

Bank of England

Fire sales of safe assets

Staff Working Paper No. 1,089

July 2024

Gabor Pinter, Emil Siriwardane and Danny Walker

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee.



Bank of England

Staff Working Paper No. 1,089

Fire sales of safe assets

Gabor Pinter,⁽¹⁾ Emil Siriwardane⁽²⁾ and Danny Walker⁽³⁾

Abstract

We use trade-level data to study price pressure effects in the UK gilt market from September to October 2022. During this period, forced sales by liability-driven investment funds (LDI funds) led to price discounts on the order of 10%, accounting for roughly half the total decline in gilt prices. Balance sheet segmentation and operational issues slowed equity injections into LDIs by well-capitalised pension investors, leading LDIs to instead sell gilts. This effect was most pronounced for pooled LDIs, which invest on behalf of multiple pension schemes, because of co-ordination problems between pensions. Hedge funds also appear to have delayed entry to time the bottom of the fire sale. Overall, our findings illustrate how capital can be slow moving internally, due to contracting frictions, and externally, due to strategic arbitrageur behaviour.

Key words: Crisis, fire sale, operational frictions, LDI, government bonds.

JEL classification: D83, D84, G01, G11, G12, G14.

(1) Bank for International Settlements. Email: gabor.pinter@bis.org

(2) Harvard Business School and NBER. Email: esiriwardane@hbs.edu

(3) Bank of England. Email: daniel.walker@bankofengland.co.uk

We are grateful to Saba Alam, Andrew Bailey, Yuliya Baranova, David Baumslag, Juliane Begenau, Sarah Breeden, Rob Czech, Leo Fernandes, Robin Greenwood, Sam Hanson, Andrew Harley, Andrew Hauser, Simon Jurkatis, Anil Kashyap, Ali Moussavi, Ioana Neamtu, Pierre Ortlieb, Piotr Piecha, Huw Pill, Matt Roberts-Sklar, Andrea Rosen, Adi Sunderam, Erik Stafford, Jeremy Stein, Boris Vallee, Dimitri Vayanos and Nick Vause for helpful comments and conversations. The views expressed in this paper are those of the authors, and not those of the Bank of England or its policy committees.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH

Email: enquiries@bankofengland.co.uk

©2024 Bank of England

ISSN 1749-9135 (on-line)

1 Introduction

From September to October 2022, UK gilt markets experienced historically unprecedented volatility, with yields on 30-year gilts rising by over 100 basis points in four days (see Figure 1a). These unprecedented market moves have largely been attributed to two factors: the September 23 announcement of expansionary fiscal policy—the so-called “mini-budget”—and forced sales by liability-driven investment funds (LDIs) that were used by UK corporate defined-benefit pension schemes for asset-liability matching (Breedon, 2022; Hauser, 2022). In response, the Bank of England (the “Bank”) executed a targeted intervention in the gilt market within a preset timeframe from September 28 to October 14 (Alexander et al., 2023).

What proportion of gilt price fluctuations was due to selling pressure from LDIs? Why couldn’t LDIs recapitalize quickly enough to avoid forced sales? What led to the scarcity of liquidity providers? The answers to these questions are important for several reasons. For one, they shed light on the mechanics of fire sales, the process of arbitrage, and by extension, the limits of market efficiency (Shleifer and Vishny, 1992, 1997). If, for instance, material and persistent mispricing can occur in a market as safe and liquid as gilts, then failures of arbitrage are likely more frequent and severe for a wide range of assets. In addition, these questions are useful for evaluating the Bank’s intervention strategy and guiding future governmental responses to fire sales.

In this paper, we provide answers using regulatory data that covers virtually all UK gilt, repo, and derivatives transactions. This data offers us the ability to observe the intricacies of a fire sale and the dynamics of arbitrage with an unusually high level of precision. We start by documenting how the balance sheet of the LDI sector evolved over this period. LDIs entered September with around £300 billion (bn) in assets, mostly in cash gilts (88%), funded by an equal blend of repo borrowing and equity from UK corporate pension schemes. Their asset-to-equity ratio, initially below 2, rose steadily throughout September and then spiked sharply following the mini-budget to 2.8, before rapidly returning to its initial level by the end of October. This swift deleveraging was achieved in large part by gilt sales. The trading data further show that gilt sales commenced immediately following the mini-budget on September 23, with LDIs subsequently selling £25.3 bn

of gilts in par terms and reducing their repo borrowing by £32.6 bn.

After confirming the timing and size of LDI sales, we estimate the impact of LDI selling on gilt prices. To understand the identification problem, consider two gilts, A and L , and assume LDIs exclusively hold L . Suppose that a negative fundamental shock like the mini-budget then hits the UK economy, lowering the price of both gilts, raising LDI leverage, and forcing LDIs to sell L . If A and L were equally exposed to the fundamental shock and there are no spillover effects as in Vayanos and Vila (2021), then the impact of LDI selling could be measured simply by comparing how much further L falls in price relative to A . However, if L was more exposed to the fundamental shock, then its price would fall more than A 's even in the absence of LDI selling. This makes it hard to distinguish the impact of LDI selling on L 's price (relative to A 's) from that of the fundamental shock.

Our solution to this identification problem is to compare two gilts with very similar maturities within the same hour on the same day, differing primarily in how much they were held by LDIs before the crisis. The basic idea is that gilts positioned nearby along the yield curve should be equally exposed to any fundamental shocks hitting the UK economy, most notably the mini-budget. We are careful to define exposure to LDIs using positions at the beginning of the month, thereby alleviating concerns that LDIs had private information about the fundamental value of the gilts they sold after the mini-budget. Our most conservative difference-in-differences estimates suggest that, at the peak of the fire sale, LDI selling led to discounts of 6.87%. This discount was completely eliminated by the end of October, weighing further against the idea that LDIs had private information and instead supporting a fire sale mechanism.¹ A back-of-the-envelope calculation using this price pressure estimate further suggests that LDIs were responsible for at least half the decline in long-dated gilt prices that occurred in the days following the mini-budget.

In light of the significant discounts caused by LDI sales, we explore the factors that delayed capital inflow into the market, both in terms of mitigating the fire sale's impact and in preventing

¹The credibility of our difference-in-difference estimate rests on the assumption that, absent LDI selling, the prices of gilts with similar maturities would have followed similar paths. We show that our estimates are robust to deviations from this parallel trends assumption using the techniques proposed by Rambachan and Roth (2023).

it from occurring in the first place. With regard to the question of why natural buyers were scarce, some may have hesitated to enter the market because they could not differentiate the effects of the mini-budget from those of the fire sale. Yet, a deep understanding of the fiscal shock was not crucial for recognizing the clear mispricing between gilts with similar maturities. Another possibility is that potential gilt buyers were uniformly distressed, a situation echoing the 2008-2009 Global Financial Crisis (Mitchell and Pulvino, 2012). However, the circumstances surrounding the UK's mini-budget were more geographically localized. Moreover, hedge funds were heavily short duration entering the crisis and thus profited greatly from the rise in gilt yields. High-frequency trading data suggests that they may have nevertheless withheld liquidity provision to the gilt market in an attempt to time the bottom of the fire sale. This behavior aligns with theories suggesting that strategic considerations may cause arbitrage capital to flow slowly to known mispricings (Abreu and Brunnermeier, 2003).

In the latter portion of the paper, we investigate why LDIs were not able to raise capital quickly enough to avoid forced sales. One obvious answer is that investors in LDIs, namely corporate pension schemes, did not have sufficient assets to recapitalize the LDIs. However, a closer examination of the combined balance sheet of the LDI-pension sector suggests that this explanation falls short. The simple reason is that the UK corporate defined-benefit pension sector is substantially larger than the LDI sector. At the height of the fire sale, our analysis shows that debt (through repo and interest rate swaps) held on LDI balance sheets amounted to no more than 15-20% of the combined LDI-pension sector's balance sheet. Furthermore, this relatively low level of financial debt was supported primarily by sovereign debt, investment-grade corporate credit, and developed-market public equities, all of which trade in relatively deep and liquid markets.

Given that UK corporate pension schemes seemingly had ample assets to fully secure the debt on LDI balance sheets, we hypothesize that the internal contracting structure between pensions and LDIs effectively created a form of slow-moving capital (Mitchell et al., 2007; Duffie, 2010; Siriwardane, 2019). More specifically, collateral agreements between LDIs and their pension investors were structured such that their balance sheets were essentially segmented—LDI debts were

legally backed only by LDI assets and not the full set of their pension investors' assets. Consequently, as the demand for collateral increased with rising LDI leverage, pensions were asked to supply collateral on very short notice. However, institutional barriers, such as the need for formal pension scheme trustee approval, impeded timely transfers, forcing LDIs to instead proceed with liquidations.

As one test of this hypothesis, we examine the behavior of pooled LDIs, investment vehicles funded by multiple pension schemes. While recapitalization frictions can affect all types of LDIs, we focus on pooled LDIs because these frictions are easier to detect in this group for a few reasons. For one, the process of recapitalizing a pooled LDI fund is likely more challenging due to the need for coordination among multiple pensions, a task that becomes increasingly difficult during times of financial stress.² Furthermore, pooled LDIs face a collective action problem, as pensions may postpone transferring collateral in anticipation of contributions from others. In line with this idea, our analysis indicates that, compared to single LDIs, pooled LDIs sold roughly 13 percentage points more of their gilt holdings in the weeks following the mini-budget, even after accounting for differences in balance sheet composition, liquidity risk (Alfaro et al., 2023), and fund-manager fixed effects.

To isolate the fire sale discount resulting from pooled-LDI sales, we use similar approaches to those applied in assessing the price pressure caused by overall LDI sales. Our most conservative difference-in-differences estimator effectively compares two gilts with similar maturities within the same hour on the same day, both held in comparable amounts by LDIs before the crisis but differing in their exposure to pooled LDIs. At the height of the fire sale, we find that sales by pooled LDIs resulted in price discounts of 9.29%. While this analysis necessarily focuses on pooled LDIs for the purposes of identification, they support our broader argument that segmentation between LDI and pension balance sheets, coupled with delays in collateral transfers, is essential for understanding why the fire sale occurred in the first place.

This study connects to research on fire sales, the limits to arbitrage, and the boundaries of

²According to surveys of LDIs conducted by the Financial Conduct Authority, pooled LDI funds typically have more than 20 smaller pension schemes invested in them.

market efficiency.³ Fire sales have been documented in several assets, both real and financial, including aircrafts (Pulvino, 1998), housing (Campbell et al., 2011), concentrated stocks (Coval and Stafford, 2007; Bian et al., 2018), convertible bonds (Mitchell et al., 2007), and corporate bonds (Ellul et al., 2011; Falato et al., 2021). A closely related strand of research has studied law of one price deviations between financial derivatives and their underlying reference entities (e.g., Mitchell and Pulvino, 2012; Du et al., 2018; Siriwardane et al., 2023). Our contribution, made possible with detailed transaction-level data, is to tightly identify a meaningful and persistent fire sale in a large, liquid, and relatively safe asset class.⁴ At its peak, LDI selling resulted in price discounts on the order of 10% that did not fully close for over a month.

In addition, the granularity of the data allows us to pinpoint two forms of slow-moving capital that contributed to the gilt market crisis.⁵ The first operated through balance sheet segmentation, which was itself a function of collateral agreements and investment mandates between LDIs and their corporate pension investors. This segmentation prevented LDIs from recapitalizing quickly enough to avoid forced sales, despite the large asset base and improving solvency of corporate pension schemes. The second type involved strategic entry by arbitrageurs (Abreu and Brunnermeier, 2003). In essence, our empirical results demonstrate how the structure of debt contracts, internal agency conflicts, collective action problems, and strategic behavior all interact to impede the movement of arbitrage capital and detract from market efficiency, even within a transparent and easily accessible asset class like gilts.

This paper also complements previous research on the causes and consequences of the 2022 gilt market crisis, much of which has focused on the liquidity mismatch between the assets and liabilities of UK pensions (Chen and Kemp, 2023; Alfaro et al., 2023; Kodres, 2023; Dunne et al., 2023). On paper, the rise in gilt yields during September 2022 was actually favorable for pension

³See Shleifer and Vishny (2011) for a survey of research on fire sales. We also connect to the large literature on downward sloping demand curves in financial markets, such as Shleifer (1986), Greenwood and Vayanos (2010), Lou et al. (2013), and many others.

