration could produce effects.

I propose that a concentrated ownership structure (fewer investors own more assets) of a government bond corresponds to higher secondary market yields. These higher yields then push up the primary market yields of debt issued to replace maturing securities, increasing future debt service payments. This paper contributes to the literature granular, security-level clarity of the relationship between ownership concentration and yield. More generally, it contributes a clear theoretical understanding of one way the ownership structure of a government bond can affect debt service costs. These findings have implications for redistribution and the tension between democratic accountability and economic credibility.

It is possible that attributes of the government borrower contribute to determining ownership concentration in the first place by affecting the investors who enter the market and the size of the positions they accrue. For example, governments with a large amount of outstanding debt who have robust repayment histories could be attractive to many types of investors, while governments who only come to market periodically or have a history of default or restructuring could only be within the risk tolerance of a small subset of specialized investors.

I address these concerns theoretically and empirically. Theoretically, I argue that while risk factors may affect an investor's decision to enter the market, the size of position (conditional on market entry) is likely determined by other factors. If an investor's position size is unrelated to their risk appetite, then the ownership concentration of the security must also be unrelated to risk, because concentration is dependent on the position size of all investors. I also implement empirical design and methodological controls to eliminate the effect of other issuer-related concerns that could affect both yield and ownership concentration.

The next section discusses existing literature that helps form the analytical basis for my theoretical framework, which is detailed in the third section. The fourth section discusses research design, the fifth discusses results, and the final section concludes.

2 Background

Why the bond market? Modern countries rely heavily on bond markets for funding from creditors that are institutional investors, individuals, central banks, sovereign wealth funds, international financial institutions, and others. Naturally, these investors have different budget constraints and investment objectives. Creditors exert bargaining power over governments, sometimes with substantial policy effects: the extreme version of this phenomenon is creditor bargaining power over a government after a default. The hold-out investors referred to above restricted Argentina's ability to pay other creditors who had agreed to a debt restructuring before repaying them the original, un-restructured debt. One of the creditors' conditions to resolving the Argentinian standoff was input in future Argentinian domestic market fund-raising (Stevenson, 2016).

2.1 Aggregated PE Knowledge

To better understand how such ownership dynamics affect governments, the political economy literature offers several lessons. This literature has a strong tradition of examining the relationship between capital and governments (Przeworski and Wallerstein, 1988), and more recently the connection between financial markets and governments. But most of the recent literature focuses on only one of the two possible causal directions: the effect of various political phenomena on financial assets (Ferrara and Sattler, 2018). Political phenomena affect, among other things, the price level (Roberts, 1990; Campello, 2015) and volatility (Bechtel, 2009) of financial markets, as well as the currency composition (Ballard-Rosa et al., 2021) and maturity structure (McDade et al., 2021) of debt issuances. However, as Ferrara and Sattler (2018: p. 21) note, the connection between politics and financial markets is bi-directional: financial markets also affect the government. This should be particularly true for government bond markets.

Nevertheless, the political economy literature offers insights about why countries make choices about certain characteristics of the debt they issue. Countries can strategically choose to issue debt denominated in local currency or foreign currency in order to minimize currency risk or achieve domestic political goals (Eichengreen and Hausmann, 1999; Ballard-Rosa et al., 2021). Along similar lines, they can choose to issue short-term or long-term debt, choose to issue a large amount of debt at once or issue smaller amounts more frequently, and choose to default or not to default (Roos, 2019). But the effects of who owns debt securities remain murky.

Two main works provide specific insight into the effects of concentrated ownership of government debt. The first argues that countries' unwillingness to default on sovereign debt derives in part from the increasing concentration of the global financial system (Roos, 2019). Because states can only really finance themselves via state-owned enterprises, taxation, or borrowing (O'Connor, 1979), such concentration imposes market discipline on debtor states by eliminating alternative financing options for countries in distress (Roos, 2019: p. 71). While Roos' analysis is insightful, it does not draw data on the ownership of particular securities, leaving room for interrogation of the mechanisms.

The second work digs deeper into the policy effects of bond market dependency. Kaplan (2013) offers a collective action explanation for how bond market indebtedness constrains fiscal policy. When faced with a fiscal situation that does not prioritize debt repayment, the small cost of market exit incentivizes bondholders to do so. Such market exit then "yield[s] a higher-risk premium quickly that translates into rising funding costs for sovereign borrowers" (Kaplan, 2013: p. 10). Countries with high bond market exposure, in this line of reasoning, are more susceptible to creditor influence, and tend to have more orthodox fiscal policy as a result.

Moreover, in a follow-up paper, Kaplan and Thomsson (2017) show that countries whose external debt is heavier on bonds exhibit greater fiscal balance. The authors conclude that because the "bond market" prefers governments to retain orthodox fiscal policy to better pay off debt, countries with more bond debt conform their fiscal policy to the position bondholders most prefer. But this work suffers from several flaws that muddy its conclusions. First, it depends on aggregated data that does not permit examination of the proposed mechanism, price pressure. Second, its conclusions depend on the assumption that bond market actors homogeneously prefer a certain kind of fiscal policy. If this assumption does not hold true, then different actors in the market would do different things in reaction to government fiscal policy, not necessarily resulting in more expensive financing. Moreover, the authors do not probe the ownership structure of the bond debt itself; they merely consider its size in relation to the issuing country's total external debt.

2.2 Heterogeneous Preferences

For the bond market to function, investors must buy and sell bonds. Aside from risk-adjusted expected rate of return, there are two reasons to do so, each of which informs the investor's perception of default risk. The first, policy preference, is an ideal point on the policy spectrum of the government's ability to repay. The second, risk preference, is tolerance over deviation from that ideal point. For movement in the market to occur, there must exist some heterogeneity among bond market investors across policy or risk preferences such that different investors buy and sell debt under the same conditions.

Nevertheless, traditional capital market models such as CAPM and Black-Scholes assume homogeneous investor preferences; some authors argue that homogeneous preferences in these models does not accurately reflect the dynamics of equity markets and instead results in predictable and repeatable market cycles (Levy and Levy, 1996; Chan and Kogan, 2002; Abbot, 2017).

If investor preferences were homogeneous, markets should exhibit certain tendencies. Mosley (2000: p. 746) theorizes that when preferences are homogeneous, the policy consequences of investor behavior in the issuing country will be greater. Mosley finds that institutional investors use the same indicators to inform their decisions, namely inflation and fiscal balance. Therefore, if preferences were homogeneous and investors use the same information to inform decisions, markets should clearly react to microeconomic policy announcements. But Mosley et al. (2020) show that prices of sovereign debt in bond markets do not systematically react to significant changes in microeconomic policy, implying that there is no entity called "the market" that reacts systematically, as a whole, to microeconomic policy changes. In fact, Brooks et al. (2019) show that higher investor uncertainty about government willingness and ability to repay does not lead to the market agreeing upon a higher risk premium for that government's debt. Instead, different actors make different decisions, leading to higher volatility of bond spreads.

The economics and finance literature finds clear support for heterogeneous preferences. Even if institutional investors generally inform their actions with the same indicators (Mosley, 2000), investor preferences vary across three general categories. The first is belief about repayment, which can manifest in preferences over policy of the issuing government (Hardie, 2006; Brock and Durlauf, 2010; Mosley et al., 2020) or beliefs about the underlying economic growth rate (Cvitanić et al., 2012; Chabakauri, 2015). The second is risk preferences (Levy and Levy, 1996; Fischer et al., 1996; Campbell and Viceira, 2001; Isaenko, 2008; Condie, 2008; Weinbaum, 2009; Sarasvathy et al., 2010; Christensen et al., 2012; Cvitanić et al., 2012; Chabakauri, 2015; Hauser and Kedar-Levy, 2018), which results in some customers exiting markets before others (Hirschman, 1970: pp 33-43). The third is investment goals derived from investor position, like time horizon (Modigliani and Sutch, 1966; Wachter, 2003; Sangvinatsos and Wachter, 2005; Chan and Kogan, 2001; Isaenko, 2008; Cvitanić et al., 2012; Wellhausen, 2015) or liquidity (Hauser and Kedar-Levy, 2018; Chen et al., 2020). All three kinds of heterogeneity contribute to making markets work.

2.3 Effects of Ownership Structures

Heterogeneity means different bonds have different creditors who enter and exit the market at different times for different reasons. Therefore each asset has a particular ownership structure, which then has an effect on its price. For example, investor movements into and out of managed investment funds can distort prices away from the fundamental values of the assets in which the fund invests (Vayanos and Woolley, 2013). Much of the scholarship on the pricing effect of the ownership structure of bonds analyzes what is called the preferred habitat hypothesis: that investors who prefer assets of a certain time horizon will propel movements in the prices of those assets (Modigliani and Sutch, 1966). Recent empirical work has found support for the preferred habitat hypothesis (Wachter, 2003; Greenwood and Vayanos, 2010), especially in relation to pension and insurance company demand for assets at the long end of the yield curve (Greenwood and Vissing-Jorgensen, 2018). The preferred habitat hypothesis is one example of how the heterogeneity that causes investors to enter and exit certain securities causes prices to move and results in concentration or dispersion.

2.4 Takeaways

Because the two main bodies of literature examining the effects of ownership structures remain largely unconnected, this phenomenon deserves another look. Political economy literature often relies on assumptions about market preference distribution and untested mechanisms driving conclusions. The finance literature analyzes these mechanisms, but stops short of security-level analysis of the pricing effects of ownership structures of government bond markets. I will attempt to bridge this gap by filling in some of the gaps in the political economy literature using tools from the finance toolbox.

3 Argument

Despite all the useful context, the literature leaves unanswered the relationship between ownership concentration and yield. The political economy literature in particular comes closest, but it does not offer empirical evidence at a granular enough level of analysis to validate its mechanisms and depends on assumptions thoroughly refuted by the finance literature. Moreover, although political economists have offered general lessons about trends in government-finance relations, there is no clear answer to how the concentration of a security's ownership affects governments. I argue that government bonds with higher ownership concentration have higher price volatility. This higher volatility results in a volatility risk premium, which translates to higher future yields. I also address the potential that the "riskiness" of a security affects both its ownership concentration and its yield by arguing that the factors that drive an investor to enter a market are different from the factors that affect the size of the position they accrue. In doing so, I contribute security-level clarity of the relationship between ownership concentration and yield. More generally, I contribute a clear theoretical understanding of one way the ownership structure of a government bond can affect yields.

3.1 Ownership and Yield

One important characteristic of a security's ownership structure is how concentrated its ownership is among its investors. This ownership concentration can change independent of the security's price. The two arise from different properties of the buy/sell transaction: changes in price come from aggregated buyer willingness (i.e. demand) to buy the same quantity on offer at a different price than the seller is offering, and changes in concentration come from a different *number* of investors willing to buy the same amount of the security for the same price as the seller is selling. Concentration must increase, decrease, or stay the same with every transaction depending on whether the buyers number less than, more than, or the same as the number of sellers.[1]

Like other measures of the ownership structure of securities, ownership concentration has been shown to affect prices. Greenwood and Thesmar (2011) show that ownership concentration makes asset price more susceptible to swings in non-fundamental flows such as movements into and out of a managed fund: securities with higher ownership concentration have higher price volatility. The intuition behind these findings is that, holding all else equal, when an asset is held primarily by several large players, any unexpected movement into or out of the security is "unlikely to be 'cancelled' by the trades of the other owners, resulting in price impact" (Greenwood and Thesmar, 2011: p. 472). But because liquidity shocks are inherently difficult to predict, the authors focus on predicting volatility.

The financial economics literature clearly establishes that higher volatility in bond prices is associated with a negative risk premium, implying that more volatile bonds have lower prices and higher yields (Almeida and Vicente, 2009; Chung et al., 2019). I extend these findings by arguing that investors demand a premium to compensate them for the volatility risk, including in situations where volatility risk is derived from a concentrated ownership structure.

Why does concentration lead to volatility? Greenwood and Thesmar (2011) show that in equity markets, more concentrated ownership corresponds to higher price volatility. The underlying reasoning is that in situations when each holder owns more of the asset, decisions to divest can result in larger quantities on the market at the same time, resulting in larger price swings. I expect similar dynamics to exist in bond markets. Although bond market prices incorporate information about the debtor's willingness and ability to repay the debt, many of the factors affecting debtor government willingness and ability to repay remain stable over time.[2] Prices change much more frequently than fundamentals do. These minute-to-minute movements in bond prices are not derived from changes in fundamentals; instead, they respond to external factors such as stock prices, investor liquidity, and changes in investor expectations about fundamentals. I expect that in a concentrated environment,

[1] Consider an example. Seller X brings 10 shares of stock ABC to market when the market price is $10/share. If 10 buyers are willing to buy one share each for $10/share, price would remain constant but concentration would decrease. If one buyer is willing to buy five shares at $10/share and another buyer is willing to buy five shares at $9/share, the price and concentration both decrease. If one buyer is willing to buy all 10 shares at $9/share, price decreases and concentration remains steady. If one buyer buys all 10 shares from Seller X as well as 10 shares from Seller Y, but is only willing to pay $9/share, concentration increases and price decreases. If the same buyer buys these 20 shares from Sellers X and Y but is willing to pay $12 for each share, concentration and price both increase.

[2] Such as domestic institutions, government composition, credit history, etc.

these movements into and out of government bonds due to non-fundamental factors will result in higher volatility because sales by large holders will be unmatched on the demand side.

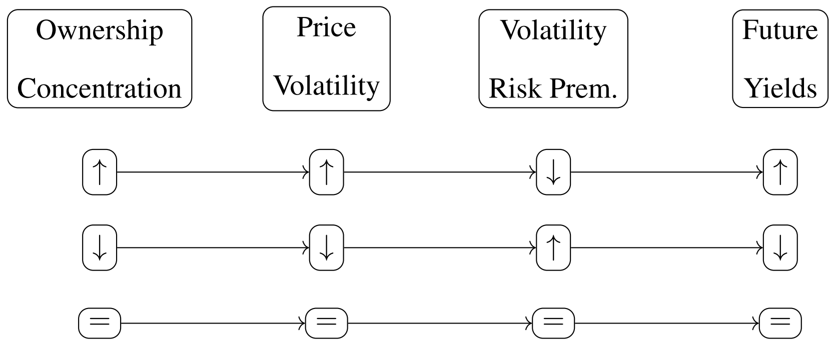

Drawing upon this reasoning, I expect securities with high ownership concentration to have higher price volatility that leads to higher secondary market yields. But why does this matter to governments? In the case of stocks, investor exit that causes a share price to drop doesn't have an immediate effect on the finances of the issuing company. But in the case of bond markets, it is possible to think of each investor having an exit *threat*: market exit via sale can negatively affect the issuing government. The political economy literature shows one high-level example of this, which relies on a more micro-level mechanism.

The high-level explanation is that bondholder exit threats are a failure of collective action that can result in "indirect influence over debtor governments" (Kaplan and Thomsson, 2017: p. 607). Analytically, exit from a bond market is more complicated than exit from other markets because even in a collective action failure the investor still retains most of his/her investment. Because each creditor has such a small share of the borrower's debt exposure, the creditor is incentivized to exit the market instead of holding their assets or providing new funds. Here, the collective good for which the creditors fail to bargain is a policy change that would increase government willingness and/or ability to repay the debt. In the bond market, decentralized creditors "benefit from their coordination problem," thereby "indirectly increas[ing] their influence over debtor governments": "if countries do not demonstrate commitment to policies that ensure debt repayment, bondholders can cut their financial ties without incurring a severe profitability shock" (Kaplan and Thomsson, 2017: p. 607).

But when governments issue bonds in the primary market, the coupon and yield to be paid upon maturity are agreed upon at the start; the secondary market is merely an appraisal of the issuing government's likely willingness and ability to carry out its promise to repay. So why would secondary market exit a-la-Kaplan and Thomsson affect the issuing government? The answer is that Kaplan and Thomsson's high-level explanation rests on a micro-level explanation: secondary market price changes affect the issuance terms of new debt, which in turn affect debt service payments.

When governments issue debt, they sell a tranche of debt securities to an underwriter at a previously agreed-upon price. The underwriter then resells the securities to actors over a secondary market, who buy the debt at market prices. In this way, the secondary market price can differ from the original price at which the government sells the debt security to the underwriter, and therefore from the amount due to the holder upon maturity. As market actors buy and sell a government debt security on the secondary market, its price and yield (face value of the debt minus the price) fluctuate accordingly. Higher demand corresponds to lower yields, and lower demand to higher yields. As debt securities reach their maturity date, the government issues new debt to take the place of the maturing debt. The secondary market yield of the maturing debt then informs the yield at which the government issues the new debt (Duffie, 2010; Lou et al., 2013; Eisl et al., 2019; Cole et al., 2020; Sigaux, 2020).

Earlier, I said I expect that securities with more concentrated ownership structures will have higher price volatility and therefore higher secondary market yields. Because secondary market prices inform the terms of new bond issuances, any pricing effects of ownership structure should affect the yields of subsequent issuances. This is important because governments with higher ownership concentration across their debt securities could, in the long run, be required to pay higher yields on debt to attract investors, leading to higher debt service levels.

3.2 Investor Selection

However, it is possible that the "riskiness" of a security could influence both its ownership concentration and its yield. If true, there is a potential that a security's riskiness could confound the relationship between concentration and yield. The relationship between yield and riskiness is fairly intuitive: investors will demand a higher premium to invest in the debt of a country if the risk of default is higher. But the relationship between riskiness and concentration merits consideration in more detail. It is possible that only certain kinds of investors are willing to buy into a security of a certain riskiness, and that riskiness also affects the investor's position. By affecting the selection of investors and the size of their positions, riskiness could affect concentration.

To resolve this endogeneity, I propose an investor-level complement to this theory, which has so far focused on the security level. An investor taking a position in a security is a two-step process: first, an investor decides to enter the market or not, and, contingent upon entering the market, the investor decides what size position to take. I define the "riskiness" of a debt security to be the market's valuation of the likelihood the security's issuer will be unwilling or unable to repay.

If a security's riskiness affects ownership concentration, it should do so by affecting both investor decisions: market entrance (the selection stage) and position size (the outcome stage). But the riskiness of the security itself is not the operative concern, nor the riskiness of the issuer, since those are constant across all investors at any given point in time. Rather, the characteristics of the investor push different investors to make different decisions using the same information.

I expect that there is enough variation in investor tactics that investors will vary significantly in their interpretation of the same information. Consider, for example, credit default swap (CDS) spreads as a measure of issuer "riskiness."[3] CDS spreads are commonly available information; all investors can use them as a benchmark of the market's expectation that an issuer might default. If spreads make a move and the market reacts, by definition some investors took that movement as a signal to sell and others took it as a signal to buy. There are two questions: whether these groups of investors who co-move all have the same risk tolerance and whether risk tolerance affected the position decisions of those who ended up entering the market.

I expect that the two processes are different – investors use different criteria to decide when to get into a market than to decide about the size of their position. While traits such as risk acceptance and the characteristics of the fund (price-to-earnings ratio, etc.) likely affect whether or not to invest in the security at all, other factors likely affect the size of the stake once they have decided to enter. For example, leverage and non-fundamental flows could affect the amount of capital available to invest (Greenwood and Thesmar, 2011); movements in other markets likely affect how hedged investors

[3]In such a security, the purchaser buys the CDS as insurance against the issuer defaulting. If the issuer defaults, the CDS provider (lender) reimburses the CDS purchaser. CDSs are usually paid for incrementally in a manner similar to an insurance premium, and their "spread" is the annual premium in relation to the notional amount insured, expressed in basis points.

wish to be; whether the investor uses fundamentals-based or quantitative investing strategies likely affects their investment tactics (Satchell and Scowcroft, 2000); whether an investor is a pension fund or an insurance company or a hedge fund can affect their preferred environment (Modigliani and Sutch, 1966); behavioral and demographic attributes can also affect portfolio construction (Frijns et al., 2008).

3.3 Hypothesis

This theoretical set-up leads to two hypotheses.

1. Ceteris paribus, a more concentrated ownership structure for a government bonds will lead to higher secondary market yields.
2. Although a security's riskiness may affect investor decisions about market entrance, I do not expect it to have a significant relationship with position size.

4 Research Design

To isolate the effect of ownership concentration on security price volatility and yield level, I undertake a three-part empirical approach. First, I examine the way investors select into the market to assess my expectation that the processes driving selection and position size are indeed different. Second, I examine the time-series relationship between ownership concentration and return volatility of a single representative security. Third, I extend these findings with a time-series cross-sectional approach for the population of bonds for which complete data is available.

4.1 Empirical Setting

I set my empirical study in California municipal bonds from 2013 to the present. This empirical setting keeps constant many variables that affect market perception of government ability and willingness to pay: the issuing entity itself remains the same, maintains largely continuous fiscal policy, and issues all its in US Dollars.

Moreover, Californian municipal bonds have several qualities that make them suitable for this study. They are numerous, cover a long period of time, have varying maturities, are widely invested in, and have been a continuous financial tool over the last decade. They are common investments for institutional investors, pension funds, and mutual funds. But even though municipal bonds are well-traded by institutional investors, they remain outside the mainstream of financial assets. Their slightly niche nature means that however well-capitalized their investors are, municipal bonds are not as good as cash; in some cases, they go days without a trade. Therefore, many municipal bonds are subject to influence by individual market actors.

4.2 Data Description

Assessment of this hypothesis has stringent data requirements: to my knowledge, this paper is the first time that a comprehensive data set of bond ownership has been used in political economy literature. First, holdings data on government bonds is quite difficult to come by. Even when procured, it is limited by the reporting requirements of the relevant jurisdictions. I source ownership data from the FactSet Standard Ownership Data Feed V5. This data describes each holder of a bond: who they are, how much of the security they hold, and more (FactSet, 2022). FactSet sources this data from regulatory filings as well as text-based data from investor websites and portfolio descriptions. I derive the ownership concentration of the security from this data.

Secondly, because my theory relies on time-series pricing of securities, I must obtain a historical security-level pricing data. Even in high-fidelity commercial data repositories like Bloomberg, such data is spotty at best. I source historical municipal bonds pricing data from the Municipal Securities Transaction Database from the Municipal Securities Rulemaking Board (MSRB) (Municipal Securities Research Bureau, 2022). I source descriptive data on each security (e.g. coupon rate, maturity date, amount outstanding) from Bloomberg (Bloomberg, 2022).

I also incorporate data to account for other possible explanations. For example, the amount of debt that a security has outstanding could affect volatility by changing the market size, offering more or less liquidity. I include security-level characteristics such as a security's yield at issue, its coupon structure, its maturity length, and its monthly close price. It is possible that securities that vary across these attributes could exhibit different volatility patterns, so I include them as explanatory variables in my time-series cross-sectional models. Furthermore, the months remaining until a security's maturity is likely related to the amount and kind of transactions in the secondary market, and therefore to security volatility. I also include it as an explanatory variable.

4.3 Empirically Accounting for Endogeneity

It is possible that a country's underlying "riskiness" affects both the ownership concentration of its debt securities and the yield of those securities, raising concerns about endogeneity. The independent variable, ownership concentration, is a function of the decisions of individual investors to enter the market or not, which is likely related to whether or not investor attributes (e.g. risk acceptance) align with the underlying riskiness of the security. The dependent variable, yield of the bond, reflects many things, chief among them the premium required to compensate investors for the possibility that the debtor is unwilling or unable to repay – that is to say, riskiness.

I take several steps to address the endogeneity concern. First and foremost, I select a research design that reduces the effect of riskiness as a confounding variable. I restrict the empirical setting to one issuer, which means the issuer's budgets, constituents, and services remain more stable than if I considered multiple issuers. This reduces variation in aspects of the government that could contribute to riskiness – if these attributes do not change over time, they cannot affect the regression results. Moreover, I examine the entire population of that issuer's securities during the time period in question in order to avoid sample selection bias.

Ideally, I would attempt to eliminate the effect of riskiness on my empirical models by including a security-level measure of California's riskiness. But the most commonly accepted measure, credit default swap (CDS) spreads, are not security-specific; they are a general indication of the creditworthiness of the issuer at a given point in time. Figure 1 shows that although CDS spreads are priced differently for debt with different maturity lengths, these spreads follow the same general pattern over time. Unlike price, volatility, and ownership concentration, CDS spreads do not vary by security over time and therefore cannot explain the security-level relationship between ownership concentration and volatility; I exclude them as a measure of riskiness of the issuer.

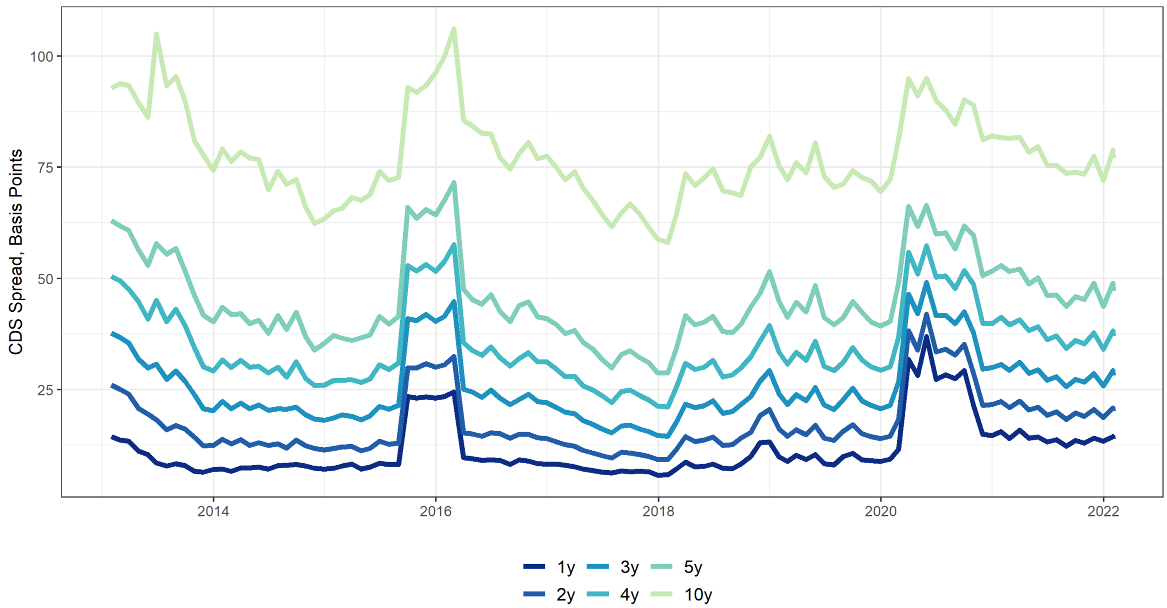

Figure 1: California CDS Spread Prices.