⁴Following Hu et al. (2013), Bicu-Lieb et al. (2020) estimate that the noise in the gilt market—a measure of liquidity based on deviations from a fitted yield curve—is about 4 basis points for the period 2010 between 2017. Hu et al. (2013) find similar levels of liquidity (~3 basis points) in the U.S. Treasury market during normal times.

⁵Grossman and Miller (1988), Gromb and Vayanos (2002), Duffie and Strulovici (2012), and Greenwood et al. (2018) are examples of theories featuring imperfect arbitrage and slow-moving capital.

solvency because the value of their liabilities—primarily pension obligations—fell more than their financial assets. Yet, in terms of liquidity, the LDI-pension sector became stressed because of the debt held on LDI balance sheets. Alfaro et al. (2023) build on this observation to propose a broader framework for understanding how hedging strategies that mitigate solvency risk can introduce liquidity risks, particularly in environments of low interest rates.

Our study underscores the pivotal role of the contracting environment in determining whether these types of liquidity mismatches ultimately lead to fire sales. In the context of the LDI crisis, our findings indicate that the liquidity pressures highlighted by Alfaro et al. (2023) may not have caused such large gilt sales had LDI debt been secured against the entire asset base of the combined LDI-sector. However, because LDI and pension-scheme balance sheets were effectively segmented, the severity of the fire sale hinged on the ability of pensions to quickly transfer collateral to LDIs. This issue particularly affected pooled LDIs, as evidenced by their more aggressive deleveraging compared to single LDIs. Moreover, balance sheet segmentation helps explain why the UK life insurance sector faced significantly less stress than the LDI-pension sector, despite a similar balance sheet composition at the outset of the crisis. This is because life insurers held their financial debt on their own balance sheets, unlike many pension schemes.

Finally, our results inform optimal policies during a fire sale, of which there are two main types (Shleifer and Vishny, 2011; Diamond and Rajan, 2009). The first involves offering loans or liquidity support to distressed sellers. However, during the gilt market crisis, the viability of this option was complicated by the substantial debt levels that were already hampering LDI balance sheets.⁶ The second response, direct asset purchases, presents its own dilemma regarding the intervention's duration. A simple strategy would be to implement asset purchases until prices align with their fundamental values, but this method is complicated by the difficulty of accurately determining when prices have normalized. Our analysis suggests that, in some cases, setting a definitive time-frame for asset purchases could be more advantageous. A predetermined intervention period not only creates incentives to resolve coordination issues that may be delaying equity injections from

⁶The complexities involved with offering short-term lending facilities to LDIs are discussed further in Section 5.

current asset owners—like those faced by LDIs, especially pooled ones—but also has the potential to draw in private-sector buyers. The Bank of England’s decision to establish clear start and end points for its intervention in the gilt market aligns with this reasoning.

The remainder of the paper is organized as follows. Section 2 provides background on LDIs and describes our main data sources. Section 3 studies how LDIs responded to financial distress induced by the mini-budget, identifies the impact of their selling on gilts prices, and provides evidence that buyers during the fire sale were either distressed or strategically delayed entry into the market. Section 4 explores why LDIs were unable to recapitalize quickly enough to avoid forced sales. Section 5 discusses the Bank’s policy response and concludes.

2 Background and Data

In this section, we provide background on LDIs, describe our data sources, and present summary statistics of our analysis sample.

2.1 Liability-Driven Investments

We start by briefly discussing LDIs and their role in the UK corporate defined-benefit (DB) pension system. Breeden (2022), Hauser (2022), and Chen and Kemp (2023) provide a more complete treatment. The concept of LDIs is closely connected to the literature on asset-liability matching (Campbell and Viceira, 2002; Brandt, 2010; van Binsbergen and Brandt, 2016) and can be understood through a simple example. Imagine a DB pension scheme that commits to providing £100 of real benefits 40 years from now. With an assumed annual inflation rate of 2% and a nominal riskless rate of 3%, the current value of the pension’s liability is £67.7 ($=100 \times [(1 + 2\%)/(1 + 3\%)]^{40}$). Should inflation rise to 2.5% without any change in nominal rates, the value of the liability would increase to £82.3. Similarly, an isolated decrease in the nominal riskless rate of 1% would increase the value of the liability to £148.3. Thus, the pension scheme is exposed to inflation and interest-rate risk through its liability structure.

Absent any funding constraints, the optimal asset allocation for the pension scheme is straightforward: it should purchase £67.7 of a 40-year, zero-coupon, inflation-linked gilt. This allocation completely safeguards the pension against any future fluctuations in interest rates and inflation, albeit at a high cost. Alternatively, a pension that is unable or willing to pay these hedging costs might opt to purchase long-dated gilts using financial leverage. However, as Alfaro et al. (2023) argue, the use of financial leverage potentially introduces a liquidity mismatch on the balance sheet of the pension scheme.

Historically, UK corporate DB pension schemes have executed these sorts of levered bond purchases indirectly by investing in so-called LDI funds (Chen and Kemp, 2023). In a typical LDI fund, a single pension or group of pensions invests capital, after which the LDI fund borrows additional capital to purchase gilts, either outright or using derivatives. The liability side of an LDI's balance sheet therefore consists of equity capital from the pensions and debt in the form of repurchase agreements and interest rate swaps (paying floating). The asset side consists of gilts, interest rate swaps (receiving fixed), and any cash equivalents. LDI funds may also use inflation swaps, though we ignored them in our illustrative example because they are a relatively small portion of the LDI balance sheet in practice. We present the balance sheet of the aggregate LDI sector in the UK below, after describing our main data sources.

Given their reliance on leverage, LDIs are vulnerable to forced sales when leverage exceeds sustainable levels. For example, if LDIs fail to meet variation margin payments, repo lenders or derivatives counterparties might contractually mandate liquidations. Similarly, internal risk management programs might compel LDIs to liquidate gilts if they deem leverage to be too high, termed “Deleveraging Events” in typical LDI prospectuses.⁷ During these events, LDI managers have some discretion over whether to raise additional equity through share issuance or to liquidate assets, with all associated costs borne by the fund's investors. In addition to internal risk management considerations, LDI prospectuses usually explicitly set minimum equity cushions, often between 5 and 10 percent of assets. In what follows, our discussion of forced sales by LDIs en-

⁷A representative LDI prospectus can be found [here](#).

compasses both types: those triggered by external parties like repo lenders and those driven by internal risk management protocols.

Throughout this paper, we assume that the only investors in our sample of LDIs are UK corporate DB pension schemes. In other words, LDIs are wholly owned by the UK corporate DB pension sector. This assumption stems from direct conversations with LDI managers and The Pensions Regulator (TPR), the principal regulatory authority for UK corporate pensions.

2.2 Data Sources

Classifying Trading Counterparties Our analysis uses several trade-level regulatory data sources in which the counterparties of a transaction are identified by their so-called Legal Entity Identifier (LEI). We map LEIs to different counterparty types (e.g., LDI funds, insurance companies, broker-dealers, etc.) using a combination of the Bank of England’s internal classification system, surveys of LDI funds, and manual searching of LEI codes. A single firm will often be associated with multiple LEIs. For example, a large asset manager may have LEIs associated with equity-focused funds, bond-focused funds, and LDIs. We link individual LEIs to parent fund managers (henceforth, fund managers) manually based on LEI names and public filings of investment funds and their managers. There are 238 LEIs associated with LDIs in our primary analysis sample and 18 unique LDI managers. Our LDI sample includes only single and pooled LDIs, the latter of which invests on behalf of multiple pension schemes. Segregated LDI mandates are not included because they are associated with the LEIs of the ultimate pension scheme investors.

Gilt Market Transactions We obtain data on gilt transactions from the MIFID II database, which is sourced by the UK Financial Conduct Authority. The MIFID II database consists of all transactions in the gilt market since 2017 and contains, among other things, the following information on each trade: (i) the buyer and seller, as identified by the LEI; (ii) the security exchanged, as identified by its International Securities Identification Number (ISIN); (iii) the transaction price;

(iv) the par value of the bond; and (v) the exact timestamp of the trade.⁸

The MIFID II database has been used previously by Czech and Pinter (2020), Pinter et al. (2024), Jurkatis (2024), and several others. We follow these studies and clean the MIFID II data as follows. We eliminate trade reports that are duplicates, lack client identifiers, or have negative reported prices. We identify and replace erroneous price entries with estimates derived from the total value of the transaction divided by the quantity of bonds traded. Our sample encompasses trades conducted both on a principal and agency basis. We supplement the MIFID II data using hourly price quotes at the ISIN level from the Thomson Reuters Eikon database.

Derivative Positions and Transactions We use the UK EMIR Trade Repository (EMIR TR) data for information on interest rate swap and inflation swap positions. The EMIR TR data come in two forms. The first is a daily “activity” file that shows transactions entered into or updated from the previous day. The second is a daily snapshot of all outstanding positions for each counterparty, including the notional amount, mark-to-market value, and maturity of each position. For each LEI, we use the latter to assemble the weekly stock of overnight index swaps (OIS) with floating legs linked to the Sterling Overnight Index Average (SONIA).

Repo Positions and Transactions We draw on the Sterling Money Market Data (SMMD) for information on the repo positions for each LEI. These data are maintained by the Bank of England and track virtually all sterling-linked money market transactions at the daily level. For each transaction, the dataset contains the borrower and lender in each repo transaction (identified by LEIs), the ISINs of the underlying collateral pool, repo rate, time stamp, maturity date and borrowed amount among others. The steps to clean the SMMD data follow our treatment of the MIFID II dataset: we focus on the client-dealer segment of the repo market and exclude contracts with missing client identifiers.⁹ We focus on repos collateralized by gilts, either nominal or inflation-linked,

⁸Further information on the MIFID II data can be found [here](#).

⁹The dealer-client segment is almost entirely settled on a bilateral basis and captures the majority of transaction volume. Unlike in the US, only a small fraction of the UK repo market is settled on tri-party basis. We exclude from our analysis interdealer transactions (that are typically settled through CCPs). See Van Horen and Kotidis (2018) and Gerba and Katsoulis (2021) for information on the UK repo market and other applications of the SMMD dataset.

and include repos with all maturity types (e.g. overnight, term and open repos). To approximate the stock of repos at a point in time, we accumulate the borrowed amounts for all outstanding contracts. The repo data is then merged with the gilt and derivatives data at the LEI-level.

Gilt Yields and Outstanding Amounts Daily zero-coupon bond yields on UK government bonds are obtained from the Bank of England. These data are based on Anderson and Sleath (2001) and include nominal and real yield curves, as well as the implied inflation term structure for the UK. Issuance and outstanding par values for each bond ISIN are taken directly from the website of the UK Debt Management office.¹⁰

2.3 Variable Definitions

Our goal is to construct the balance sheet of LDI funds through time. In what follows, we index LDI funds by their LEI i and denote $f(i)$ as manager of fund i . Time is indexed by t and is weekly because derivative positions are observed at the end of each week in our EMIR TR data. The subscript m is used to describe market values and the subscript q is used for par or notional values. Cash holdings do not appear on the balance sheet because we do not observe shares of money market funds and other cash-like assets.

The asset side of the balance sheet consists of all long gilt positions, including those created synthetically using interest rate swaps. For each LDI i , we compute the stock of gilt positions in the cash market by accumulating all purchases and sales in the MIFID II transaction data since 2017, the start of the data. While this approach introduces some error in estimating gilt positions, Section A.1.1 of the Internet Appendix demonstrates that our estimates align closely with surveys conducted by the Financial Conduct Authority (FCA) of the top LDI managers. We define $G_{i,t}^m$ and $G_{i,t}^q$ as the market and par values, respectively, of i 's gilt positions at time t . To reduce measurement error, we exclude 16 LDI accounts with a non-positive estimated gilt stock as of September 2, as this indicates an insufficient transaction history to accurately estimate their stock of gilts. These

¹⁰Specifically, the information on issuance is taken from the DMO's [issuance calendar](#), and information on gilts outstanding is taken from the DMO's [Gilts in Issue](#) section.

accounts are small, holding no swap positions and constituting just 1% of total outstanding repo borrowing by LDIs.