Moreover, because CDS spreads are the same for the whole market at any given time, they cannot actually cause any investor decisions. Rather, decisions are caused by the way that different investors process the same information. Another way I address endogeneity is to analyze the way that investor attributes, including risk acceptance, drive investor decisions about market entry and position size. Such an analysis helps contextualize the relationship between investor risk acceptance and concentration and therefore contributes to my main theoretical concern: the relationship between concentration and yield. An investor's choice about whether to enter the market for a particular security, the "selection effect," is a separate decision from how big a position to take once the investor enters the market. These two decisions could be driven by the same factors or two separate sets of factors.

I compare these two sets of drivers. Any discrepancy between the drivers of selection and position size would help provide evidence that although investors with certain attributes (e.g. risk accepting) may select into the market, those investors do not automatically develop large positions. Such a result would suggest that the processes driving selections and positions are different. A lack of relationship between investor risk acceptance and position size suggests a limit on the relationship between risk and concentration. More specifically, if an investor's risk acceptance is related to selection but not to position, then risk likely does not affect concentration and therefore is likely not a confounding variable.

To carry out such an analysis, I merge the security-level data set used in the first part of the analysis with a data set describing the holders of debt (FactSet, 2022). Specifically, the data includes attributes of specific funds such as the price-to-earnings (PE) ratio, price-to-books (PB) ratio, dividend-yield ratio, market beta, and other descriptors of the fund's portfolio. This data set covers some 150,000 funds, only several dozen of which have holdings large enough to report during any given period. One disadvantage of this data set is that it is a static snapshot of the most recent values for a given fund. Ideally, I would have all these attributes in time-series, but this may not be a large disadvantage because funds oftentimes pick portfolio attributes in advance of launching the fund and only rebalance periodically at the margins. As a result, I do not expect that the static nature of this data set will contaminate my results. Nonetheless, I do exclude some metrics that are time-dependent, such as those that illuminate the price movements of the fund's portfolio. This does not hamper testing my theory, however, because I am theoretically interested in more static attributes like general risk profile, as reflected in indicators PE ratio and beta.

4.4 Variable Definitions

The data set starts with descriptive data on all Californian municipal bonds issued after 2002 from Bloomberg, containing information such as issuance dates, maturity dates, coupon rates, ratings, yield at issue, and more (for 4,709 securities). MSRB has pricing data for 3,999 of these securities, but FactSet only has ownership data for 1,713 of them.[4] There are 1,661 securities that have both pricing and holdings data, only 399 of which had non-null holdings data. Of these 399, only 136 have data on the amount of debt outstanding at any given time, which is necessary to calculate percentage ownership. Only 114 have $h_{im} \in [0,1]$.

This description of the data set immediately an issue: the overall lack of holdings data. Holdings data is much harder to come across than security details or pricing information. Moreover, holdings data is only available from 2013 to the present. Although holdings data is available individually for some bonds before 2013, this early holdings data is inconsistent across time, type of bond, and

[4] Ownership data is only available after 2013, and some securities had matured by then.

often within one bond. Some of this inconsistency can be explained by issuers calling a bond before its maturity date, resulting in some bonds having reported holdings data for only a subset of their original maturity.

I use the Hirschman-Hirfindahl Index (HHI) to measure my main explanatory variable, ownership concentration.[5] For security i in month t with total amount outstanding o_{it}, I calculate the ownership concentration h_{it} across all owners $j = 1, 2, \ldots n$, where holder j holds a amount of the security, to be

$$h_{it} = \sum_{j=1}^{n} \left(\frac{a_{jit}}{o_{it}} \right)^2 . \tag{1}$$

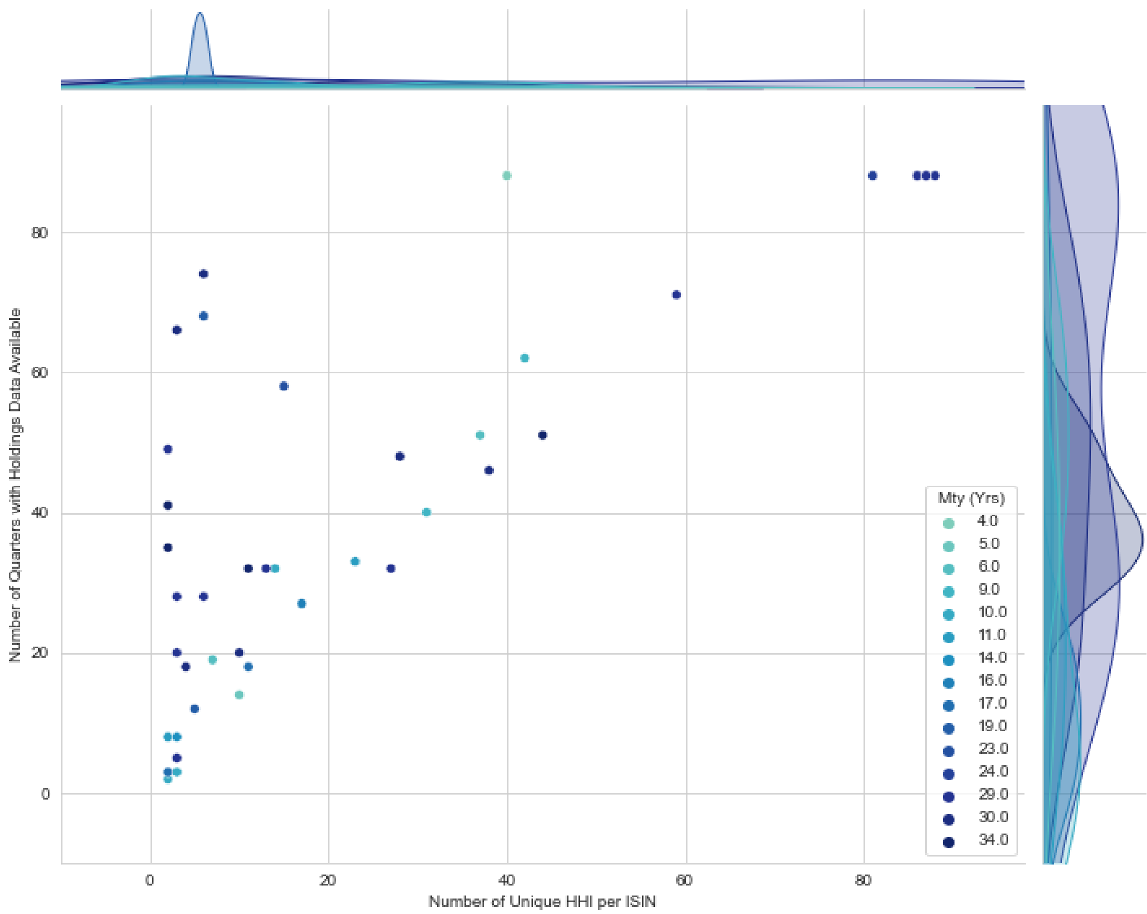

Figure 2: Data availability for securities with holdings data.

Figure 2 shows that the ownership concentration of many bonds does not change much over time.

[5]In line with (Cetorelli et al., 2007; Peltonen et al., 2014; Boermans, 2015), among others.

The x-axis shows the number of unique values of HHI that a security has had; higher values indicate more changes in ownership concentration. The y-axis shows the number of months for which a security has available holdings data; higher values indicate more data. The vertical cluster going up the left-hand side shows that it is fairly common for securities to exhibit stable ownership over time.

Several factors likely contribute to this phenomenon. Bond investors can hold a position over a long period of time because they seek conservative long-term returns or because they are passive investors (Sangvinatsos and Wachter, 2005; Sushko and Turner, 2018). Moreover, only some holders are required to report, so those that do report can often be institutions who hold stable positions. Lastly, bond investors often do not pay careful attention to the contract terms of their bonds and their positions may remain constant as a result (Kahan and Klausner, 1997; Gulati and Scott, 2012; Gulati and Kahan, 2018; Kahan and Gulati, 2021). Even though we may know some of its causes, this lack of variation in h_{it} poses an analytical challenge.

However, there are some securities that have both a long time-series of available data and a variation in h_{it}. A security's price volatility is typically measured with respect to a trailing time frame: the annualized standard deviation of the logged daily price differences of the past n trading days. For price p, I calculate the n-day price volatility v_{it}^{n} of security i to be

$$v_{it}^{n} = \sigma\left(\left[\forall d \in [t-n,t] \mid \ln\left(\frac{p_{id}}{p_{id-1}}\right)\right]\right) * \sqrt{(}252). \tag{2}$$

Figure 4 shows h_{it} plotted over time, at a monthly level, alongside the daily volatility v_{it}^{10}, v_{it}^{90} for what I will call my showcase bond: CUSIP 13063A5G5. This security is a 30-year bond issued in 2009, maturing in 2039, rated AA-, and is the largest bond California currently has outstanding at 3 billion USD. Figure 3 shows h_{it} plotted alongside its price. Because price, yield, and volatility are reported at the daily level; for the time-series cross-sectional analysis I take the monthly average to match the level of analysis of the holdings data.

For the investor selection analysis, I merge the security-level data set with a data set describing the holders of debt (FactSet, 2022). Specifically, the data includes attributes of specific funds such as the PE ratio, PB ratio, Dividend-Yield ratio, market beta, and various descriptors of the price momentum of the fund's portfolio. This data set covers some 150,000 funds, only several dozen

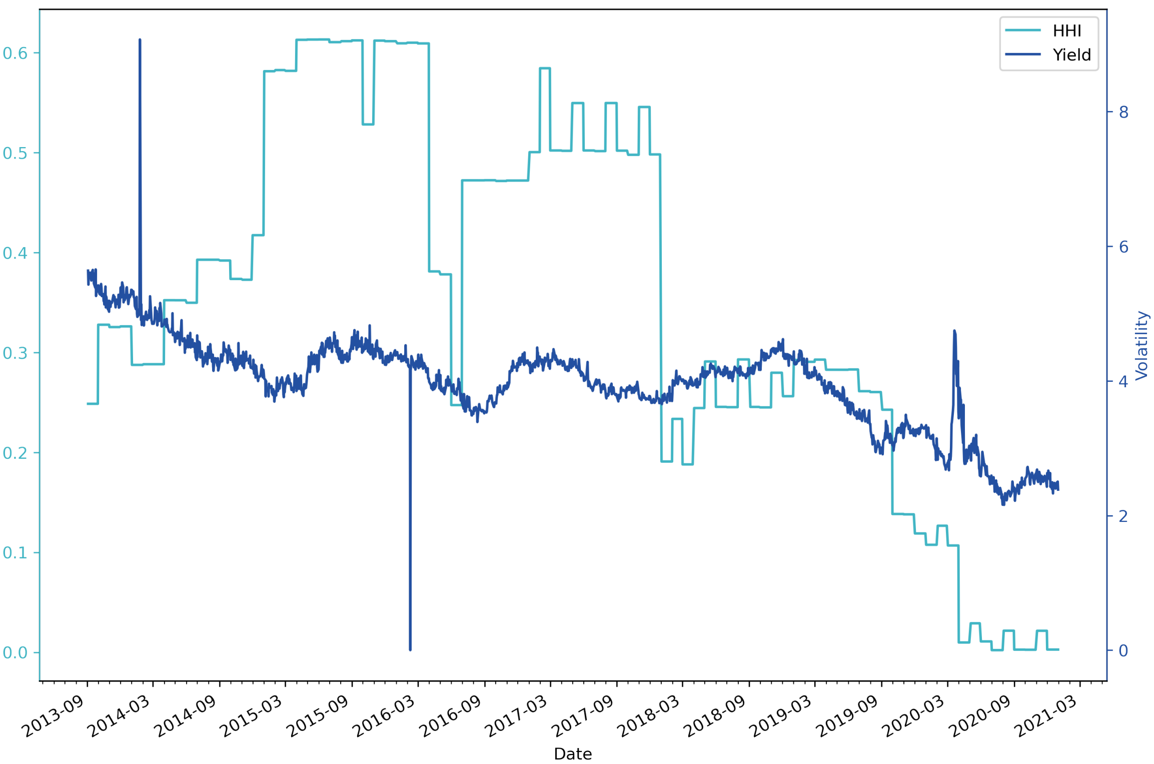

Figure 3: Price and ownership concentration over time.

of which have holdings in a given security large enough to report during any given period. One disadvantage of this data set is that it is a static snapshot of the most recent values for a given fund. Ideally, I would have all these attributes in time-series, but this may not be a large disadvantage because funds oftentimes pick portfolio attributes in advance of launching the fund and only rebalance periodically at the margins. As a result, I do not expect that the static nature of this data set will contaminate my results. Nonetheless, I do exclude some metrics that are time-dependent, such as those that illuminate the price movements of the fund's portfolio. This does not hamper testing my theory, however, because I am theoretically interested in more static attributes like general risk profile, as reflected in indicators PE ratio and beta.