There are two main types of derivatives that appear as assets on LDI balance sheets, OIS swaps and inflation swaps. Throughout this paper, we treat the pay-floating leg of OIS swaps as a liability and the receive-fixed leg as an asset. This convention follows naturally from the fact that a swap in which an LDI receives fixed and pays floating (SONIA) is equivalent to borrowing at the floating rate and using the funds to purchase gilts. We define $L_{i,t}^m$ and $L_{i,t}^q$ as the respective market and notional values of the fixed leg of the swap, both of which are reported in EMIR TR. For inflation swaps, we adopt the convention that the receive-floating leg is an asset and the pay-fixed leg is a liability. $C_{i,t}^m$ and $C_{i,t}^q$ are defined as the respective market and notional values of the receiving leg of the swap.¹¹ The market value of assets for i at time t is therefore $A_{i,t}^m = G_{i,t}^m + L_{i,t}^m + C_{i,t}^m$. Similarly, par assets are defined as $A_{i,t}^q = G_{i,t}^q + L_{i,t}^q + C_{i,t}^q$.

The liability side of the balance sheet consists of borrowing via repurchase agreements, the pay-floating leg of OIS swaps, the pay-fixed leg of inflation swaps, and equity. We denote the amount of net borrowing via repo by $B_{i,t}^q$ and assume it is priced at par, meaning $B_{i,t}^m = B_{i,t}^q$. As discussed above, the amount of synthetic borrowing via OIS swaps equals the net notional amount of positions in which an LDI pays the floating leg, less any initial margin required to initiate the swap. We do not cleanly observe initial margin amounts in our EMIR TR data, but conversations with market participants suggests they were small or negligible during our sample. Thus, we set the amount of synthetic borrowing via OIS swaps equal to the notional value of the swap, denoted $S_{i,t}^q = L_{i,t}^q$.¹² We further assume that the market value of the floating leg in swaps equals its par value $S_{i,t}^m = S_{i,t}^q$, a natural assumption given the floating-rate nature of the synthetic debt. Similarly, the par value of the liability created by inflation swaps is set to $W_{i,t}^q = C_{i,t}^q$, the notional amount of the pay-fixed leg. We approximate the market value of the pay-fixed leg in an inflation swap

¹¹The EMIR TR reports the net market value of swaps from the perspective of each counterparty. The market value of the fixed leg of OIS swaps can therefore be inputted using the notional amount in the swap and the net market value. We approximate the market value of the receive-floating leg of inflation swaps as the notional amount of the swap plus the net market value, where net market values are defined such that they increase with realized inflation.

¹²For an initial margin h per dollar of notional, the amount of borrowing via swaps would then be $S_{i,t}^q = (1 - h)L_{i,t}^q$.

as $W_{i,t}^m = W_{i,t}^q$, which is accurate so long as the inflation swap rate is close to maturity-matched nominal interest rates. The market value of debt is therefore $D_{i,t}^m = B_{i,t}^m + S_{i,t}^m + W_{i,t}^m$ and equals the par value under our assumptions. This means that the equity of LDI i equals $E_{i,t}^m = A_{i,t}^m - D_{i,t}^m$.

Table 1 reports summary statistics on the 238 LDIs in our analysis sample. The purpose of the table is to show how the LDI sector’s balance sheet looked going into the crisis, so all of the summary statistics are as of September 2, 2022. In aggregate, the LDI sector had around £300 bn of assets entering the crisis, 88% of which was direct holdings of gilts in the cash market. The aggregate debt-to-asset ratio was 49%, implying an asset-to-equity ratio of 1.9. Around 72% of the aggregate debt was in the form of repo and the rest in OIS swaps. Figure 2a visualizes the aggregate balance sheet, with blue and green bars representing assets and orange and red bars representing debt. In addition to aligning with FCA surveys of LDIs, these estimates also largely agree with data collected by the Bank of Ireland on LDIs domiciled in Ireland (Dunne et al., 2023).

Table 1 also reveals that the size of LDIs was heavily skewed, with a few large players dominating the market. For example, the average LDI fund held £1.25 bn of assets whereas the median held only £0.32 bn. There was also considerable variation in the amount of leverage used by LDIs. The debt-to-asset ratio of the 25th fund was 0.26 compared to 0.89 for the 75th percentile fund. Some of the larger debt-to-asset ratios we observe are driven by measurement error in assets, particularly for LDIs with short transaction histories.

The last thing to note from the table is the distribution of pooled vs single LDI funds. Pooled LDIs make up 35% of our analysis sample by count, but only 19% by assets. This reflects the well-known fact that pooled LDI funds are generally used by smaller corporate pension schemes (Breedon, 2022; Hauser, 2022).

3 Price Pressure from LDI Selling

In this section, we proceed in three steps. First, we show how leverage within the LDI sector gradually rose throughout September 2022, spiking after the mini-budget announcement. It then

quickly returned to its initial level by early October, fueled mainly by asset sales in the gilt market and the closure of repo positions. Second, we measure the impact of LDI selling on gilt prices in several ways, estimating a fire sale discount of 6.87%. Finally, we identify potential liquidity providers during the fire sale and argue that their delayed entry was at least partly driven by a desire to time the bottom of the market.

3.1 LDI Behavior in Response to Financial Distress

Figure 2b depicts how the leverage—measured as the ratio of assets to equity—of the aggregate LDI sector evolved from the beginning of September through the end of October. The time-series of leverage in the plot is constructed in the same manner as in Table 1 and Figure 2a, with all positions are marked-to-market.

Figure 2b clearly shows the increased stress faced by the LDI sector through the month of September. Aggregate leverage steadily increased from 1.9 at the start of the month to just under 2.4 leading up to the mini-budget announcement. Following the announcement, leverage rose sharply, peaking at around 2.8 within three weeks, before returning to its original level by the end of October. These leverage ratios suggest that the equity capital of LDIs declined from 51% of assets to a low of 36% at the height of the crisis.

To develop a sense of the main factors driving LDI leverage, Figure 3 plots the cumulative flows of gilts, OIS and inflation swaps, and repo for the LDI sector from September to October. Flows in the plot are based on par values and indexed to value of zero on September 23, 2022, the day on which the mini-budget was announced. We use par values to isolate active balance sheet adjustments made by LDIs as opposed to passive changes arising from changes in gilt prices. Interestingly, in the weeks leading up to the mini-budget, LDIs purchased £5.7 bn of gilts in par terms and simultaneously raised their repo borrowing by £17.0 bn. Thus, the increase in LDI leverage during the first few weeks of September seen in Figure 2b was driven by an increase in repo borrowing and a fall in gilt prices (see Figure 1).

Figure 3 further shows that LDI selling began precisely after the announcement of the mini-

budget on September 23. Over the subsequent five weeks, LDIs sold roughly £25.3 bn worth of gilts, 30% of which occurred in the five day window between the mini-budget and the Bank's first intervention in the gilt market on September 28. In Figure A1 of the Internet Appendix, we further show that the velocity and size of these sales was abnormal relative to the historical trading behavior of LDIs.

The red line in Figure 3 suggests that LDIs used the proceeds from selling gilts to pay down debts, namely their repo borrowing. From the onset of the mini-budget to the end of October, LDIs reduced repo borrowing by £32.6 bn. The magnitude of this decline was large. At the time of the mini-budget, LDIs had £120.5 bn of repo outstanding, implying that they exited 27% of their outstanding repo positions between September 23 and October 31.

Despite the clear efforts by LDIs to delever, Figure 3 indicates that their leverage nonetheless continued to rise for almost three weeks following the mini-budget. The reason why is that gilt prices were falling rapidly during this period. For example, Figure 1 shows that 30-year gilts rose by almost 100 basis points in the 5 days after mini-budget, driving down the value of existing LDI assets and pushing up LDI leverage.

Two clear takeaways emerge from Figures 2b and 3. First, the mini-budget announcement seems to have been a catalyst for LDI selling, a point we build on in later sections when estimating the effect of LDI selling on gilt prices. Second, the rapid deleveraging by LDIs that occurred in the weeks after the mini-budget resulted from a mix of gilt sales, reductions in repo borrowing, and a rebound in gilt prices. Overall, our analysis suggests that 73% of the debt reduction by LDIs from September 23 to October 31 was financed by gilt sales, with the remainder coming from equity injections or existing cash reserves.¹³ The lack of equity injections and its implications are discussed in detail in Section 4.

¹³These estimates suggest that LDIs raised £8.8 bn of equity during this period. This is very likely a lower bound because we are not able to infer equity injections that are invested in cash or cash-equivalents, as we do not observe the latter. See Dunne et al. (2023) for estimates of equity injections between September and October for Irish-domiciled LDIs.

3.2 Price Pressure Estimates

3.2.1 Conceptual Framework

We now estimate the effect of LDI selling on gilt prices. To understand the identification problem and our proposed solution to it, suppose that the return of gilt i is given by:

$$r_{i,t} = a_i U_t + b_i F_t \quad (1)$$

where U_t is a fundamental factor and F_t is a fire-sale factor that is driven by LDI selling. F_t can be thought of as being negative when LDI selling pressure puts downward pressure on gilt prices. a_i captures gilt i 's exposure to fundamental shocks, such as the mini-budget, and could stand in duration or other gilt characteristics. The loading b_i captures gilt i 's exposure to LDI selling. It could be positive because LDIs directly sell i or because LDI selling of other gilts has spillover effects on i 's price, as in Vayanos and Vila (2021).

The central identification challenge is to isolate the component of gilt i 's return that is driven by the fire sale, $b_i F_t$. It is clear from the reduced-form model in (1) that this is not possible from i 's return alone. However, suppose it is possible to find two gilts, i and j , that: (i) have the same exposure $a_i = a_j$ to the fundamental shock U_t ; and (ii) have differential exposure to the fire-sale factor F_t , such that $b_i > b_j \geq 0$. In this case, the return difference between the two gilts equals:

$$r_{i,t} - r_{j,t} = (b_i - b_j) F_t. \quad (2)$$

When $b_j = 0$, the return differential between the two bonds fully recovers the component of gilt i 's return attributable to LDI selling, $b_i F_t$. Even if $b_j > 0$, the return differential still provides a lower bound on the total effect of the fire sale on gilt i . In other words, $b_i F_t \geq r_{i,t} - r_{j,t}$, so long as $b_i > b_j \geq 0$.

3.2.2 Difference-in-Difference Estimates

Equation (2) motivates our identification strategy, which involves comparing gilts with similar maturities but varying degrees of exposure to LDI selling. The basic idea is that gilts with similar maturities have comparable exposure to fundamental shocks (a_i 's), most notably the mini-budget. Within the set of gilts that have a similar maturity, we proxy for high- b_i gilts based on those that were more heavily held by LDIs entering September.

We implement this identification strategy using the following difference-in-difference regression model:

$$\log(p_{b,h,t}) = \gamma_b + \alpha_{h,t,m} + \sum_w \beta_w \times 1(t \in w) \times D_b + \varepsilon_{b,t}, \quad (3)$$

where $p_{b,h,t}$ is the median transaction price of gilt b in hour h of day t . To construct the panel of transaction prices $p_{b,h,t}$, for the regression, we first assign all transactions that occur between 5 pm on date t and 8 am on date $t + 1$ to the hour between 7 and 8 am GMT on $t + 1$.¹⁴ We then compute the median transaction price $p_{b,h,t}$ of bond b in hour h on date t . This means median transaction prices are available for hours $h \in [7, 8, \dots, 16]$ of each day.

γ_b in regression (3) is a fixed effect for bond b . $\alpha_{h,t,m}$ is a maturity-by-hour-by-date fixed effect and is constructed by sorting bonds into m bins based on maturity, then interacting fixed effects date and hour with indicator variables based on maturity-bin assignment. We set $m = 5$ for our baseline analysis. $1(t = s)$ is an indicator variable for whether date t is in week w .

D_b in the regression is an indicator variable for whether bond b is heavily held by LDIs. To construct it, we first compute the fraction H_b of bond b held by LDIs as of September 1, 2022. Data on the outstanding amount of each bond comes from the historical issuance reports that are produced by the UK Debt Management Office (DMO). Because these data are in par values, we use the par value of gilt b that is held by LDIs when computing H_b .¹⁵ D_b is then defined based on whether H_b is in the top tercile.

¹⁴Standard trading hours in the gilt market are from 8:00 AM to 4:30 PM.

¹⁵Figure A2a of the Internet Appendix confirms that H_b varies meaningfully across bonds. For example, as of the beginning of September, LDIs held 2.8% of the bonds in the lowest tercile of H_b and 43.3% in the highest tercile.