I include all investors in my selection data set. This includes all "funds" (e.g. mutual funds, exchange-traded funds, pension funds. etc.) who have taken a position in a security large enough to require public reporting. Naturally, this results in the vast majority of funds without a position in a

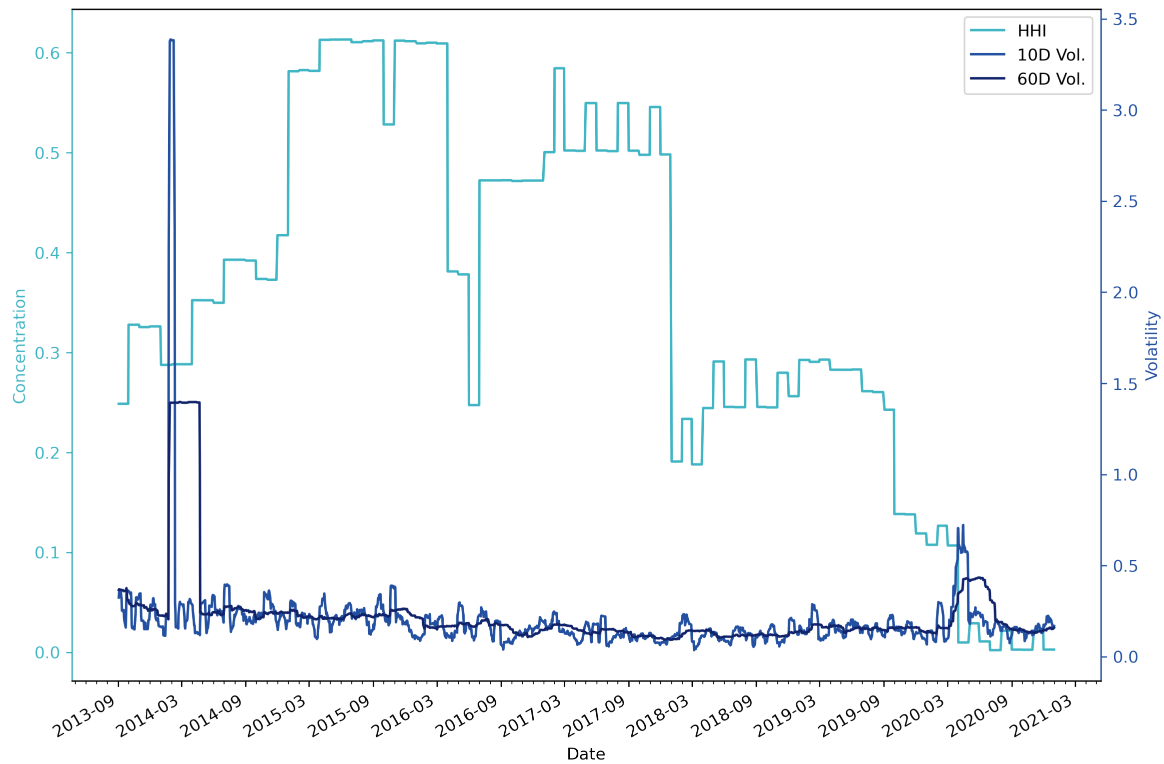

Figure 4: The volatility and ownership concentration over time.

given Californian bond during a given month. I focus on two major attributes of these funds: beta and PE ratio. The beta of a fund is a measure of its volatility in relationship to the overall market.[6] A PE ratio is a measure of a stock, but the aggregate PE ratio of a fund is the weighted average of the PE ratios of the stocks in its portfolio. Although not related to the bonds it holds, a fund's PE ratio is a loose measure of its risk. With a range set in advance of the fund's launch, the PE ratio has two attributes that makes it appealing for inclusion in this analysis. First, it is relatively static over time. Second, even though it does not pertain directly to bonds, it is a general measure of how a fund positions itself with regard to growth and value, and therefore risk.

[6]Values of beta greater than one mean the fund is more volatile than the market, a beta of one means that it is equally as volatile as the market, and values between zero and one mean the fund is less volatile than the market. Negative betas mean the fund exhibits an inverse relationship with the market (consider gold and the S&P 500).

5 Empirical Models and Results

I conduct three empirical analyses. The first assesses whether riskiness affects both investor selection into the market and investor's position size contingent upon entering the market. The second uses an approach new to Political Science literature to consider the effect of concentration on price volatility of the largest bond that California currently has outstanding. The third approach extends these analyses to a more conventional time-series cross-sectional analysis.

5.1 Investor Selection Results

To recap from earlier, an investor purchasing a position in a security is a two-step process consisting of a decision to enter the market or not and, contingent upon entering the market, a separate decision about the size of position to accrue. I expect that selection into the market is driven by risk attributes and position size is driven by different concerns. once they are in the market. Results supporting my hypothesis would show that the investor attributes driving market entrance are indeed different from those driving position size. Such results would provide basis to conclude that a security's "riskiness" is unrelated to the size of position that investors take, and therefore that riskiness cannot affect ownership concentration. Such findings would alleviate concerns about endogeneity and permit moving on to a security-level analysis of the relationship between ownership concentration and volatility.

To adjudicate my hypothesis, I employ a Heckman selection model to understand the two stages of my theorized investor-level decision-making process. In the first stage (the selection stage), investors decide whether or not to take a position in the security in question. In the second stage, those investors who have decided to take a position decide what size position to take.

There are some implementation issues, however. First, my data set is a panel data set of individual securities over time. The panel nature of this is complicating; it is difficult to estimate selection models for panel data. Because the data set of investors is so large, I settle on focusing on the dynamics of a representative security. I will call the security in question my showcase bond, introduced above. Second, the time-series nature of the data poses a concern: it is possible that an investor's lagged position predicts their current position. I include lagged position as a predictor in both stages of the analysis. I estimate separate models for each month and consider the effects of the variables over time.

I estimate a two-stage selection model, where the first stage predicts whether investor i enters the market at time t, given by $z_{it} \in \{0,1\}$, and the second stage predicts the investor's position in the security conditional on market presence, given by y_{it}^*:

$$z_{it} = \gamma X_i' + u_i, \tag{3}$$

$$y_{it}^* = \beta X_i' + \varepsilon_i, \tag{4}$$

where X_i is a vector of investor attributes (including position in the security at $t-1$), and y_{it}^* only observed if $z_{it} = 1$. I include as independent variables several attributes of portfolios: PE ratio, PB ratio, and Dividend-Yield ratio. This last measure has no relationship with the particular time value of a particular security and therefore is not post-treatment; it is merely a description of how much a fund's portfolio pays out in dividends relative to its return.

However, I expect the PE ratio to absorb most of the cross-investor variation in selection. Because it is important to consider the distribution of PE ratios across the market, I include in X_i a term for the squared PE ratio. My hypothesis will find support from a statistically significant relationships in γ but not β. More specifically, I expect to find that γ contains statistically significant relationships between PE ratio and market selection, with a positive coefficient for the first-order term and a negative relationship for the squared term. These choices reflect my expectation that PE ratios are normally distributed across the market.

Table 1 shows the results for both selection and outcome stages of the model for the latest time period in the data set, December 2020. The results indicate that investors with higher PE ratios are more likely to have a position, but the effect of PE ratio on position size is statistically indistinguishable from zero. The results of this regression indicate a statistically significant quadratic relationship between PE ratio and market selection, suggesting that a fund's placement within the market is closely related to how appropriate they view the security in question. For these investors, there is a sweet spot

Table 1: Regression Results, Heckman Selection, December 2020

	Dependent variable:	
	Investment Binary	Position Size
	probit Selection	*OLS* Outcome
	(1)	(2)
Beta	−0.00001	−129,769.5000
	(0.0014)	(280,298.6000)
PE Ratio	0.0429***	526,093.0000
	(0.0157)	(631,892.2000)
PE Ratio Sq.	−0.0005*	−6,571.7570
	(0.0003)	(7,539.9630)
PB Ratio	0.0054	−14,407.4900
	(0.0041)	(162,773.8000)
PB Ratio Sq.	−0.1834	−2,542,321.0000
	(0.1194)	(2,749,789.0000)
Dividend Yield	0.000000***	5.1343
	(0.000000)	(5.8762)
Lagged Holding		11,224,022.0000
		(15,878,167.0000)
Inv. Mills Ratio	−4.1126***	−48,562,448.0000
	(0.2292)	(68,718,866.0000)
Observations	116,928	21
R^2		0.9695
Adjusted R^2		0.9530
Log Likelihood	−181.8180	
Akaike Inf. Crit.	377.6360	
Residual Std. Error		216,018.7000
F Statistic		58.9808***
Note:		*p<0.1; **p<0.05; ***p<0.01

around 30: the average PE ratio for the Standard and Poor's 500 index in March 2022 is 24.56, implying that funds that focus on California municipals have a higher PE ratio on average than the market (Nasdaq, 2022). This suggests that Californian municipals are more likely to be in the portfolios of funds that have a growth mindset.

But Table 1 shows just a snapshot in time. Figure 5 shows how the effect of PE ratio on selection changes over time. Although not every time period shows a statistically significant relationship, those that do usually have a strong positive effect on selection. The squared value of the PE ratio also has a persistent, negative, statistically significant relationship with selection over time. Taken together, these results provide strong evidence that the attributes of an investor are closely tied to whether or

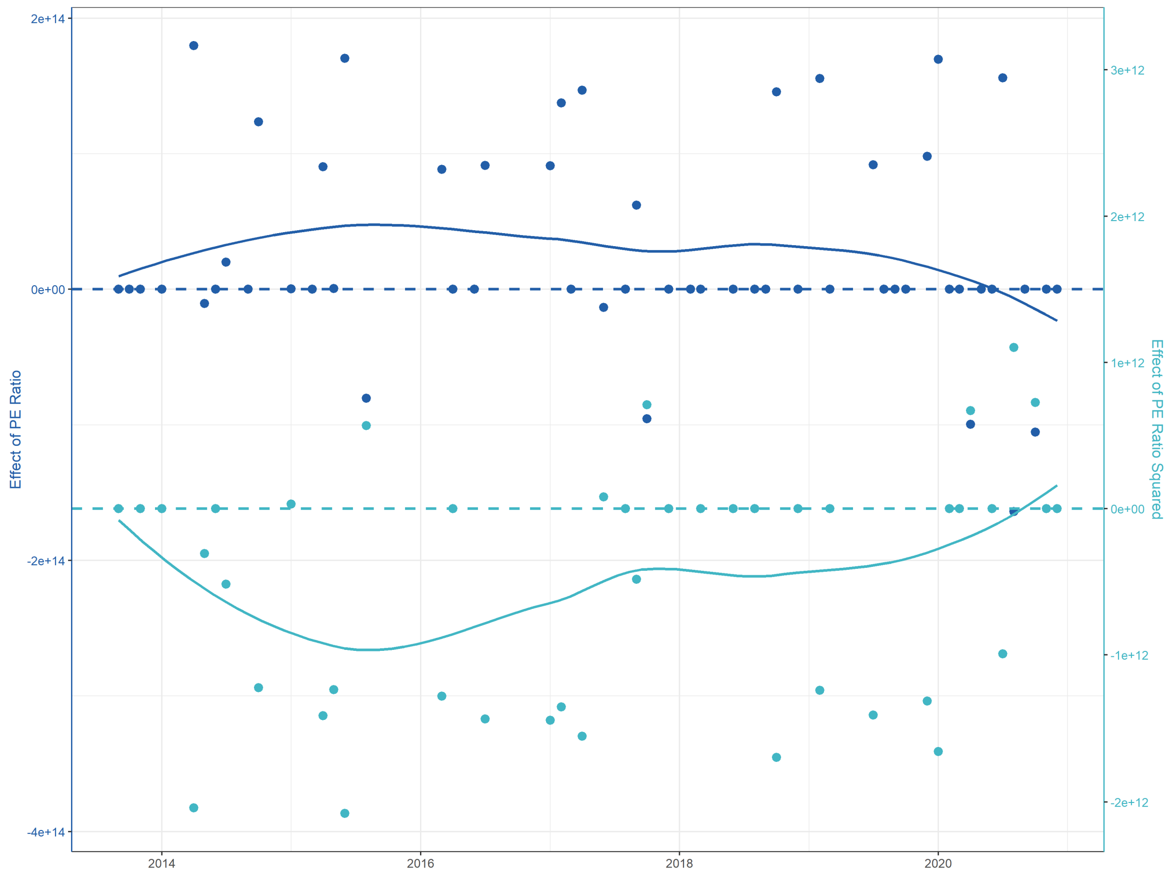

Figure 5: The Effect of PE Ratio and PE Ratio Squared on selection into the market over time.

not that investor enters the market for a security. [7]

Figure 6 shows that the non-relationship between PE ratio and portfolio size is persistent over time. Although the number of observations is far lower for the outcome model because of the limited number of investors who are required to report holdings, an inconclusive relationship is persistent over time. This suggests that other attributes affect portfolio allocation decisions, which is consistent with my expectations.