The key focus in regression (3) are the coefficients $\{\beta_w\}$. Because of the maturity-by-hour-by-date fixed effect $\alpha_{h,t,m}$, their identification comes from effectively comparing the prices of two gilts with similar maturities, traded in the same hour of the same day, differing only in their pre-crisis exposure to LDIs. The inclusion of the fixed effect $\alpha_{h,t,m}$ is critical to our empirical strategy because LDIs typically led longer-duration gilts (Internet Appendix Figure A2b), which are more sensitive to any news affecting the level of interest rates. Unobserved bond-level traits, such as issue size or offering date, are accounted for by the bond fixed effect γ_b . As with any difference-in-difference estimator, the β_w 's trace out the causal impact of LDI selling on gilt prices under the assumption of parallel trends: absent LDI selling, bonds in the same maturity bin would have followed similar price trends from early September through the end of October.

Column (1) of Table 2 contains estimates of regression (3) using data from the first full week of September through the end of October. The estimated coefficients in the table are normalized such that $\beta = 0$ for the week in which the mini-budget was announced (2022w38). They are plotted in Figure 4a along with 95% confidence intervals based on standard errors that are double-clustered by bond and date-hour. The plot shows that prior to the mini-budget, bonds held heavily by LDIs followed a similar price path to those that were not. The coefficients prior to the mini-budget are all close to zero and statistically insignificant, and we cannot reject the null hypothesis that they are equal ($p = 0.83$). The fact that prices for the two gilt groups follow similar trends prior to the mini-budget supports our assumption of parallel post-trends.

During the first two weeks after the mini-budget, the average fire sale discount widens to 3.3% and is statistically significant at conventional confidence levels. It rises to 6.3% ($t = -4.76$) in the third week before steadily declining. By the end of the month, there is no statistically detectable difference between the price of bonds with low- and high-LDI exposure. This pattern of decay reinforces the idea that the estimated price gaps in first weeks after the announcement were driven by forced selling as opposed to information revelation. For instance, one might worry that LDIs had private information about the fundamental value of the bonds they sold, perhaps because they better understood the implications of the mini-budget for specific gilt issues. In this case, their

selling simply transmitted information into bond prices and should not be interpreted as causing a fire sale. However, as argued by Coval and Stafford (2007), this channel would also predict that the price gap between bonds with high and low LDI exposure should remain permanent, in contrast to Figure 4a. Moreover, recall that H_b (and hence D_b) is defined using LDI holdings at the beginning of September, well before the onset of the crisis.

While the lack of pre-trends in Figure 4a is consistent with parallel post-trends, it does not ensure its validity. In recent work, Rambachan and Roth (2023) develop methods for assessing how β_w estimates and their confidence intervals would adjust in the face of various violations of the parallel trends assumption. We implement their procedure in Figure 4b, focusing on how the estimated β_w in the third week following the mini-budget would shift if the trend low and high-LDI exposure bonds had continued linearly from before announcement. Reassuringly, even under this violation, the estimated peak weekly price impact is largely unchanged and, if anything, is slightly larger. This robustness simply reflects the parallel pre-trends embedded in Figure 4a.

The fixed effect $\alpha_{h,t,m}$ is central to our empirical strategy because it ensures that price pressure effects are identified by comparing bonds with similar duration. In column (1) of Table 2, $\alpha_{h,t,m}$ is based on sorting bonds into quintiles based on their maturity. However, it is possible that bonds in the same maturity-quintile may still differ meaningfully in duration. Given LDIs generally held longer-duration bonds, any such variation might over-inflate the estimated size of the fire sale discount. To alleviate this concern, columns (2) and (3) show estimates of regression (3) when $\alpha_{h,t,m}$ is defined by assigning bonds into more granular maturity bins. The standard errors naturally widen since the identifying variation comes a narrower set of bonds. Nonetheless, the point estimates are largely unchanged and the peak fire sale discount of around 5% is statistically significant when dividing bonds into ten or fifteen maturity bins. In Internet Appendix A.1.4, we also find similar estimates when using hourly price quotes instead for transaction prices.

Regression (3) estimates the average fire sale discount over the course of each week. However, Figure 1 shows there is substantial price variation within each week, particularly the one in which the Bank first intervened. To better capture these higher-frequency dynamics, we reestimate re-

gression (3) by allowing the β 's to vary daily, as opposed to weekly. Figure 4c plots the resulting coefficient estimates and their associated 95% confidence bands, again based on standard errors that are clustered by bond and date-hour.¹⁶

The first thing to notice from the plot is that the sharp drop in the relative prices of high-exposure LDI bonds begins after the announcement of the mini-budget on September 23, precisely when LDIs began selling gilts (Figure 3). While it might seem tempting to ascribe this decline to the announcement's revelation of bond fundamentals, recall that the coefficients are identified by comparing bonds of similar durations—those presumably equally affected by the mini-budget. Three business days later, on September 27, bonds heavily exposed to LDIs traded at discounts of 6.87% ($t = -3.53$). The fire sale discount disappears when the Bank first intervened in the gilt market, only to reopen in the following days. By October 11 it had widened to 7.77% ($t = -6.24$), after which it steadily declined until the end of the month.

Figure 4c further shows that there are no visual pre-trends prior to the mini-budget announcement. The magnitudes of the pre-announcement coefficients are generally small and we cannot reject the null hypothesis that they are all equal at a 95% confidence level ($p = 0.07$). To further probe the sensitivity of results to violations of parallel trends, we again follow Rambachan and Roth (2023) and assess how our estimate of the fire sale discount on September 27 would change if the price trend between low and high-LDI exposure bonds had continued linearly through the mini-budget announcement. Figure 4d shows that the estimated discount on the day before the Bank's intervention is largely unchanged under this scenario.

Given richness of the fixed effects in regression (3), the fire sale discount of 6.87% in Figure 4c should be thought of as a lower bound on the true effect of LDI selling on gilt prices. To illustrate why, let us return to the conceptual framework laid out at the beginning of the subsection. In particular, suppose there are three gilts: i , j , and k . Gilt k has no exposure to LDI selling, gilt j has medium exposure, and gilt i has high exposure (i.e., $b_i > b_j > b_k = 0$). Let m_x be the maturity of gilt x and assume that the maturity structure is $m_k < m_j = m_i$. Further assume that all three gilts are

¹⁶In Figure A4 of the Internet Appendix, we show that the difference-in-difference estimates are similar when we interact the date-by-hour-by-maturity fixed effect $\alpha_{h,t,m}$ with an indicator for whether gilt b is a nominal instrument.

equally exposed to a fundamental shock that hits the UK economy ($a_i = a_j = a_k$) and that after the shock, the returns of the gilts are $r_k = -10\%$, $r_j = -20\%$, and $r_i = -25\%$. In this case, it would be reasonable to conclude that LDIs caused a fire sale discount of 15% [= $r_k - r_i$]. However, because our identification strategy relies on comparing bonds with a similar maturity, it would suggest a discount of 5% [= $r_j - r_i$]. Moreover, if there are spillover effects from LDI selling onto gilt k ($b_k > 0$), as is often the case in models of preferred habitat investors (Vayanos and Vila, 2021), the estimated impact of LDI selling on prices would be even more understated. Thus, while our difference-in-differences approach controls for any potential heterogeneity in shock exposure due to bond maturity, it does so at the cost of conservatively estimating the effect of LDI selling.

3.2.3 Portfolio Sorts

Given the conservative nature of our difference-in-difference estimates, we now use simple portfolio sorts to develop a sense of the total potential impact of LDI selling on gilt prices. In particular, we sort gilts into terciles based on their value of H_b , resulting in about 30 bonds per portfolio. Figure 5a then plots the cumulative return of the low- and high-LDI exposure portfolios through time. Cumulative returns at the bond level are computed using end-of-day price quotes from the Thompson Reuters Eikon database and are relative to closing prices on September 22, the day before the mini-budget. Portfolio returns are calculated using an equal-weighted average of the returns from each bond in the portfolio.

Figure 5a shows a sharp decline in the prices of high-LDI exposure bonds relative to low ones following the announcement of the mini-budget. Within a matter of days, high-LDI exposure gilts had fallen 23.0% from their level just prior to the mini-budget, whereas low-LDI exposure gilts experienced a decline of 6.3%. The performance difference between the two portfolios implies that LDI selling caused a fire sale discount of 16.7%. This figure is likely an upper bound on the true fire sale impact because gilts held by LDIs were almost surely more exposed to the fundamental shock of the mini-budget by virtue of their longer duration. These duration differences also explain why the implied fire sale discount of 16.7% in Figure 5a is somewhat larger than our more conservative

difference-in-difference estimate of 6.87%.

3.3 Where were the natural buyers?

Given the large fire sale discounts documented in Section 3.2, a natural next question is where were the natural buyers? We now offer a partial explanation, with a focus on the week following the mini-budget announcement. Before moving forward, it's important to emphasize the complexities involved in addressing this question. Providing a comprehensive answer necessitates understanding the motives and constraints of all potential buyers in the market, information that is notoriously challenging to obtain in practice. Therefore, the evidence we present below should be viewed as indicative and not conclusive.

With this caveat in mind, we start by investigating the types of institutions that could in principle have provided liquidity to the gilt market when LDIs began selling after the announcement of the mini-budget. To do so, we compute the par value of gilts, net swap notional (receiving fixed), and the sum of the two for the following institutions: (i) hedge funds; (ii) dealers; (iii) insurance companies; and (iv) asset managers (non-LDIs). With the exception of dealers and insurance companies, the stock of gilt positions for each institution type are computed as in Section 2.3, by cumulating gilt flows in the MIFID II dataset. For dealers, we use Form BT (item 32D) collected by the Bank of England. Form BT arguably provides a more accurate estimate of their gilt stock because dealers participate in primary auctions with the DMO and MIFID excludes such transactions. The cost of using Form BT is that it may include gilt stocks of the banking arms of dealers and it reports gilt stocks in market, not par values. For insurance companies, we use Bank of England regulatory data covering life and non-life insurance companies' asset holdings. For all institution types, the stock of derivative positions in the EMIR TR.

Table 3 reports the notional value of gilt positions at the beginning of September, several weeks before the mini-budget. A few key patterns emerge from the table. Notably, hedge funds were very well-positioned to profit from the rise in gilt yields that happened throughout September: they held short positions of roughly £200 bn in gilts, driven mostly by swap positions. Although

dealers also had a relatively small net short position in gilts, the amount of their shorting is likely overstated because their holdings are valued at market prices in Form BT and gilt prices had been declining throughout 2022. Additionally, when including swap positions from banks that own dealers, the combined dealer-bank sector was net long about £100 bn gilts. Perhaps unsurprisingly, insurance companies and asset managers were positioned in the same long-duration direction as LDIs, rendering them less likely to provide liquidity when LDIs started selling. This is especially true given many gilt investors have preferences for certain types of bonds (Giese et al., 2024), as in the preferred habitat model of Vayanos and Vila (2021).

Given that the hedge fund sector was very likely unconstrained, why was there not more price competition when LDIs were forced to sell in the days after the mini-budget? One hypothesis stems from the model of Abreu and Brunnermeier (2003), who study the conditions under which a bubble in asset prices can occur even in the presence of well-capitalized and rational arbitrageurs. The main insight of their model is that bubbles can persist in equilibrium because an arbitrageur may not trade against an overvalued asset if it thinks the bubble will expand further, thereby increasing its total profits. This belief can be rational if other arbitrageurs are not yet aware the asset is overvalued. Thus, the lack of certainty about when selling pressure will reach a critical mass creates rational incentives for arbitrageurs to remain dormant.

It is easy to imagine many of the same forces governing fire sales. For example, an unconstrained arbitrageur who is aware of a fire sale may still refrain from deploying capital if it believes the fire sale is going to deepen further. The optimal time for such an arbitrageur to enter is just before the fire sale ends, which in Abreu and Brunnermeier (2003) amounts to timing the point at which other arbitrageurs become sufficiently informed about the fire sale.