An investor's PE ratio has a statistically significant relationship with the selection stage but not with the outcome stage. This is consistent with my theoretical expectations: when a fund is created, investors decide what the fund's profile will be with regard to asset class, risk tolerance, and returns.

[7]Figures 5 and 6 have been filtered to only show the time periods with statistically significant relationships. For full results, see Figures 7 and 8 in the Appendix.

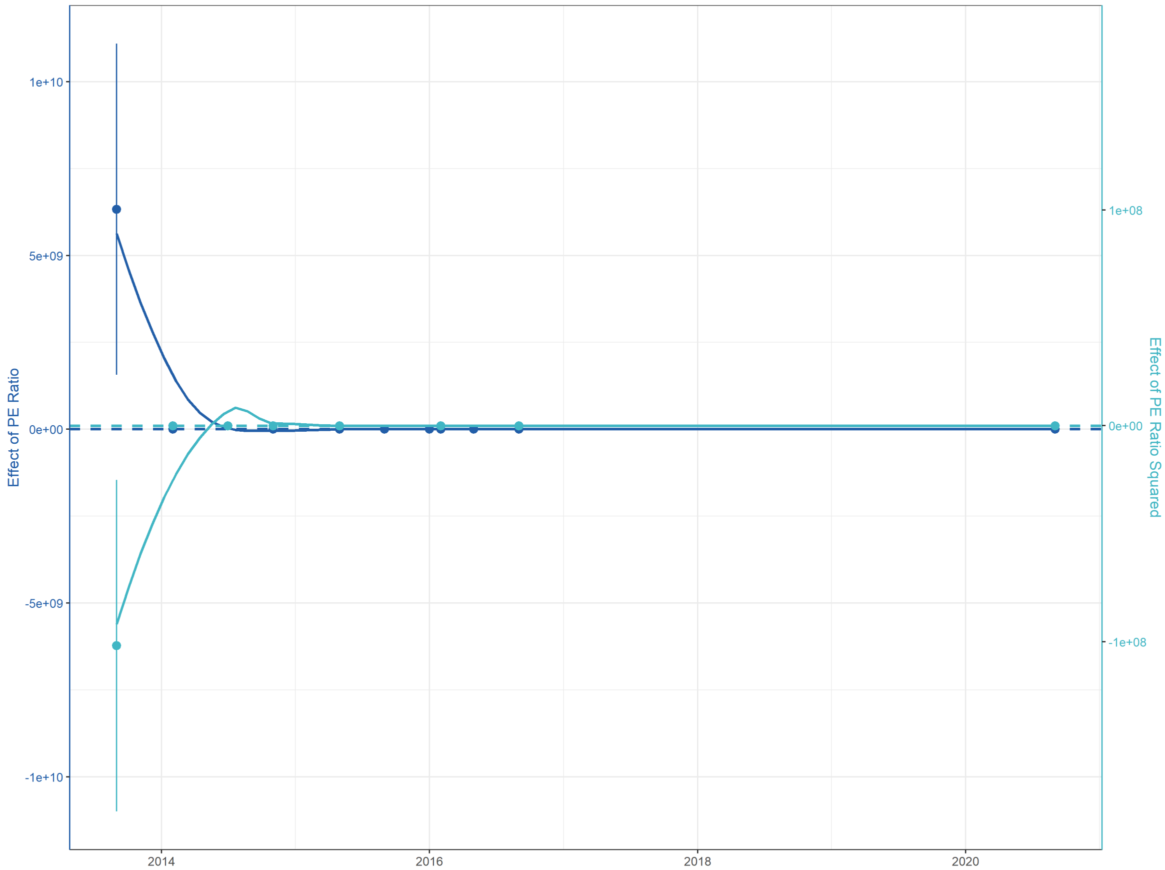

Figure 6: The Effect of PE Ratio and PE Ratio Squared on investor position size over time.

Such decisions are made well before particular securities are considered for inclusion in the portfolio. This explains why funds with a certain PE ratio are more likely to invest in Californian municipals. Furthermore, the results of the outcome stage of the selection model indicate that these same fund attributes are unrelated to position size. This supports my expectation that position size is determined by other factors aside from the risk matching between the investor and the security in question.

Moreover, these results provide support to my theoretical expectation that the "riskiness" of a debt security is not related to the size of position an investor takes. Without this relationship, it is not possible for riskiness to confound the results of the prior stage of analysis. These results undergird the empirical support for my expectation about the relationship between concentration and yield.

5.2 GARCH-MIDAS

5.2.1 Analysis

Now that there is some baseline support that ownership concentration is not affected by a risk match between a security and investors, I move on to examining the relationship between concentration and yield volatility.

Ownership concentration is reported at the monthly level but security prices fluctuate daily. To account for these different sampling frequencies of h_{it} and v_{it}^{n}, I use a technique common in the finance literature: a GARCH-MIDAS approach. GARCH-MIDAS[8] models combine two kinds of empirical models to suit this analysis situation (Engle et al., 2013). By assuming an ARMA (AutoRegressive Moving Average) structure for the error variance, GARCH models account for two tendencies of financial pricing data: autoregression and oscillation between periods of high and low volatility (Bollerslev, 1986).[9] MIDAS models allow for inclusion of variables sampled over different time periods, such as holdings data sampled monthly and pricing data sampled daily (Ghysels et al., 2004). GARCH-MIDAS models combine the two approaches by modeling short-term fluctuations of a GARCH component "around a time-varying long-term component that is a function of (macroeconomic or financial) variables" (Conrad and Kleen, 2020: p. 19). GARCH-MIDAS models have been increasingly popular in the finance literature to use macroeconomic variables to explain volatility of stocks (Girardin and Joyeux, 2013; Asgharian et al., 2013; Wang et al., 2020), cryptocurrencies (Conrad et al., 2018), exchange rates (You and Liu, 2020), and commodities (Pan et al., 2017), and tend to outperform other specifications (Conrad and Kleen, 2020: p. 20).[10]

Because of the intense computational requirements of the GARCH-MIDAS model, it is not possible to include a panel of covariates, so I consider only the relationship between ownership concentration and return, leaving inclusion of covariates for the time-series cross-sectional analysis in the next section. For log return on day i in month t, the GARCH-MIDAS approach follows the process given by:

[8] Generalized AutoRegressive Conditional Heteroskedasticity-MIxed DAta Sampling.

[9] Where volatility is a function of both the prior volatility and the value of the error in the prior period. This is opposed to the ARCH model, which assumes the error follows an AR model.

[10] See also Conrad and Kleen (2020); Virk et al. (2021).

$$r_{it} - E_{i-1,t}(r_{it}) = \sqrt{\tau_t g_{it}}\varepsilon_{it}, \forall i = 1,2,\ldots,N_t, \tag{5}$$

$$\varepsilon_{it}|\psi_{i-1,t} \sim N(0,1), \tag{6}$$

where N_t is the number of trading days in month t, $E_{i-1,t}()$ is the conditional expectation given information up to time $(i-1)$, and $\psi_{i-1,t}$ represents the information set up until day $i-1$ of period t. The volatility has two separate components: the short-term component g_{it} and the long-term component τ_t, initially assumed to be fixed for month t. Because I expect ownership concentration to influence short-term volatility, I choose a mean-reverting unit-variance GJR-GARCH(1,1) process for the short-term component g_{it}:

$$g_{it} = (1-\alpha-\gamma/2-\beta) + (\alpha+\gamma\mathbb{1}_{\{\varepsilon_{i-1,t}<0\}})\frac{\varepsilon_{i-1,t}^2}{\tau_t} + \beta g_{i-1,t}, \tag{7}$$

where $\alpha, \beta > 0$ and $\alpha+\beta < 1$. The long-term component τ_t follows the smoothed realized volatility of the MIDAS regression:

$$\tau_t = m + \theta \sum_{k=1}^{K} \phi_k(\omega_1,\omega_2) RV_{t-k}, \tag{8}$$

where RV_t denotes the fixed time realized volatility RV at time t:

$$RV_t = \sum_{i=1}^{N_t} r_{it}^2. \tag{9}$$

$\phi_k(\omega_1,\omega_2)$ is the function that defines the weighting scheme of MIDAS filters parameterized via the Beta weighting scheme (Conrad and Kleen, 2020):

$$\phi_k(\omega_1,\omega_2) = \frac{(k/K)^{\omega_1-1}(1-k/K)^{\omega_2-1}}{\sum_{j=1}^{K}(j/K)^{\omega_1-1}(1-j/K)^{\omega_2-1}}, \tag{10}$$

where the weights sum to one:

$$\sum_{\ell=1}^{K} \phi_\ell(\omega_1,\omega_2) = 1.$$

I follow the convention established in the literature by Engle (1982) and others in regressing

$$X_t = \sum_{j=1}^{12} \alpha_{jt} D_{jt} + \sum_{j=1}^{12} \beta_j X_{t-j} + \varepsilon_t, \tag{11}$$

where D_{jt} is a monthly dummy variable. The squared residuals ε_t^2 are taken as the proxy of volatility of macroeconomic variable X_t.

The GARCH-MIDAS approach can account for the mixed sampling frequencies of the data points, but cannot account for the panel structure of the data. Moreover, the varying availability of historical holdings data does not facilitate a GARCH-MIDAS approach for every security. I accordingly run a GARCH-MIDAS model for the spotlight bond mentioned above as a proof of concept. This approach obviously lacks the desirable quality of generalizability across the entire universe of Californian municipal bonds, but does showcase the time-series relationship between ownership concentration and volatility.

My hypotheses would find support from statistically significant estimations of α, β, γ, and θ.

5.2.2 Results

As a result, I present models for the two bonds with the longest time series of complete data. I analyze the same bond as in the selection stage, the showcase bond, which has 88 months of price and ownership data available.

Table 2 shows the GARCH-MIDAS output for the showcase bond.[11] Using the Bollerslev-Wooldridge reported in the "OPG SE" column as the reference standard errors, the parameters $\alpha, \beta,$ and γ are statistically significant, meaning that the model fits the data. The α and β terms sum to close to 1, confirming the existence of a strong volatility persistence effect.

My primary theoretical expectation is that securities with higher concentration will lead to higher short-term volatility because flow-induced trades happen over the short term. The statistically significant positive γ parameter means that a higher h has a greater short-term effect on volatility than a decrease does; in other words, h has an asymmetrical effect on this bond's returns.

[11] Using the mfGARCH R package (Kleen, 2020), with $K = 38$..

Table 2: GARCH-MIDAS Results, Spotlight Bond

Term	Estimate	Rob.SE	P-value	OPG SE	OPG P-value
μ	-0.032	0.028	0.260	0.034	0.350
α	0.201	0.139	0.149	0.048	0.000
β	0.521	0.213	0.014	0.070	0.000
γ	0.170	0.099	0.086	0.086	0.047
m	0.332	0.214	0.120	0.197	0.093
θ	-859.594	519.090	0.098	489.145	0.079
$w2$	319.369	44.618	0.000	301894.818	0.999

The results in table 2 also permit conclusions about the base long-term volatility m and the effect of h on long-term volatility, given by θ. The estimated value of m, the intercept of the long-term component of the volatility, is negative and statistically significant, suggesting that base long-term volatility is positive. The negative coefficient estimate for θ implies that higher values of h lead to less long-term volatility of returns. Although this seemingly contradicts my hypothesis, the statistical significance on these estimates is lower, giving less confidence in these results. At minimum, it merits further analysis.

The Variance Ratio (Engle et al., 2013) is often used to quantify the relative importance of the long-term and short-term volatility on returns, and is defined by:

$$VR = \frac{var(log(\tau_t))}{var(log(\tau_t g_t))} \tag{12}$$

For this specification, we find $VR = 10.71$, implying that 10.71% of expected variation in returns can be attributed to variation in h. These results provide preliminary support for my hypothesis: changes in h affect the volatility of returns of the spotlight bond. However, the diverging signs for the coefficient estimates of the short-term and long-term volatility of returns suggests that volatility over different time frames merits further consideration in the time-series cross-sectional analysis.

5.3 Time-Series Analysis

I next attempt to widen the scope of the analysis by considering the entirety of Californian bonds with available data. Because these bonds are all accessible by the same investors, it is possible that their prices, volatilities, and ownership patterns co-vary. To account for such a possibility, I test for cross-sectional dependence.[12] I find the securities are indeed cross-sectionally dependent, requiring an appropriate empirical approach.

Further estimation difficulties could arise from the varying reporting cadences of the data points: holdings data is reported monthly and price data daily. I take monthly averages of price data by security, and use time-series cross-sectional models to examine the effect of ownership concentration on a monthly average of n-day running volatility of log returns (denoted v_{it}^{n}). This approach has two advantages over the GARCH-MIDAS approach: it permits the examination of multiple securities at once and also allows for inclusion of covariates.

I test two specifications. First, I use a within-security approach to measure the effect of ownership concentration on the level of volatility. Second, given the statistically significant relationship between h and volatility in the GARCH-MIDAS results, I use a first-difference approach. The empirical model for the within approach is given by:

$$v_{it}^{n} = \alpha + \beta h_{it} + \gamma Z_{it} + \mu_i + \varepsilon_{it}, \tag{13}$$

My hypothesis would find support from a positive value of β, which would imply that higher levels of h are associated with higher volatilities v_{it}^{n}.

For the approach where I first-difference the time-variant variables, removing the time-invariant components of the regression, the empirical model is given by:

$$\Delta v_{it}^{n} = \beta \Delta h_{it} + \gamma \Delta Z_{it} + \Delta \mu_i + \Delta \varepsilon_{it}, \tag{14}$$

where $\Delta v_{it}^{n} = v_{it}^{n} - v_{i,t-1}^{n}$ and Z_{it} is the vector of aforementioned time-varying control variables. Instead of measuring the actual level of volatility estimated in the first time-series approach, the first-difference approach shown in Equation 14 measures the monthly *change* in volatility.

As a last step, I examine the effect of volatility on price and yield of a bond. I again use a first-difference approach to the empirical model:

[12] Using the PLM package from R (Croissant and Millo, 2008). Because the data set has a much larger n than T, I use both scaled LM and Pesaran's CD tests.

$$\Delta y_{it} = \beta \Delta v_{it}^{n} + \gamma \Delta Z_{it} + \Delta \mu_i + \Delta \varepsilon_{it}, \tag{15}$$

where $\Delta v_{it}^{n} = v_{it}^{n} - v_{i,t-1}^{n}$, etc. and Z_{it} is the vector of aforementioned time-varying control variables. Instead of measuring the actual level of the yield, the first-difference approach estimates monthly *change* in yield. I expect positive and statistically significant values of β, which would suggest that there is a positive volatility risk premium.

5.4 Time-Series Results

5.4.1 Predicting Volatility

The results based on Equation 13 do not show any statistically significant relationship between ownership concentration and volatility. Perhaps this is because, according to my theory, the driver of price variation is *changes* in the ownership concentration, which propel movements in price when holders buy and sell. I therefore focus on the results of the first difference specification, shown in Equation 14.

Table 3 shows the results of the first-difference model. Here, the results show the relationships between Δv_{it}^{n} and h_{it} are positive and statistically significant for $n = 10, 30$. Substantively, this means that an increase in the ownership concentration as measured by HHI corresponds to, on average across all bonds, an increase in ten-day and 30-day volatility. Interpreting the size of these coefficients is tricky, though; because HHI is measured on a scale from zero to one, a movement from zero to one is not very meaningful in the real world.[13] Moreover, volatilities also usually fall between zero and one. So, a one-one-hundredth of a unit change in HHI would correspond to a 0.12-unit increase in ten-day volatility, which is quite large. Changes in HHI have an effect almost twice as large on the ten-day volatility as the 30-day. Notwithstanding the variation in effect size, these two positive effects support my hypothesis.

[13]Zero would mean an infinitesimally small amount of the asset is held by an infinite number of holders, and one would mean that one person holds everything.

Table 3: The Effect of Ownership Concentration on Bond Price Volatility, FD Models, Robust SE

	DV: Number of Days Rolling Volatility Δv_{it}^{n}					
	3d	5d	10d	30d	60d	90d
	(1)	(2)	(3)	(4)	(5)	(6)
HHI	−5.3209	2.3315	11.9268***	6.8989*	5.0483	4.5194
	(5.4463)	(3.3506)	(3.6124)	(3.7384)	(3.4326)	(3.1180)
Pct. OS Known	0.6641*	0.0309	−0.8095***	−0.5639**	−0.4013*	−0.3643*
	(0.3706)	(0.2335)	(0.2412)	(0.2485)	(0.2259)	(0.2047)
Months to Maturity	−73.5956***	19.4323**	135.7924***	94.5575***	81.0514***	75.3684***
	(14.1408)	(8.8040)	(30.6310)	(22.4655)	(20.8633)	(19.8434)
Months to Maturity Sq.	−12.3064***	8.5741***	34.4706***	25.0800***	22.9463***	22.6083***
	(3.0597)	(1.8382)	(6.0283)	(4.4444)	(4.0405)	(3.9237)
Constant	0.0066***	−0.0007	−0.0099***	−0.0071***	−0.0060***	−0.0053***
	(0.0014)	(0.0009)	(0.0028)	(0.0021)	(0.0019)	(0.0018)

Note: *p<0.1; **p<0.05; ***p<0.01

Interestingly, the percentage we know of a bond's outstanding ownership is negatively correlated with volatilities measured in windows ten days or greater. This suggests that the more a bond is owned by funds we observe (i.e. funds that have stringent reporting requirements such as mutual funds, pension funds, insurance companies, etc.), the lower its volatility. Intuitively, this makes sense: securities owned by institutions likely are being held for long-term purposes and therefore see less churn in ownership, and correspondingly low volatility. On the other hand, securities that have large positions by hedge funds and others not required to report are likely more actively managed and pursued for short-term, activist, or other investment strategies that produce volatile prices.

Besides characteristics of the ownership structure, other variables hold statistically significant relationships with v_{it}^n. For all values of n, volatility follows a quadratic path over the maturity term of the bond. But this relationship looks different for different values of n. For $n > 3$, the initial effect of months to maturity on v_{it}^n is positive: volatility is higher when bonds are further away from maturity. But as the number of months until maturity decreases, volatility drops. This makes intuitive sense because as a bond gets closer to maturity, it is closer to redemption; therefore its price should converge to its face value. But for $n = 3$, the opposite is true. The initial effect of months to maturity on v_{it}^3 is negative: volatility is lower when bonds are further from maturity. As maturity approaches, volatility slowly approaches zero. This likely captures some variation in the way that lower values of n forget past data sooner.

5.4.2 Predicting Yield

Table 4 shows that increases in volatility, no matter what time window is considered, correspond strongly and statistically significantly with increases in yield. This supports my hypothesis: higher volatility corresponds to yield premia. A one-tenth of a unit increase in the annualized standard deviation of log day-over-day returns corresponds to at least a 2.8 basis point premium on yield. This effect is largest for 30-day volatility, where the overall effect is a 10 basis point increase.

Table 4: The Effect of Volatility on Yield, Robust SE

	DV: Monthly Average Close Yield					
	(1)	(2)	(3)	(4)	(5)	(6)
3-day Vol.	0.28***					
	(0.02)					
5-day Vol.		0.54***				
		(0.03)				
10-day Vol.			1.00***			
			(0.03)			
30-day Vol.				0.68***		
				(0.02)		
60-day Vol.					0.60***	
					(0.02)	
90-day Vol.						0.58***
						(0.03)
Months to Maturity	−88.35*	−91.40*	−94.16*	−86.71*	−84.30*	−82.71*
	(47.11)	(48.67)	(50.08)	(47.00)	(46.63)	(46.48)
Months to Maturity Sq.	−40.69***	−41.34***	−41.51***	−40.36***	−40.10***	−40.27***
	(9.12)	(9.37)	(9.63)	(9.13)	(9.07)	(9.05)
Constant	−0.01***	−0.01**	−0.01**	−0.01***	−0.01***	−0.01***
	(0.005)	(0.005)	(0.01)	(0.005)	(0.005)	(0.005)

Note: *p<0.1; **p<0.05; ***p<0.01

6 Conclusion

In conclusion, I find evidence that government bonds with higher ownership concentration are likely to have more volatile prices and higher yields. This connection helps shine a light on the pricing effect of the ownership structure of debt. Although some authors have studied the relationship between highly aggregated debt data and fiscal policy, and others have studied the determinants of bond volatility, none have yet linked the two fields of study.

I address this fundamental unanswered question in this paper. Testing my hypotheses on Californian municipal bonds issued since 2002, I find that a bond's volatility is likely to be higher if its ownership is more concentrated in the hands of fewer investors. This effect likely occurs because such a security is more susceptible to large pricing effects from large inflows and outflows. I also find evidence that higher security-level volatility corresponds to a yield premium to persuade creditors to invest.

It is possible, however, that a debt security's underlying "riskiness" drives both its yield and the set of investors willing to purchase it. I find evidence supporting my expectation that this is not the case; a risk match between an investor and a security may drive the investor to take a position in a security, but the size of the investor's position is driven by other things. Therefore, riskiness cannot drive a security's ownership concentration and cannot confound the relationship between ownership concentration and yield. I have taken further steps in empirical design and empirical methodology to reduce concerns about endogeneity, which nonetheless remain important to keep in mind when interpreting these results.

Each of the three empirical methods I use has its drawbacks. The investor selection model relies on time-invariant attributes of investors. Even though the investor attributes I consider do not change very much over time, I would still prefer to have a (regretfully unavailable) time-series data set. My second empirical analysis, the GARCH-MIDAS approach, does not permit the inclusion of covariates, which can be troublesome when considering the multitude of factors that likely affect security volatility and return. Furthermore, its inability to accommodate a panel approach propelled me to pursue a more conventional time-series cross-sectional analysis. The panel approach itself has its own drawbacks: it regresses a monthly metric on smoothed monthly measures of daily price volatility for many securities, several of which do not have much variation in ownership concentration over time. The biggest drawback of the empirical approach, however, is data availability. Ideally, I would have access to complete holdings data unencumbered by the variation in reporting requirements. But the only way I know of to find such data is through clearinghouses, which charge far more for data access than I could dream to pay. However, taken as a whole, the weight of the evidence supports the hypothesis that securities with higher ownership concentration have more volatile prices and higher yields.

This paper's theory and findings bear on the relationship between the concentration of debt holding and power, which has implications for the democratic peace literature and markets peace literature. For example, is there something called a debt peace? Does China have power over the United States because it holds a large amount of US government bonds? Or does the US instead have control over a portion of the Chinese balance sheet? Does the structure of debt ownership confer power or weakness?

This research has one major implication for governments. Other research has shown that government bond yields on primary markets are influenced by secondary market price dynamics such as the one this paper depicts. This paper suggests that governments whose debt is owned by a more concentrated group of investors could end up paying more in debt service costs in the long run.

Possible future work could include replication of this study with more high-fidelity holdings data. Also, an examination of the different aspects of the ownership structure of government debt would be productive; different kinds of debt holders have different goals, which likely means that they have divergent policy tolerances. A deeper understanding of these concepts could arise from similar investigations.

A Appendix

A.1 Selection Effect Graphs

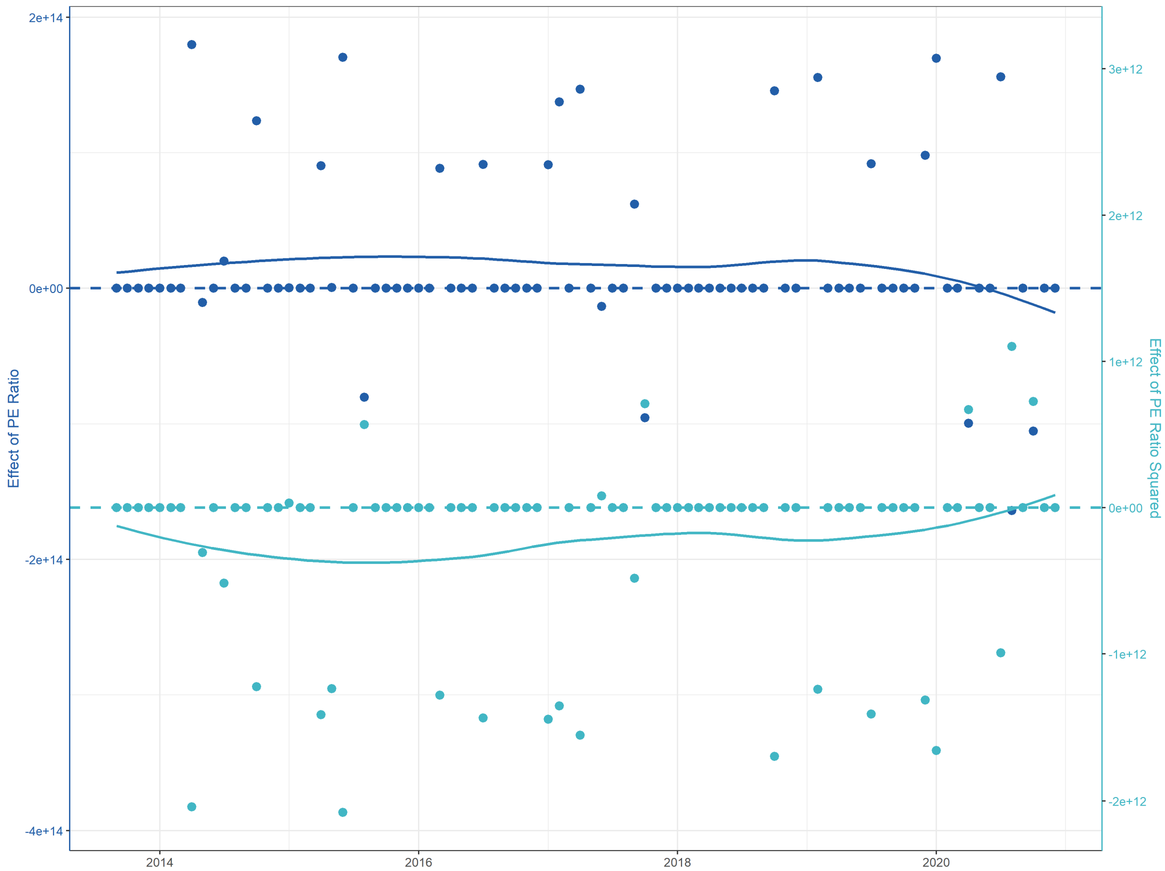

Figure 7: The Effect of PE Ratio and PE Ratio Squared on investor selection, over time.