It is perhaps harder, however, to imagine that a timing mechanism of this sort was operant during September 2022. After all, the gilt-market crisis was very public, making it doubtful that arbitrageurs were differentially informed about the existence of a fire sale. More importantly, specialized knowledge was not needed to recognize that large price wedges between gilts of similar maturities represented a profitable trading opportunity, even amidst uncertainty about the fiscal

impact of the mini-budget. A more plausible scenario is that arbitrageurs delayed entry because they were trying to time when, if ever, the Bank of England might intervene by purchasing gilts directly in the open market.

To be clear, any such timing strategy does not imply that arbitrageurs possessed private information about the Bank’s policy plans. The volatility in the gilt market was highly publicized and market participants may have thought that the Bank would intervene if it turned into a fire sale, given its financial stability objectives. As early as September 26, Governor Andrew Bailey issued a public statement stating that the “Bank [was] monitoring developments in financial markets very closely in light of the significant repricing of financial assets.” Moreover, according to an October 2022 letter from Deputy Governor Sir Jon Cunliffe to Mel Stride, Chairman of the Treasury Committee, the Bank received several pieces of market intelligence as early as September 23 that LDIs would be forced to liquidate very large quantities of gilts if market conditions worsened.¹⁷ It seems reasonable to think that at least some hedge funds were privy to some or all of the same market intelligence, perhaps through their dealer networks, allowing them to time their purchases when prices troughed.

To test this hypothesis, we study high-frequency trading behavior in the hours leading up to the Bank’s first intervention on September 28. At 11 am on the 28th, the Bank announced that it would “carry out temporary purchases of long-dated UK government bonds,” with the purpose of “restor[ing] orderly market conditions.” They further announced that its purchase program would commence at 3 pm later that day and run through October 14.¹⁸ With these event-times in mind, Figure 6a shows hourly cumulative flows in the gilt market starting on the morning of September 27 for hedge funds, dealers, and LDIs. Time on the x -axis of the plot is expressed in the number of trading hours relative to 11 am UK time, when the Bank announced its purchase program. To make the plot more readable, we collapse trading in each day to so that all flows occur between the hours of 7 am and 5 pm, assigning overnight trades to the 7-8 am hourly window. This choice

¹⁷Governor Bailey’s announcement on September 26th can be found [here](#). Depute Governor Sir Jon Cunliffe’s letter on October 5 can be found [here](#).

¹⁸The formal announcement from the Bank can be found [here](#).

reflects the fact that standard trading hours in the gilt market are between 8 am and 4:30 pm.

The plot shows that LDIs engaged in significant selling of gilts during the thirty-hour period leading up to the Bank's intervention. On September 27 alone, LDIs offloaded £4.64 bn worth of gilts, which represented about 20% of their total sales from the time of the mini-budget announcement through the end of October (see Figure 3). Dealers absorbed around half of these sales, purchasing £2.47 bn of gilts, with hedge funds absorbing the rest.

Though hedge funds had been active at absorbing some of the LDI gilt sales on the 26th and 27th, the pace of their liquidity provision surged early on September 28, with £3.76 bn acquired in the hours leading up to the Bank's intervention announcement. To put this into context, in the two weeks after the mini-budget, hedge funds acquired a total of £8.60 bn in gilts, meaning 44% of their purchases during this period were made in the small window just before the Bank's announcement. Moreover, Figure 6b reveals that the bulk of the sector's gilt purchases on the morning of the 28th were made by a limited number of hedge funds. To prepare this plot, we first calculate gilt flows for all hedge funds within a given window. The figure then plots the proportion of total gilt purchases attributed to the top k hedge fund buyers, among those that engaged in buying. On the morning of September 28, the top five hedge funds were responsible for 66% of the sector's total purchases. By comparison, the top five hedge funds made up 43% of purchases in the week following the mini-budget (excluding the 28th morning) and 47% of hedge fund purchases on September 27. Together, the patterns in Figure 6 support the idea that hedge funds strategically provided liquidity to the market because they were attempting to time the bottom of the fire sale.

A different interpretation is that hedge fund purchases simply mirror LDI sales and do not reflect strategic timing. However, Figure 6 illustrates that hedge funds predominantly bought gilts from dealers on September 28, who reduced their gilt inventories by £2.21 bn prior to the Bank's announcement. The significant volume of purchases from dealers during this window is consistent with hedge funds proactively initiating these transactions. The fact that the pace of LDI selling remained relatively stable on the morning of the 28th, whereas hedge fund buying increased rapidly, further supports this interpretation. Overall, the behavior of hedge funds during this period aligns

with the wait-and-see strategy described by Abreu and Brunnermeier (2003). Figure 6c, which shows the average hourly price of long-dated gilts, suggests such a strategy was highly profitable: hedge funds who purchased gilts on the morning of the 28th would have earned about a 40% unlevered return if they then sold them at the end of the trading day.¹⁹

Stepping back, it's important to emphasize that the evidence from Figure 6 only hints at the reasons why liquidity provision failed to fully counteract the initial price pressures from LDI sales. While we've outlined strategic considerations, there are many other frictions that could have deterred buyers from swiftly entering the market. For instance, even well-capitalized hedge funds might delay providing liquidity out of fear that the fire sale could intensify, especially if they use leverage or face risk limits (Shleifer and Vishny, 1997). The key takeaway, regardless of the specific frictions, is that these barriers were significant enough to contribute to the substantial fire sale discounts documented in Section 3.2.

4 The Mechanism: Why Couldn't LDIs Recapitalize?

The preceding section detailed the rapid gilt sales by LDIs under financial distress, which resulted in significant fire sale discounts. This section aims to understand why LDIs were forced to sell in the first place. We begin by analyzing the consolidated balance sheets of LDIs and UK pensions, demonstrating that their combined leverage was not particularly high at the onset of the crisis, nor did it significantly increase at its peak. This leads us to hypothesize that LDIs were forced to sell because their debt was not collateralized by the full spectrum of UK pension assets. Consequently, LDIs required large collateral transfers from their pension investors when their leverage spiked, yet institutional constraints prevented these transfers from happening swiftly enough to prevent forced sales. We explore the implications of this hypothesis about internal slow-moving capital through the lens of pooled LDI funds, where such frictions are expected to be most pronounced.

¹⁹Long-dated gilts are defined as those whose maturity is at least 20 years. The threshold of 20 years was chosen because this was the original minimum maturity gilt that the Bank offered to buy.

4.1 The Consolidated LDI-Pension Balance Sheet

In Section 3.1 we examined how the balance sheet of LDIs evolved during September 2022. To better understand why LDIs delevered through asset sales and not equity injections, we now shift our attention to the combined balance sheet of LDIs and their primary investor, the UK corporate sector. Going forward, we refer to this combined or consolidated balance sheet as belonging to the LDI-Pension sector.

Our analysis relies on data from the UK Pension Protection Fund (PPF), a statutory entity that compensates members of qualifying corporate DB pension schemes in cases of employer insolvency. Each year, the PPF releases the “Purple Book,” which provides a comprehensive snapshot of the solvency, demographic composition, and portfolio composition of the UK corporate pension sector. The Purple Book draws on supervisory data reported to TPR, the primary regulator of UK pension schemes. We supplement these annual reports with monthly data on pension assets and obligations from the January 2024 edition of the PPF 7800 index.²⁰

Importantly, data in the Purple Book and PPF 7800 Index reflect the consolidated balance sheet of the LDI and pension sectors. This conclusion follows from guidelines provided by the TPR to its reporting pension-scheme members. For example, on their website, the TPR considers a pension that allocates 10% of its assets to an LDI fund focusing solely on nominal gilts. According to the TPR, if this LDI fund employs leverage to purchase £1.80 in gilts for every £1 of equity invested, then the pension should classify 18% of its portfolio as invested in nominal gilts and -8% as cash.²¹ Cash positions in the Purple Book therefore reflect any LDI leverage.

Figure 7 summarizes this balance sheet at the end of August 2022. The value of assets and pension obligations are based on the August 2022 entry of the PPF 7800 index. The composition of assets is based on allocations in the 2023 release of the Purple Book. We use the 2023 edition because asset allocations in the Purple Book are generally reported with a lag. For example, according to the 2023 Purple Book (p21), 75.9% of its asset allocations are dated between April 1,

²⁰The Purple Book and PPF 7800 Index can be found, respectively, [here](#) and [here](#).

²¹This example is taken directly from the TPR’s [website](#).

2021 and March 31, 2022. The bulk of the remaining allocations (21.2% of total reporting plans) are dated between April 1, 2022 and September 30, 2022. This means that our asset allocations may be slightly biased by any portfolio rebalancing that occurred during the month of September 2022.²²

Figure 7 shows that, at the end of August 2022, the LDI-Pension sector held £1581 bn of total financial assets. This sum included £661 bn of government debt, £377 bn of corporate debt (fixed rate), and £191 bn of public equities.²³ The remaining £354 bn of LDI-Pension assets were invested in what we label as alternatives, consisting mainly of private equity, hedge funds, property, and annuities.

Unsurprisingly, the largest liability of the LDI-Pension sector is in the form of pension obligations, which were valued at £1184 bn under s179 valuation standards.²⁴ Financial debt, again including any repo and swap borrowing by LDIs, comprised a much smaller portion of liabilities (£77 bn). Together, these numbers mean that U.K corporate pensions entered September with a funding surplus of £320 bn. The surplus primarily resulted from the increase in gilt yields over 2022, combined with the longer duration of pension obligations compared to that of their assets.

Figure 7 shows no cash assets because financial leverage is net of any cash. Table 1 provides a way to approximate the combined cash position of LDIs and pensions. The table shows that LDIs entered September with about £145 bn of debt in the form of repo borrowing, pay-floating swaps, and pay-fixed inflation swaps, implying a cash position of around £68 bn for the combined LDI-Pension sector.

The most striking insight from Figure 7 is the LDI-Pension sector's relatively low level of financial leverage prior to the gilt-market crisis. At the beginning of September 2022, net financial debt constituted just 5% of total LDI-pension assets. Another way to think about the amount of leverage

²²This bias is likely to be quite small. To see why, assume that plans are equally reporting between April 1, 2022 and September 30, 2022 are equally likely to report in a given month. In this case, about 3.5% (=21.2%/6) of plans would have reported in September 2022.

²³Our asset total for August 2022 is larger than the £1503.9 reported in the PPF 7800 Index because we consider negative cash holdings as a liability.

²⁴According to the TPR valuation [guidelines](#), liabilities should be discounted using the yield curve high-quality corporate bonds, usually AA-rated or higher.

is to focus solely on the £145 bn of debt on LDI balance sheets. Although this debt amounted to approximately 49% of LDI assets (Table 1), it represented only 9% of the aggregate assets of the combined LDI-Pension sector. This disparity underscores the considerable size difference between the UK pension sector's balance sheet and that of the LDI sector.

The use of LDI-Pension leverage is also not particularly excessive when considering the composition of the sector's assets. Table 4 shows that £257 of the £377 bn of corporate debt held by LDI-Pensions was rated investment-grade, with the remainder classified as high-yield or private credit. In terms of public equities, £21 bn was allocated to UK stocks, £149 bn to developed markets, and £21 bn to emerging markets. These figures mean that 69% of LDI-Pension sector capital was invested in sovereign debt, developed-market public equities, or investment-grade credit. From the viewpoint of a portfolio or hedge fund manager, employing 5-10% leverage on such a portfolio would not be considered large by any reasonable market standard, even during times of stress.

One limitation of this analysis is that it relies on balance sheet data from the end of August 2022, thus ignoring the decline in gilt prices leading up to the announcement of the mini-budget on September 23. To address this concern, we use a back-of-the-envelope calculation to assess the impact of falling gilt prices on LDI-Pension assets. Specifically, we regress changes in LDI-Pension assets on changes in 30-year gilt yields over the 2-year period between August 2020 through August 2022. This regression implies a total dollar-duration of about £160 bn. With 30-year gilt yields rising about 60 basis points from September 1 to September 2022, this duration suggests LDI-Pension assets would have decreased by around £96 bn during this timeframe. Under such conditions, LDI borrowing would still account for less than 10% of LDI-Pension assets. Even with a 180 basis point surge in 30-year yields – a scenario mirroring the crisis's zenith – LDI borrowing would not exceed 15% of LDI-Pension assets.