A.2 Additional Time Series Results

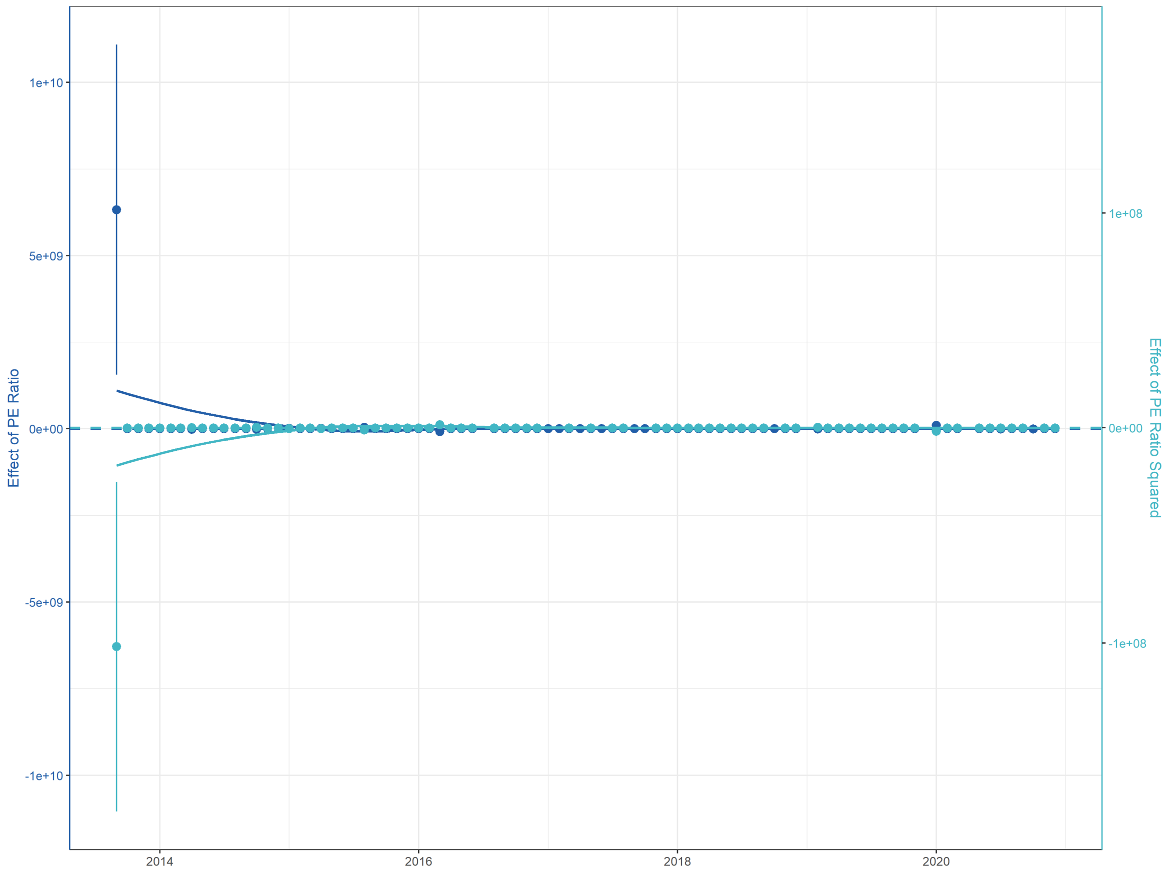

Figure 8: The Effect of PE Ratio and PE Ratio Squared on investor position size, over time.

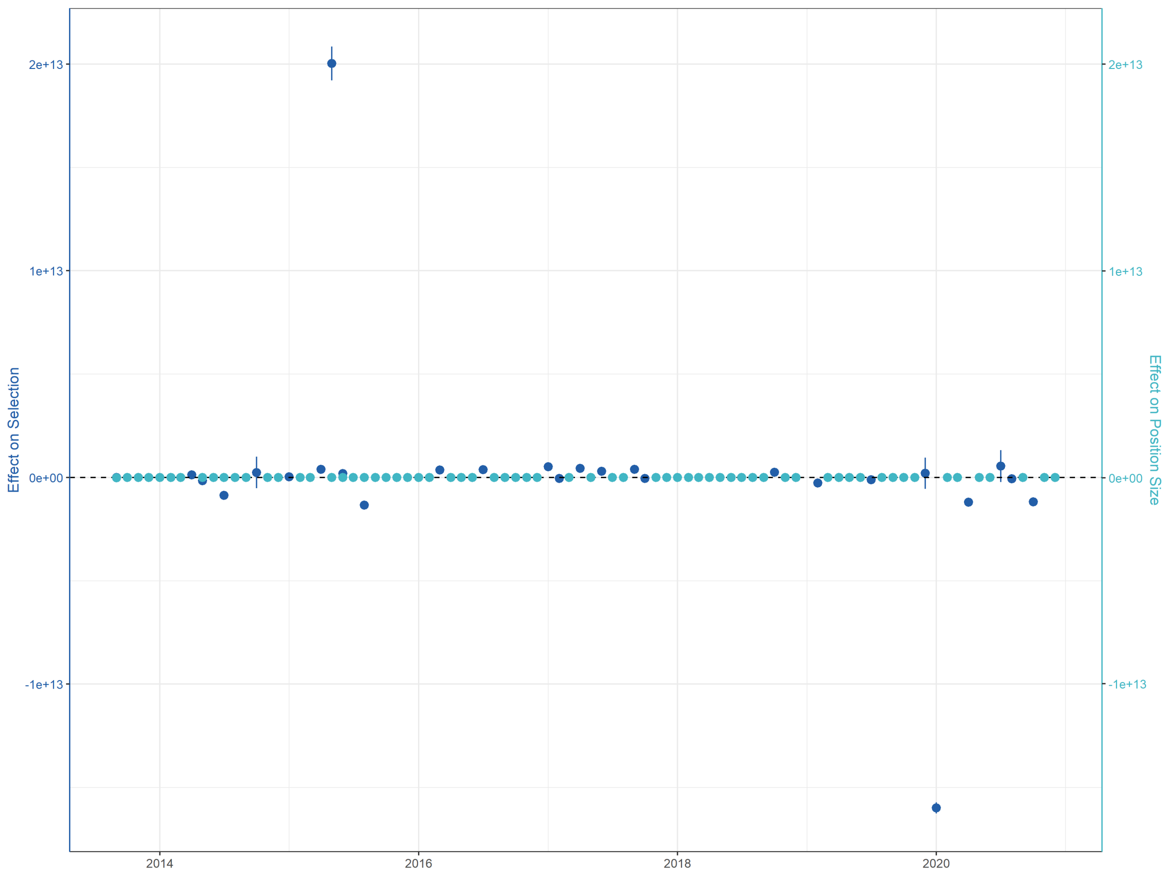

Figure 9: The Effect of Dividend Yield Ratio on investor selection and position size, over time.

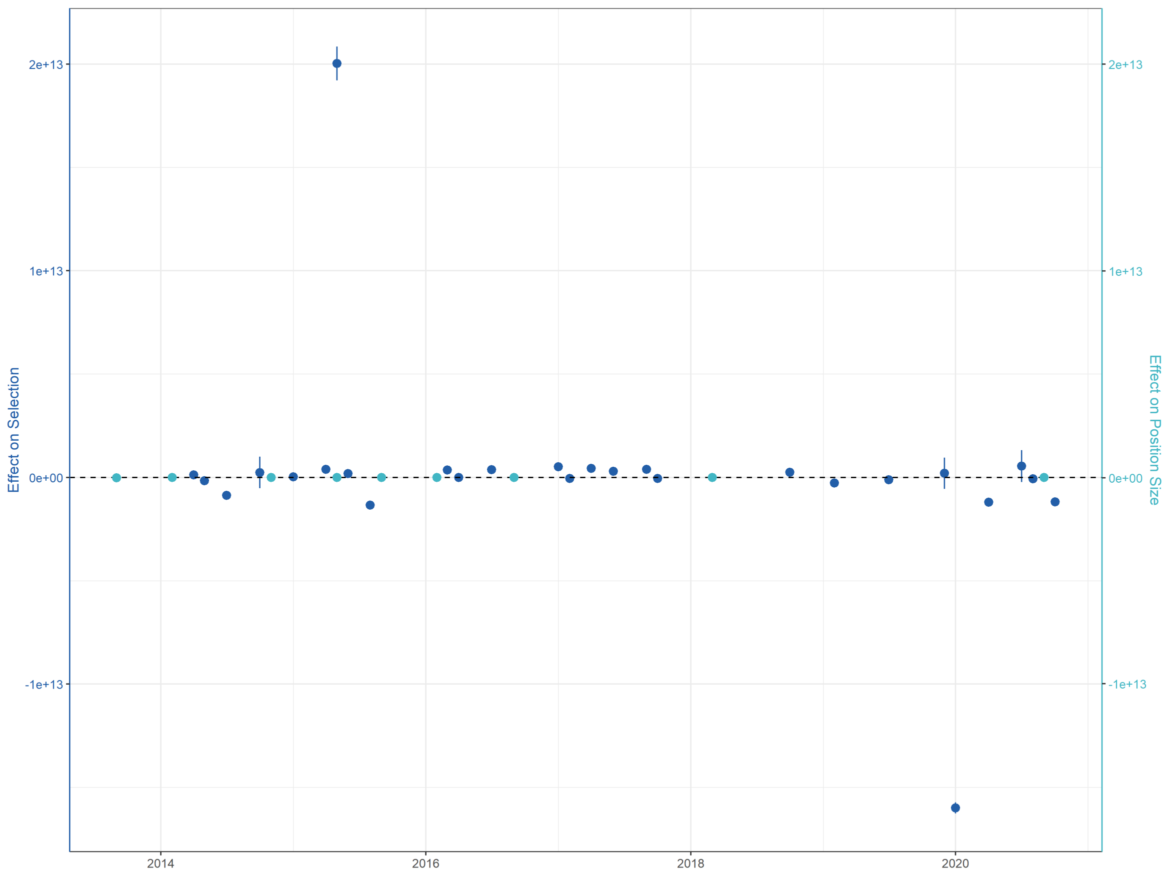

Figure 10: The Effect of Dividend Yield Ratio on investor selection and position size, over time. Statistically significant results only.

Table 5: The Effect of Ownership Concentration on Bond Price Volatility, Within Models

	DV: Number of Days Rolling Volatility Δv_{it}^{n}					
	3d	5d	10d	30d	60d	90d
	(1)	(2)	(3)	(4)	(5)	(6)
HHI	0.0000	−0.0000	−0.0000	0.0000	−0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Pct. OS Known	−0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Amt. Outstanding	0.0000***	0.0000***	0.0000***	0.0000***	0.0000***	0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Mty. Size	0.0000***	0.0000***	0.0000***	0.0000***	0.0000***	0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Months to Maturity Sq.	−0.0000	−0.0000	0.0000	−0.0000	−0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Observations	2,493	2,493	2,493	2,493	2,493	2,493
R^2	0.1676	0.1578	0.0633	0.2342	0.2084	0.0632
Adjusted R^2	0.1001	0.0895	−0.0127	0.1721	0.1442	−0.0128
F Statistic	85.4500***	149.0904***	95.4925***	89.3366***	126.4709***	35.3654***

Note: *p<0.1; **p<0.05; ***p<0.01

Table 6: The Effect of Ownership Concentration on Bond Price Volatility, Within Models, Robust SE

	DV: Number of Days Rolling Volatility Δv_{it}^{n}					
	3d	5d	10d	30d	60d	90d
	(1)	(2)	(3)	(4)	(5)	(6)
HHI	0.0000	−0.0000	−0.0000	0.0000	−0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Pct. OS Known	−0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Amt. Outstanding	0.0000***	0.0000***	0.0000***	0.0000***	0.0000***	0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Mty. Size	0.0000***	0.0000***	0.0000**	0.0000***	0.0000***	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Months to Maturity Sq.	−0.0000	−0.0000	0.0000	−0.0000	−0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)

Note: *p<0.1; **p<0.05; ***p<0.01

Table 7: The Effect of Ownership Concentration on Bond Price Volatility, FD Models