Discussion In analyzing the causes of the gilt crisis, several observers have highlighted the liquidity mismatch inherent in the balance sheets of LDI-pension schemes (e.g., Alfaro et al., 2023). This mismatch became particularly evident when gilt yields spiked in September 2022, leading to

a decrease in the value of pension obligations and pushing the funding ratio of UK pensions to a twenty-year high (Figure A5). Despite the reduction in total liabilities, the liquidity needs of the sector actually rose because of the margin calls triggered by LDI leverage. Although this liquidity mismatch played a clear role during the crisis, both Figure 7 and Table 4 imply that, in theory, such a mismatch should not have been so problematic. The underlying reason is straightforward: UK corporate pensions possessed sufficient high-quality assets to fully collateralize the debt on their LDI investments' balance sheets.

This simple observation immediately suggests that LDIs were forced to sell assets due to contracting frictions with their pension investors. To understand the nature of the contracting friction that we have in mind, consider a pension with £100 of capital that wishes to purchase £150 of gilts using margin debt. The simplest arrangement for the pension would be to hold the margin debt on its own balance sheet, as depicted in Portfolio 1 of Figure 8. A second, more complex arrangement would be for the pension to invest £30 of capital into an LDI fund that purchases the £150 using £120 of margin debt. This situation is depicted as Portfolio 2 in Figure 8.

On a consolidated basis, Portfolio 1 and Portfolio 2 might look the same, but their reaction to market downturns could vary significantly. Should there be a 20% decline in gilt prices, Portfolio 1 would remain stable; its £50 of margin debt would still be far exceeded by the £120 value of its gilts. Conversely, the same drop would push the LDI in Portfolio 2 into technical default, since its margin debt is collateralized only by its own assets, not those of its pension owner. Theoretically, the pension could transfer £70 of its cash reserves to bolster the LDI's balance sheet. Yet, institutional barriers, such as the need for board approval, might impede timely collateral replenishment, forcing the LDI to instead sell its gilt holdings.

The preceding example illustrates how the siloed or segmented nature of the LDI's balance sheet, coupled with procedural delays in transferring additional collateral, can create a form of slow-moving capital and lead to forced sales (Duffie, 2010; Shleifer and Vishny, 2011)). We will now provide two pieces of evidence suggesting that these dynamics are essential for understanding why LDIs were forced to sell after the mini-budget, despite the substantial asset base of the LDI-

Pension sector relative to its financial debt (Figure 7). First, UK life insurers, who entered the crisis with a balance sheet composition similar to that of the LDI-Pension sector, did not experience nearly as much stress. We attribute this difference to life insurers maintaining their financial debt directly on their balance sheets, in contrast to the inherent segmentation between LDIs and pension schemes. Second, pooled LDIs, which face large frictions to collateral transfers due to coordination issues, sold gilts and reduced repo borrowing much more than single LDIs with similar balance sheets.

4.2 UK Life Insurers

In Internet Appendix A.2, we offer a detailed comparison of the balance sheets of life insurers and the LDI-Pension sector. Our analysis, which uses regulatory data available to the Bank, reveals several key findings. First, the balance sheet composition of life insurers and LDI-Pensions were comparable on the eve of the crisis, which is not surprising given insurers serve a similar economic function to DB pensions. For example, both held around 60% of their assets in fixed-income securities, though life insurers held more corporate bonds (30%) and pension schemes held more sovereign debt (42%). Additionally, life insurers held long-duration assets, with 60% of their bonds having maturities of over 10 years and more than 30% exceeding 30 years. Consequently, the assets of UK life insurers were significantly exposed to interest rate risk at the onset of the gilt-market crisis. As with the pension-LDI sector, the net financial debt of the insurance sector was quite small (<10%) relative to their total assets.

Second, solvency measures for the insurance sector also rose with gilt yields. Specifically, Figure A10a in the Internet Appendix illustrates that the capital coverage ratio for the insurance sector closely mirrored the solvency ratio of the LDI-Pension sector in the quarters leading up to the crisis. During the first half of 2022, solvency ratios for both sectors increased steadily alongside rising government debt yields globally. Both sectors then reached decade-long highs in solvency following the spike in gilt yields that followed the mini-budget.

Third, like LDI-pension schemes, life insurers nonetheless faced increasing liquidity needs as

their solvency improved. In 2022Q3 alone, life insurers made nearly £70 billion in variation margin payments related to their derivative positions. This number likely understates the true liquidity needs of the sector given it ignores any variation margin payments related to repo financing.

Given the clear similarities between the life insurance and LDI-pension sectors, one might expect similar behaviors following the mini-budget. Nonetheless, Figure A11 of the Internet Appendix indicates that this was not the case. While LDIs aggressively sold gilts and reduced repo borrowing (Figure 3), life insurers sold almost no gilts, slightly reduced their swap exposure, and modestly increased repo borrowing. Balance sheet segmentation can help explain this divergent behavior. Unlike pension schemes, life insurers held their financial debts (e.g., from derivatives) directly on their own balance sheets. Consequently, when gilt yields spiked after the mini-budget, life insurers were able to manage collateral calls and avoid large forced sales, as they were insulated from the issues of balance sheet segmentation that plagued LDIs and pension schemes.

4.3 The Behavior of Pooled LDIs

Next, we further highlight the importance of balance sheet segmentation by studying the behavior of pooled and single LDIs. Relative to single LDIs, it is natural to expect collateral transfers from pension schemes to LDIs to be slower for pooled vehicles. This is because recapitalizing a pooled LDI fund requires coordination among several pensions, which is likely to be harder during times of stress. Recapitalizing a pooled LDI is further complicated by a collective action issue, as pensions may hesitate on margin calls to see if others contribute first. In what follows, we differentiate between pooled and single LDIs to identify the effects of these logistical hurdles, acknowledging that balance sheet segmentation makes both types of LDIs subject to the recapitalization frictions outlined in Section 4.1.

We are specifically interested in how pooled and single LDI funds differentially adjusted their balance sheet after the announcement of the mini-budget. To isolate the effect of the pooled LDI structure, the ideal experiment would compare trading behavior of two LDIs that are identical in every aspect, except for their pooled structure. We attempt to approximate such an experiment

using a difference-in-difference regression of the form:

$$y_{it} = \lambda_i + \alpha_{m(i),d,t} + \Gamma X_{it} + \sum_w \theta_w \times 1(t = w) \times Pool_i + \varepsilon_{it} \quad (4)$$

where y_{it} is the cumulative amount of gilts purchased, measured par terms, or repo opened by LDI i in week t . All values of y_{it} are scaled by their initial level on September 2, restricted to be above -1, and winsorized at the 99% tail. Gilt flows are measured in par terms to strip out variation in gilt prices and isolate active portfolio adjustments. Large negative values for y_{it} indicate that LDI i sold a large fraction of its gilt positions or closed a large amount of repo. λ_i in regression (4) is an LDI fixed effect and $\alpha_{m(i),d,t}$ is a fixed effect based on the intersection of LDI manager $m(i)$, the quartiles of the initial duration of i 's gilt holdings, and time. X_{it} is a vector of lagged characteristics of LDI i in week t , including the three polynomials of leverage (debt-to-assets) and the ratio of repo to assets.²⁵ $1(t = w)$ is an indicator variable for whether week t equals w and $Pool_i$ is an indicator variable for if LDI i is pooled.²⁶ The regression is estimated using weekly data.

The coefficients of interest in regression (4) are the $\{\theta_w\}$, which capture the impact of the pooled LDI structure on gilt and repo activity. Their identification comes from comparing the behavior of pooled and single LDIs that share a manager and started with similar asset duration at a given point in time ($\alpha_{m(i),d,t}$), after controlling for time-invariant LDI-level differences like initial portfolio composition (λ_i) and time-varying characteristics like leverage (X_{it}). Because of the manager-by-duration-by-time fixed effect $\alpha_{m(i),d,t}$, the estimated $\{\theta_w\}$ are purged of any common influence that manager identity or initial portfolio duration may have had on LDI trading behavior. For instance, the difficulty in recapitalizing pooled LDIs may not stem from their structural characteristics but rather from investors' uncertainty regarding the skill of their managers, as in Shleifer and Vishny (1997). Alternatively, managers of pooled LDIs could vary in their internal risk limits or pooled LDIs may have held larger duration gilts entering the crisis, prompting their subsequent

²⁵We explore alternative specifications that avoid the inclusion of covariates while still accounting for balance sheet composition in Internet Appendix A.1.5.1.

²⁶There are a handful of cases in which an LDI manager places trades for all of its accounts using a single LEI, rendering it impossible to determine flows from pooled funds. We classify these LEIs as single LDIs, which should bias us against finding any differential behavior between pooled and single accounts.

selling. $\alpha_{m(i),d,t}$ absorbs these sorts of effects.

The estimated θ_w trace out the causal impact of the pooled LDI structure on trading behavior under the following assumption about parallel trends: pooled and single LDIs sharing a manager would have followed similar trends absent the announcement of the mini-budget, conditional on observable characteristics (e.g., leverage). Here, we have assumed that the mini-budget was the shock that pushed pooled LDIs into financial distress, though it is of course possible that the distress happened earlier.

Figure 9a plots estimates of $\{\theta_w\}$ when the outcome variable is gilt purchases, again relative to their initial stock on September 2. The full regression results are also presented in Table 5. The estimated coefficients in the table are normalized such that $\theta_w = 0$ for the week in which the mini-budget was announced (2022w38). The plot includes 95% confidence intervals, calculated using standard errors are double-clustering by LDI and week.

In the weeks leading up to the mini-budget, Figure 9a offers little visual evidence that pooled and single LDIs behaved differently in terms of gilt purchases. The magnitude of the point estimates for θ_w prior to the mini-budget are economically small and neither is statistically different from zero. However, a formal statistical test rejects the null hypothesis of no pre-trends at a 95% confidence levels ($p = 0.01$). We revisit potential violations of parallel trends below.

In the week following the mini-budget, the data indicates that pooled LDIs sharply reduced their gilt holdings, selling 11.63 pp ($t = -5.33$) more than their non-pooled counterparts. The relative rate of sales between the two groups then converged after the first week of the crisis. By the end of October, five weeks after the mini-budget, pooled LDIs had cumulatively liquidated 9.49 pp ($t = -2.33$) more of their gilt holdings compared to single LDIs.

Figure 9b asks how this cumulative estimate would change under a linear violation of parallel trends. In particular, it presents the adjusted point estimate based on Rambachan and Roth (2023) under a counterfactual in which differential gilt trading between pooled and single LDIs had continued linearly through the mini-budget announcement. In this case, the figure shows that pooled LDIs would still have sold meaningfully more gilts than single LDIs by the end of October.

Figure 9c displays the different-in-difference estimates $\{\theta_w\}$ when the outcome variable captures repo activity. Throughout the crisis period, pooled LDIs consistently reduced their repo borrowing, with the pace of these closures increasingly over time. By the end of October, they had closed 32.81 pp ($t = -6.04$) more of their repo compared to single LDIs, a pattern consistent with facing larger recapitalization frictions.

In Figure 9c, there is some evidence of differential pre-trends in repo activity between pooled and single LDIs, thus complicating the causal interpretation of our estimates. A potential explanation for the downward trend observed in the plot might be that pooled LDIs were under stress even before the mini-budget announcement, likely due to rising gilt yields throughout September. To better understand how a failure of parallel trends would impact our results, we adjust the point estimate of $\theta_w = 32.81$ pp in the fifth week of the crisis using the methods from Rambachan and Roth (2023). As before, we consider how the point estimate would change if the pre-trend had continued linearly through the announcement of the mini-budget. Unsurprisingly, the point estimate in this counterfactual would decline in magnitude. Still, even under this linear violation of parallel trends, LDIs would have reduced their repo borrowing by nearly 20 pp more than their non-pooled counterparts.

Overall, the results in Figure 9 support the view that the structure of pooled LDIs made it harder for them to recapitalize than single LDIs, hence why they sold more gilts and closed repo more aggressively. In principle, this interpretation could be complicated by any differences in capital structure between pooled and single LDIs, yet all of our regressions control for leverage and debt composition. Moreover, because we also control for ex-ante differences in gilt duration across the two groups (via the fixed effect $\alpha_{m(i),d,t}$), it is unlikely that differences in ex-ante liquidity risk can explain our results (Alfaro et al., 2023). Another possibility is pooled LDIs were unable to recapitalize as quickly because their pension investors were less sophisticated, perhaps because they tend to be smaller in overall size. However, we argue in Internet Appendix A.1.5.2 why sophistication is not the driving force behind Figure 9.

4.4 Price Pressure from Pooled LDIs

We now estimate the effect of selling by pooled LDIs on gilt prices in two ways. Our estimation strategies mirror those employed in Section 3.2, where we measured price pressure effects from all LDIs, not just pooled ones. Both methodologies are based on the following hypothetical scenario: imagine two gilts, N and P , with nearly identical maturities and equal holdings by LDIs. However, P is exclusively held by pooled LDIs, while N is held solely by single LDIs. Following a significant economic shock, such as the mini-budget, we can assess the price impact of any pooled LDI selling by comparing the price trajectory of P relative to N .

4.4.1 Difference-in-Difference Estimates

Our first strategy is a difference-in-difference regression that is similar to regression (3):

$$\log(p_{b,h,t}) = \gamma_b + \alpha_{h,t,m,l} + \sum_d \beta_d \times 1(t = d) \times D_b^p + \varepsilon_{b,t}, \quad (5)$$

where D_b^p is an indicator variable for whether bond b is heavily held by pooled LDIs. To construct D_b^p , we use the fraction P_b of each bond b held by pooled LDIs as of September 1. Bonds heavily held by pooled LDIs are those in the top tercile of P_b . $1(t = w)$ is an indicator for when date t equals d .

The main difference between regression (5) and (3) is the fixed effect, $\alpha_{h,t,m,l}$. In regression (3), we used a maturity-by-date-by-hour fixed effect. In regression (5), we augment it with an additional interaction term based on terciles of H_b , the fraction of bond b held by LDIs. That is, we create $\alpha_{h,t,m,l}$ by interacting indicators for whether bond b is in a given H_b -tercile with indicators for its maturity-decile, date, and hour. This means that the β_d 's in regression (5) are effectively identified by comparing bonds with similar maturities, held by LDIs in similar amounts, within the same hour and date.

The β_d 's trace out the causal impact of selling by *pooled* LDIs under the following assumption about parallel trends: absent pooled LDI selling, bonds in the same maturity bin and those that

were held similarly by all LDIs would have followed similar price trends from early September through the end of October. Because D_b^p is defined based on LDI holdings at the beginning of September, the regression estimates should not reflect private information that pooled LDIs may have received about the fundamental value of bonds they sold during the month.

Figure 10a presents the estimated β_d 's along with their 95% confidence interval, which are calculated using standard errors that are double-clustered by bond and date-by-hour. These estimates indicate that on September 28, the day of the Bank's initial intervention, selling pressure from pooled LDI funds caused a fire sale discount of 9.29% ($t = -3.01$). The intervention then seemingly eliminated the discount, only for it to reemerge in the weeks that followed. The estimated price pressure then disappears permanently by October 20, consistent with a fire sale.

The credibility of price pressure estimates in Figure 10a of course depend on whether our assumption of parallel trends holds. Early in the month, there is some evidence that bonds held by pooled LDI funds traded at a relative premium. However, by September 9, two weeks before the mini-budget, this premium had largely disappeared. Still, to probe the sensitivity of our estimates to this potential violation of parallel trends, we again use the results from Rambachan and Roth (2023). Figure 10b shows how the estimated fire sale discount on September 28 would change if the trend between bonds with high and low exposure to pooled LDIs had continued linearly from before announcement. Under this violation, the peak fire sale discount is still meaningful at 4.0% ($t = -2.38$).

4.4.2 Portfolio Sorts

For reasons discussed in Section 3.2, the difference-in-difference estimates of the effect of selling by pooled LDI selling likely understates the true impact. This is because the estimate is based on very narrow comparisons of bonds that have similar maturities and total exposure to all LDIs, thereby stripping out any the impact of pooled LDI selling on the broad level of the yield curve. As in Section 3.2.3, we now use simple portfolio sorts to get a sense of the total possible effect of pooled LDI selling. Specifically, we sort gilts into portfolios into terciles based on P_b , the fraction

of gilt b held by pooled LDIs as of September 1, 2022. We then compute the cumulative equal-weighted average return of each portfolio from September 22, 2022 through October 2022. Returns are based on end-of-day quotes from the Thompson Reuters Eikon database and relative to closing prices on September 22.

Figure 5b illustrates the performance disparity between the high and low- P_b portfolios. Following the mini-budget, gilts in the low- P_b portfolio declined by approximately 5.3%, compared to a decline of 25.0% in the high- P_b portfolio. These results indicate that selling by pooled LDIs contributed to a peak fire sale discount of 19.8% on September 27, the day prior to the Bank's intervention. Consistent with our previous findings, the return differential of 19.8% shown in Figure 5b vanishes after the Bank's initial intervention, reemerges in subsequent weeks, and then dissipates completely by the end of October. The relatively transient nature of this price disparity cuts against the notion that pooled LDIs trading after the mini-budget reflected private information about gilt fundamentals, especially since the portfolio assignments are based on positions at the beginning of September.

The implied fire sale discount of 19.8% in Figure 5b is larger than our difference-in-difference estimate of 9.29% for two reasons. First, because of their relatively higher duration, high- P_b gilts were more exposed to the fundamental shock of the mini-budget. Second, high- P_b gilts may also have high exposure to all types of LDIs, so the decline in their prices is not solely attributable to selling by pooled LDIs. Our difference-in-difference strategy effectively controls for both channels. For this reason, we view 19.8% as an upper bound on the impact of pooled LDI selling.

Overall, our analysis suggests that, at the crisis's peak, the selling pressure exerted by pooled LDIs led to gilt market discounts ranging from about 10-20%. These findings are consistent with the idea that selling pressure was highest in pooled LDIs because they faced larger barriers to recapitalization. More generally, our results support our broad hypothesis about why LDIs sold gilts, despite the LDI-Pension sector's substantial asset base (Section 4.1). In effect, collateral agreements meant that LDI balance sheets were partially segmented from those of their pension investors. Consequently, this segmentation, coupled with institutional frictions within pensions,

created a form of internal slow-moving capital that prevented LDIs from recapitalizing quickly enough to avoid liquidations.

5 Concluding Remarks

This paper documents a significant and persistent fire sale in the market for UK gilts during the period between September and October 2022. At its peak, selling pressure by LDIs conservatively resulted in price discounts on the order of 10%. LDIs were forced to sell despite the large asset base of their pension investors, suggesting that LDI debt was collateralized only by LDI assets. This contract structure meant that intensity of forced sales depended on the ability of pensions to transfer collateral or cash to LDIs on short notice, a task that was most difficult for pooled LDIs. Consistent with this mechanism, pooled LDIs sold gilts and reduced repo more aggressively than single LDIs, and the gilts they held were subject to larger fire sale discounts.

Price Pressure vs Fundamental Shocks Following the announcement of the mini-budget, the entire level of the gilt yield curve rose dramatically. How much of this level shift was driven by LDI selling versus the fundamental shock of the mini-budget? As a rough way to answer this question, consider gilts with maturities between ten and thirty years, where most LDI selling was concentrated (see Figure A8). After the mini-budget, the average price of these bonds dropped by about 15%. Our estimates from Section 3.2 indicate that LDI selling led to price discounts of 6.87%, accounting for approximately half of the decline in gilt prices. Furthermore, our analysis in Section 4.4.1 shows that sales by pooled LDIs resulted in discounts of 9.29%, suggesting that 60% of the price reduction was due to the fire sale. The remaining price movements can be attributed to fundamental shocks, most notably the fiscal policy changes in the mini-budget.

A potential concern with this back-of-the-envelope calculation is that our identification strategy relies on comparing gilts closely positioned along the yield curve. In principle, it is possible for these types of relative-value spreads to widen without significantly affecting the overall level of gilt yields. For example, in equity markets, changes in the spot-futures arbitrage spread for the

S&P 500 are only weakly correlated ($\rho = -6\%$) with returns on the S&P 500.²⁷

As a simple check of whether this is a relevant case for our particular episode, Figure 11 displays our price pressure estimates, which again are based on relative-value spreads between gilts of similar maturity, alongside the cumulative average return of *all* gilts from September 22 onward. If LDI selling pressure caused relative-value spreads to widen without significantly affecting the overall level of yields, the two series should show a weak correlation, akin to the relationship between S&P 500 spot-futures arbitrage spreads and S&P 500 returns. However, the figure shows that the level of gilt yields strongly tracks our price pressure estimates. In fact, a regression of average gilt returns on our price pressure estimates yields an R^2 of 76%, suggesting that the bulk of variation in the average level of gilt yields during this time was driven by LDI selling.²⁸ In Internet Appendix A.1.6, we make a similar argument using gilt-OIS swap spreads.

Policy Implications On September 28, the Bank of England intervened in the gilt market, offering to purchase long-dated nominal gilts only until October 14. On October 11, inflation-linked gilts were added to the Bank’s operations. In the end, the Bank purchased £19.3 bn of gilts, of which £12.1 bn were nominal and £7.2 bn were inflation-linked.

Our results help rationalize the design of the Bank’s interventions during the fire sale. As discussed in Diamond and Rajan (2009) and Shleifer and Vishny (2011), one option in this situation was to provide liquidity to distressed sellers using short-term lending facilities. A downside of this strategy is that it would necessarily add more short-term debt to LDIs, who were already burdened by high leverage. Of course, the Bank could have offered short-term loans directly to pensions, enabling them to then recapitalize their LDI investments. However, this strategy would have posed its own challenges, particularly for pooled funds, due to the need to coordinate with numerous pension scheme trustees who may not have been prepared for such an event.

A second intervention option, which is what the Bank ultimately pursued, was to directly pur-

²⁷This correlation is based on arbitrage spreads from Siriwardane et al. (2023) between 1/1/2010 and 12/16/2022.

²⁸This argument implicitly assumes that fundamental shocks to gilts are uncorrelated with any LDI selling pressure that lead to a widening of relative-value spreads. This assumption is arguably more plausible after the announcement of the mini-budget and before the resignation of Prime Minister Liz Truss on October 20, 2022. During this window, we still obtain an R^2 of 78%.

chase securities. But how long should such an intervention last? Our analysis suggests that setting a definitive timeframe was potentially advantageous for the Bank given the nature of the fire sale. As discussed in Section 4.1, pensions collectively held sufficient collateral to fully cover the debt on LDI balance sheets, which, in theory, should have minimized the necessity for LDIs to liquidate their gilt holdings. The primary challenge lay in the transfer of collateral between pensions and LDIs, a process further complicated by the need for pooled funds to coordinate capital injections from multiple pension investors (Section 4.3). By putting a fixed window on its purchases, the Bank created incentives for pensions to resolve these operational and coordination problems, along with any other institutional frictions (e.g., trustee approval, due diligence, investment consultations), in a timely manner.

In the aftermath of the gilt market crisis, there has also been considerable policy debate about the future regulation of LDIs. Many proposals have centered on limiting leverage and requiring liquidity buffers. For instance, in 2023, the Financial Policy Committee (FPC) of the Bank of England recommended to The Pensions Regulator that LDIs should maintain liquidity buffers capable of withstanding yield shocks of at least 250 basis points. Following this recommendation, regulatory bodies in Ireland and Luxembourg, where most LDIs are domiciled, have instituted comparable resilience standards.

Our findings speak to the issue of LDI regulation because they suggest that balance sheet segmentation and operational frictions were fundamental drivers of the crisis. Therefore, regulatory reforms that aim to better integrate the balance sheets of pensions and their LDI investments could be effective at averting future crises. The FPC's recommendations on streamlining operational processes at pensions and LDIs are aligned with this thinking. Further measures could include establishing automated collateral transfer facilities during periods of market stress (Kodres, 2023) and encouraging pension schemes to create repo facilities that can quickly generate liquidity from existing assets. Additionally, encouraging pensions to invest in LDIs through segregated mandates would naturally facilitate better integration of their balance sheets.

This discussion raises a broader question: why were LDI funds originally structured in a man-

ner that seemingly did not prioritize preventing fire sales? Given the mutual benefits for both LDIs and their pension investors in avoiding such scenarios, the private sector's failure to adapt its contracting and operational arrangements in advance of the crisis is puzzling. These considerations are crucial for understanding how decisions regarding capital structure and contract design, made well before crises emerge, can be optimized to avert future fire sales.