	DV: Number of Days Rolling Volatility Δv_{it}^{n}					
	3d	5d	10d	30d	60d	90d
	(1)	(2)	(3)	(4)	(5)	(6)
HHI	-5.3209	2.3315	11.9268***	6.8989***	5.0483***	4.5194***
	(4.4114)	(3.4294)	(2.7853)	(2.5991)	(1.9113)	(1.6265)
Pct. OS Known	0.6641**	0.0309	-0.8095***	-0.5639***	-0.4013***	-0.3643***
	(0.3299)	(0.2565)	(0.2083)	(0.1944)	(0.1430)	(0.1216)
Months to Maturity	-73.5956***	19.4323	135.7924***	94.5575***	81.0514***	75.3684***
	(19.1687)	(14.9015)	(12.1026)	(11.2936)	(8.3050)	(7.0673)
Months to Maturity Sq.	-12.3064**	8.5741**	34.4706***	25.0800***	22.9463***	22.6083***
	(5.4864)	(4.2650)	(3.4640)	(3.2324)	(2.3770)	(2.0228)
Constant	0.0066**	-0.0007	-0.0099***	-0.0071***	-0.0060***	-0.0053***
	(0.0028)	(0.0022)	(0.0018)	(0.0017)	(0.0012)	(0.0010)
Observations	2,397	2,397	2,397	2,397	2,397	2,397
R^2	0.0084	0.0033	0.0560	0.0330	0.0461	0.0577
Adjusted R^2	0.0068	0.0016	0.0544	0.0314	0.0445	0.0561
F Statistic	5.0904***	1.9598*	35.4704***	20.4050***	28.8699***	36.6211***

Note: *p<0.1; **p<0.05; ***p<0.01

Table 8: The Effect of Volatility on Yield

	DV: Monthly Average Close Yield					
	(1)	(2)	(3)	(4)	(5)	(6)
3-day Vol.	0.28***					
	(0.02)					
5-day Vol.		0.54***				
		(0.02)				
10-day Vol.			1.00***			
			(0.03)			
30-day Vol.				0.68***		
				(0.03)		
60-day Vol.					0.60***	
					(0.04)	
90-day Vol.						0.58***
						(0.05)
Months to Maturity	−88.35***	−91.40***	−94.16***	−86.71***	−84.30***	−82.71***
	(22.81)	(22.49)	(21.86)	(22.56)	(22.86)	(22.93)
Months to Maturity Sq.	−40.69***	−41.34***	−41.51***	−40.36***	−40.10***	−40.27***
	(6.98)	(6.88)	(6.69)	(6.90)	(7.00)	(7.02)
Constant	−0.01***	−0.01***	−0.01***	−0.01***	−0.01***	−0.01***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Observations	10,856	10,856	10,856	10,856	10,856	10,856
R^2	0.03	0.05	0.11	0.05	0.02	0.02
Adjusted R^2	0.03	0.05	0.11	0.05	0.02	0.02
F Statistic	102.07***	208.45***	434.92***	183.78***	84.66***	63.68***

Note: *p<0.1; **p<0.05; ***p<0.01

References

Abbot, T. (2017). Heterogeneous risk preferences in financial markets. In Paris December 2016 Finance Meeting EUROFIDAI-AFFI.

Almeida, C. and Vicente, J. (2009). Identifying volatility risk premia from fixed income asian options. Journal of Banking & Finance, 33(4):652–661.

Asgharian, H., Hou, A. J., and Javed, F. (2013). The importance of the macroeconomic variables in forecasting stock return variance: A garch-midas approach. Journal of Forecasting, 32(7):600–612.

Ballard-Rosa, C., Mosley, L., and Wellhausen, R. (2021). Coming to terms: The politics of sovereign bond denomination. International Organization.

Bechtel, M. M. (2009). The political sources of systematic investment risk: lessons from a consensus democracy. The Journal of Politics, 71(2):661–677.

Bloomberg (2022). Bloomberg terminal data.

Boermans, M. A. (2015). Market concentration in the euro area bond markets. an application with granular sectoral securities holdings statistics. In IFC workshop on Combining micro and macro statistical data for financial stability analysis. Experiences, opportunities and challenges. Warsaw. Poland, pages 14–15.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of econometrics, 31(3):307–327.

Brock, W. A. and Durlauf, S. N. (2010). Adoption curves and social interactions. Journal of the European Economic Association, 8(1):232–251.

Brooks, S. M., Cunha, R., and Mosley, L. (2019). Elections, ideology and experience? sovereign bond investors and government change. In Political Economy of International Organization Meeting.

Campbell, J. Y. and Viceira, L. (2001). Who should buy long-term bonds? American Economic Review, 91(1):99–127.

Campello, D. (2015). The politics of market discipline in Latin America: globalization and democracy. Cambridge University Press.

Cetorelli, N., Hirtle, B., Morgan, D. P., Peristiani, S., and Santos, J. A. (2007). Trends in financial market concentration and their implications for market stability. Economic Policy Review, 13(1).

Chabakauri, G. (2015). Asset pricing with heterogeneous preferences, beliefs, and portfolio constraints. Journal of Monetary Economics, 75:21–34.

Chan, Y. L. and Kogan, L. (2001). Heterogeneous preferences and the dynamics of asset prices. NBER Working Paper, (8607).

Chan, Y. L. and Kogan, L. (2002). Catching up with the joneses: Heterogeneous preferences and the dynamics of asset prices. Journal of Political Economy, 110(6):1255–1285.

Chen, X., Huang, J.-Z., Sun, Z., Yao, T., and Yu, T. (2020). Liquidity premium in the eye of the beholder: An analysis of the clientele effect in the corporate bond market. Management Science, 66(2):932–957.

Christensen, P. O., Larsen, K., and Munk, C. (2012). Equilibrium in securities markets with heterogeneous investors and unspanned income risk. Journal of Economic Theory, 147(3):1035–1063.

Chung, K. H., Wang, J., and Wu, C. (2019). Volatility and the cross-section of corporate bond returns. Journal of Financial Economics, 133(2):397–417.

Cole, H. L., Neuhann, D., and Ordonez, G. (2020). Information spillovers in sovereign debt markets.

Condie, S. (2008). Living with ambiguity: prices and survival when investors have heterogeneous preferences for ambiguity. Economic Theory, 36(1):81–108.

Conrad, C., Custovic, A., and Ghysels, E. (2018). Long-and short-term cryptocurrency volatility components: A garch-midas analysis. Journal of Risk and Financial Management, 11(2):23.

Conrad, C. and Kleen, O. (2020). Two are better than one: Volatility forecasting using multiplicative component garch-midas models. Journal of Applied Econometrics, 35(1):19–45.

Croissant, Y. and Millo, G. (2008). Panel data econometrics in r: The plm package. Journal of statistical software, 27(2).

Cvitanić, J., Jouini, E., Malamud, S., and Napp, C. (2012). Financial markets equilibrium with heterogeneous agents. Review of Finance, 16(1):285–321.

Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. The Journal of finance, 65(4):1237–1267.

Eichengreen, B. and Hausmann, R. (1999). Exchange rates and financial fragility.

Eisl, A., Ochs, C., Osadchiy, N., and Subrahmanyam, M. G. (2019). The linkage between primary and secondary markets for eurozone sovereign debt: Free flow or bottleneck? Technical report, mimeo.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. Econometrica: Journal of the econometric society, pages 987–1007.

Engle, R. F., Ghysels, E., and Sohn, B. (2013). Stock market volatility and macroeconomic fundamentals. Review of Economics and Statistics, 95(3):776–797.

FactSet (2022). Ownership datafeed v5.

Ferrara, F. M. and Sattler, T. (2018). The political economy of financial markets.

Fischer, A., Arnold, A., and Gibbs, M. (1996). Information and the speed of innovation adoption. American Journal of Agricultural Economics, 78(4):1073–1081.

Frijns, B., Koellen, E., and Lehnert, T. (2008). On the determinants of portfolio choice. Journal of Economic Behavior & Organization, 66(2):373–386.

Ghysels, E., Santa-Clara, P., and Valkanov, R. (2004). The midas touch: Mixed data sampling regression models.

Girardin, E. and Joyeux, R. (2013). Macro fundamentals as a source of stock market volatility in china: A garch-midas approach. Economic Modelling, 34:59–68.

Greenwood, R. and Thesmar, D. (2011). Stock price fragility. Journal of Financial Economics, 102(3):471–490.

Greenwood, R. and Vayanos, D. (2010). Price pressure in the government bond market. American Economic Review, 100(2):585–90.

Greenwood, R. M. and Vissing-Jorgensen, A. (2018). The impact of pensions and insurance on global yield curves. Harvard Business School Finance Working Paper, (18-109).

Gulati, M. and Kahan, M. (2018). Cash America and the structure of bondholder remedies. Capital Markets Law Journal, 13(4):570–586.

Gulati, M. and Scott, R. E. (2012). The three and a half minute transaction: Boilerplate and the limits of contract design. University of Chicago Press.

Hardie, I. (2006). The power of the markets? the international bond markets and the 2002 elections in brazil. Review of International Political Economy, 13(1):53–77.

Hauser, S. and Kedar-Levy, H. (2018). Liquidity might come at cost: The role of heterogeneous preferences. Journal of Financial Markets, 39:1–23.

Hirschman, A. O. (1970). Exit, voice, and loyalty: Responses to decline in firms, organizations, and states, volume 25. Harvard university press.

Isaenko, S. (2008). The term structure of interest rates in a pure exchange economy where investors have heterogeneous recursive preferences. The Quarterly Review of Economics and Finance, 48(3):457–481.

Kahan, M. and Gulati, M. (2021). Contracts of inattention. Law & Social Inquiry, pages 1–26.

Kahan, M. and Klausner, M. (1997). Standardization and innovation in corporate contracting (or" the economics of boilerplate"). Virginia Law Review, pages 713–770.

Kaplan, S. B. (2013). Globalization and austerity politics in Latin America. Cambridge University Press.

Kaplan, S. B. and Thomsson, K. (2017). The political economy of sovereign debt: global finance and electoral cycles. The journal of politics, 79(2):605–623.

Kleen, O. (2020). mfgarch: Mixed-frequency garch models. r package version 0.2.0. Technical report.

Levy, M. and Levy, H. (1996). The danger of assuming homogeneous expectations. Financial Analysts Journal, 52(3):65–70.

Lou, D., Yan, H., and Zhang, J. (2013). Anticipated and repeated shocks in liquid markets. The Review of Financial Studies, 26(8):1891–1912.

McDade, T. R., Mosley, L., and Rosendorff, B. P. (2021). Political shocks and sovereign borrowing. Working Paper.

Modigliani, F. and Sutch, R. (1966). Innovations in interest rate policy. The American Economic Review, 56(1/2):178–197.

Mosley, L. (2000). Room to move: International financial markets and national welfare states. International organization, 54(4):737–773.

Mosley, L., Paniagua, V., and Wibbels, E. (2020). Moving markets? government bond investors and microeconomic policy changes. Economics & Politics.

Municipal Securities Research Bureau (2022). Municipal securities transactions data.

Nasdaq (2022). PE Ratios of the S&P 500 Index.

O'Connor, J. (1979). The fiscal crisis of the state. Transaction Publishers.

Pan, Z., Wang, Y., Wu, C., and Yin, L. (2017). Oil price volatility and macroeconomic fundamentals: A regime switching garch-midas model. Journal of Empirical Finance, 43:130–142.

Peltonen, T. A., Scheicher, M., and Vuillemey, G. (2014). The network structure of the cds market and its determinants. Journal of Financial Stability, 13:118–133.

Przeworski, A. and Wallerstein, M. (1988). Structural dependence of the state on capital. American Political Science Review, 82(1):11–29.

Roberts, B. E. (1990). Political institutions, policy expectations, and the 1980 election: a financial market perspective. American Journal of Political Science, pages 289–310.

Roos, J. E. (2019). Why Not Default? Princeton University Press.

Sangvinatsos, A. and Wachter, J. A. (2005). Does the failure of the expectations hypothesis matter for long-term investors? The Journal of Finance, 60(1):179–230.

Sarasvathy, S. D., Dew, N., Velamuri, S. R., and Venkataraman, S. (2010). Three views of entrepreneurial opportunity. In Handbook of entrepreneurship research, pages 77–96. Springer.

Satchell, S. and Scowcroft, A. (2000). A demystification of the black-litterman model: Managing quantitative and traditional portfolio construction. Journal of Asset Management, 1(2):138–150.

Sigaux, J.-D. (2020). Trading ahead of treasury auctions. ECB Working Paper Series. Available at SSRN 2789988.

Stevenson, A. (2016). How argentina settled a billion-dollar debt dispute with hedge funds. New York Times.

Sushko, V. and Turner, G. (2018). The implications of passive investing for securities markets. BIS Quarterly Review, March.

Vayanos, D. and Woolley, P. (2013). An institutional theory of momentum and reversal. The Review of Financial Studies, 26(5):1087–1145.

Virk, N., Javed, F., Awartani, B., et al. (2021). A reality check on the garch-midas volatility models. Technical report.

Wachter, J. A. (2003). Risk aversion and allocation to long-term bonds. Journal of Economic Theory, 112(2):325–333.

Wang, L., Ma, F., Liu, J., and Yang, L. (2020). Forecasting stock price volatility: New evidence from the garch-midas model. International Journal of Forecasting, 36(2):684–694.

Weinbaum, D. (2009). Investor heterogeneity, asset pricing and volatility dynamics. Journal of Economic Dynamics and Control, 33(7):1379–1397.

Wellhausen, R. L. (2015). Bondholders vs. direct investors? competing responses to expropriation. International Studies Quarterly, 59(4):750–764.

You, Y. and Liu, X. (2020). Forecasting short-run exchange rate volatility with monetary fundamentals: A garch-midas approach. Journal of Banking & Finance, 116:105849.