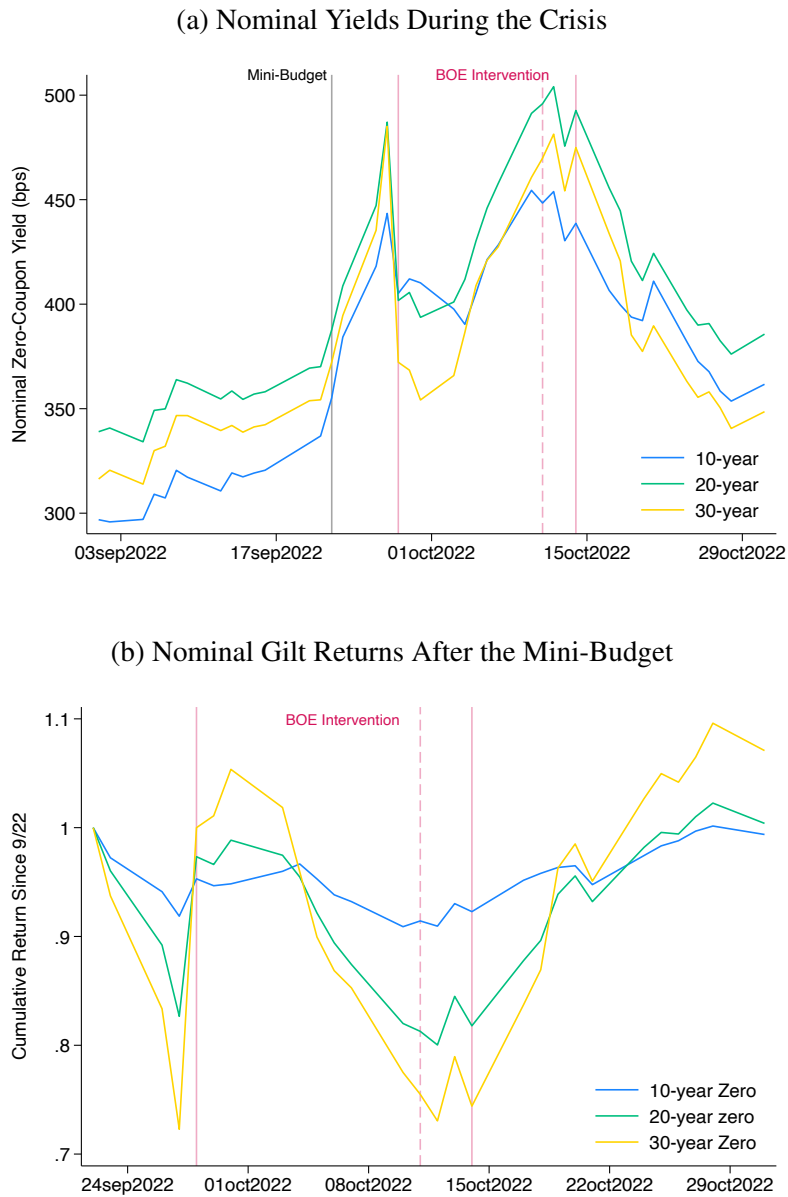
References

- Abreu, D. and M. K. Brunnermeier (2003). Bubbles and crashes. *Econometrica* 71(1), 173–204.
- Alexander, P., R. Fakhoury, T. Horn, W. Panjwani, and M. Roberts-Sklar (2023). Financial stability buy/sell tools: a gilt market case study.
- Alfaro, L., S. Bahaj, R. Czech, J. Hazell, and I. Neamt (2023). Lash risk and interest rates. Technical report, Bank of England.
- Anderson, N. and J. Sleath (2001, March). New estimates of the UK real and nominal yield curves. Bank of England working papers 126, Bank of England.
- Bian, J., Z. He, K. Shue, and H. Zhou (2018, September). Leverage-induced fire sales and stock market crashes. Working Paper 25040, National Bureau of Economic Research.
- Bicu-Lieb, A., L. Chen, and D. Elliott (2020). The leverage ratio and liquidity in the gilt and gilt repo markets. *Journal of Financial Markets* 48, 100510.
- Brandt, M. W. (2010). Chapter 5 - portfolio choice problems. In Y. AIT-SAHALIA and L. P. HANSEN (Eds.), *Handbook of Financial Econometrics: Tools and Techniques*, Volume 1 of *Handbooks in Finance*, pp. 269–336. San Diego: North-Holland.
- Breeden, S. (2022). Risks from leverage: how did a small corner of the pensions industry threaten financial stability? Speech, Bank of England.
- Campbell, J. Y., S. Giglio, and P. Pathak (2011, August). Forced sales and house prices. *American Economic Review* 101(5), 2108–31.
- Campbell, J. Y. and L. M. Viceira (2002, 01). Introduction. In *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.
- Chen, R. and E. Kemp (2023). Putting out the nbfire: Lessons from the uk’s liability-driven investment (ldi) crisis. Technical report, International Monetary Fund.
- Coval, J. and E. Stafford (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics* 86(2), 479–512.
- Czech, R. and G. Pinter (2020, November). Informed trading and the dynamics of client-dealer connections in corporate bond markets. Bank of England working papers 895, Bank of England.
- Diamond, D. W. and R. G. Rajan (2009, April). Fear of fire sales and the credit freeze. Working Paper 14925, National Bureau of Economic Research.
- Du, W., A. Tepper, and A. Verdelhan (2018). Deviations from covered interest rate parity. *The Journal of Finance* 73(3), 915–957.
- Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. *The Journal of Finance* 65(4), 1237–1267.

- Duffie, D. and B. Strulovici (2012). Capital mobility and asset pricing. *Econometrica* 80(6), 2469–2509.
- Dunne, P., A. Ghiselli, F. Ledoux, and B. McCarthy (2023, September). Irish-Resident LDI Funds and the 2022 Gilt Market Crisis. Financial Stability Notes 7/FS/23, Central Bank of Ireland.
- Ellul, A., C. Jotikasthira, and C. T. Lundblad (2011). Regulatory pressure and fire sales in the corporate bond market. *Journal of Financial Economics* 101(3), 596–620.
- Falato, A., A. Hortacsu, D. Li, and C. Shin (2021). Fire-sale spillovers in debt markets. *The Journal of Finance* 76(6), 3055–3102.
- Gerba, E. and P. Katsoulis (2021, December). The repo market under Basel III. Bank of England working papers 954, Bank of England.
- Giese, J., M. Joyce, J. Meaning, and J. Worlidge (2024). Do preferred habitat investors exist? evidence from the uk government bond market. *Economics Letters* 234, 111462.
- Greenwood, R., S. G. Hanson, and G. Y. Liao (2018, 04). Asset Price Dynamics in Partially Segmented Markets. *The Review of Financial Studies* 31(9), 3307–3343.
- Greenwood, R. and D. Vayanos (2010). Price pressure in the government bond market. *The American Economic Review* 100(2), 585–590.
- Gromb, D. and D. Vayanos (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics* 66(2), 361–407. Limits on Arbitrage.
- Grossman, S. J. and M. H. Miller (1988). Liquidity and market structure. *The Journal of Finance* 43(3), 617–633.
- Hauser, A. (2022). Thirteen days in October: how central bank balance sheets can support monetary and financial stability. Speech, Bank of England.
- Hu, G. X., J. Pan, and J. Wang (2013). Noise as information for illiquidity. *The Journal of Finance* 68(6), 2341–2382.
- Jurkatis, S. (2024). An approach to cleaning mifid ii corporate bond transaction reports.
- Kodres, L. (2023). Too many cooks, not enough risk management: Gilt market dysfunction and liability-driven investment (ldi). *SSRN*.
- Lou, D., H. Yan, and J. Zhang (2013, 06). Anticipated and Repeated Shocks in Liquid Markets. *The Review of Financial Studies* 26(8), 1891–1912.
- Mitchell, M., L. H. Pedersen, and T. Pulvino (2007, May). Slow moving capital. *American Economic Review* 97(2), 215–220.
- Mitchell, M. and T. Pulvino (2012). Arbitrage crashes and the speed of capital. *Journal of Financial Economics* 104(3), 469–490. Market Institutions, Financial Market Risks and Financial Crisis.

- Pinter, G., C. Wang, and J. Zou (2024). Size discount and size penalty: trading costs in bond markets. *The Review of Financial Studies Forthcoming*.
- Pulvino, T. C. (1998). Do asset fire sales exist? an empirical investigation of commercial aircraft transactions. *The Journal of Finance* 53(3), 939–978.
- Rambachan, A. and J. Roth (2023, 02). A More Credible Approach to Parallel Trends. *The Review of Economic Studies* 90(5), 2555–2591.
- Shleifer, A. (1986). Do demand curves for stocks slope down? *The Journal of Finance* 41(3), 579–590.
- Shleifer, A. and R. W. Vishny (1992). Liquidation values and debt capacity: A market equilibrium approach. *Journal of Finance* 47(4), 1343–1366. Reprinted in Michael J. Brennan, ed., *The Theory of Corporate Finance*, Edward Elgar Publishing Company, 1996.
- Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *The Journal of Finance* 52(1), 35–55.
- Shleifer, A. and R. W. Vishny (2011). Fire sales in finance and macroeconomics. *Journal of Economic Perspectives* 25(1), 29–48.
- Siriwardane, E. N. (2019). Limited investment capital and credit spreads. *The Journal of Finance* 74(5), 2303–2347.
- Siriwardane, E. N., A. Sunderam, and J. Wallen (2023). The rise of alternatives. *SSRN*.
- van Binsbergen, J. H. and M. W. Brandt (2016). *Optimal Asset Allocation in Asset Liability Management*, Chapter 8, pp. 147–168. John Wiley and Sons, Ltd.
- Van Horen, N. and A. Kotidis (2018, August). Repo market functioning: the role of capital regulation. Bank of England working papers 746, Bank of England.
- Vayanos, D. and J.-L. Vila (2021). A preferred-habitat model of the term structure of interest rates. *Econometrica* 89(1), 77–112.

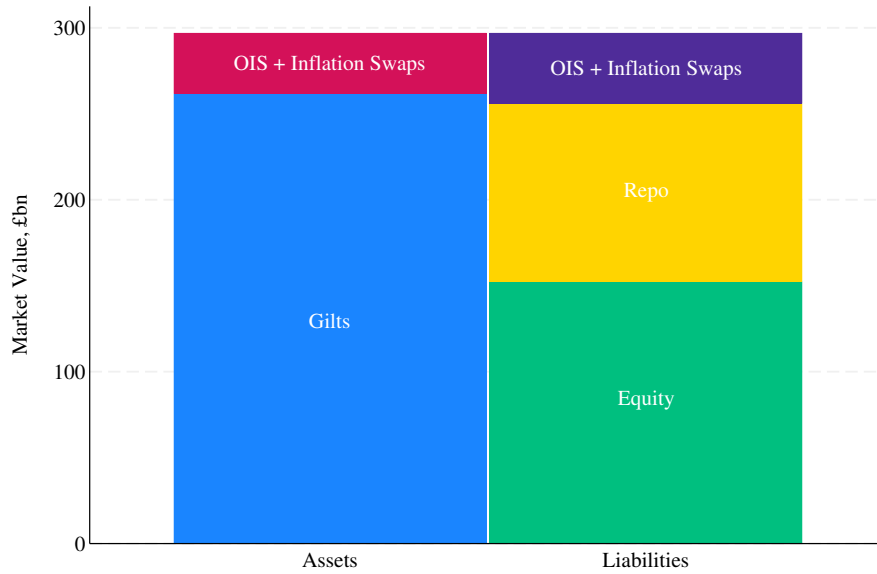
Figure 1: Zero-Coupon Gilt Yields and Prices




Notes: Panel (a) shows the yields of constant-maturity 10, 20, and 30 zero-coupon nominal gilts, respectively, from 9/1/2022 through 10/31/2022. Panel (b) shows the cumulative gross return of these same instruments from 9/22/2022 onward. The orange vertical line in the panel (a) is 9/22/2022, the day before the announcement of the mini-budget. In both panels, the first solid red vertical line is on 9/28/2022, the first day in which the Bank of England intervened in the gilt market. The dotted vertical red line on 10/11/2022 marks the day on which the Bank began purchasing index-linked gilts. The last vertical red line marks when the Bank’s intervention program ended on 10/14/2022. Data is daily and taken directly from the zero-coupon yield curve on the Bank of England’s website.

Figure 2: LDI Balance Sheet

(a) Aggregate LDI Balance Sheet (9/2/2022)

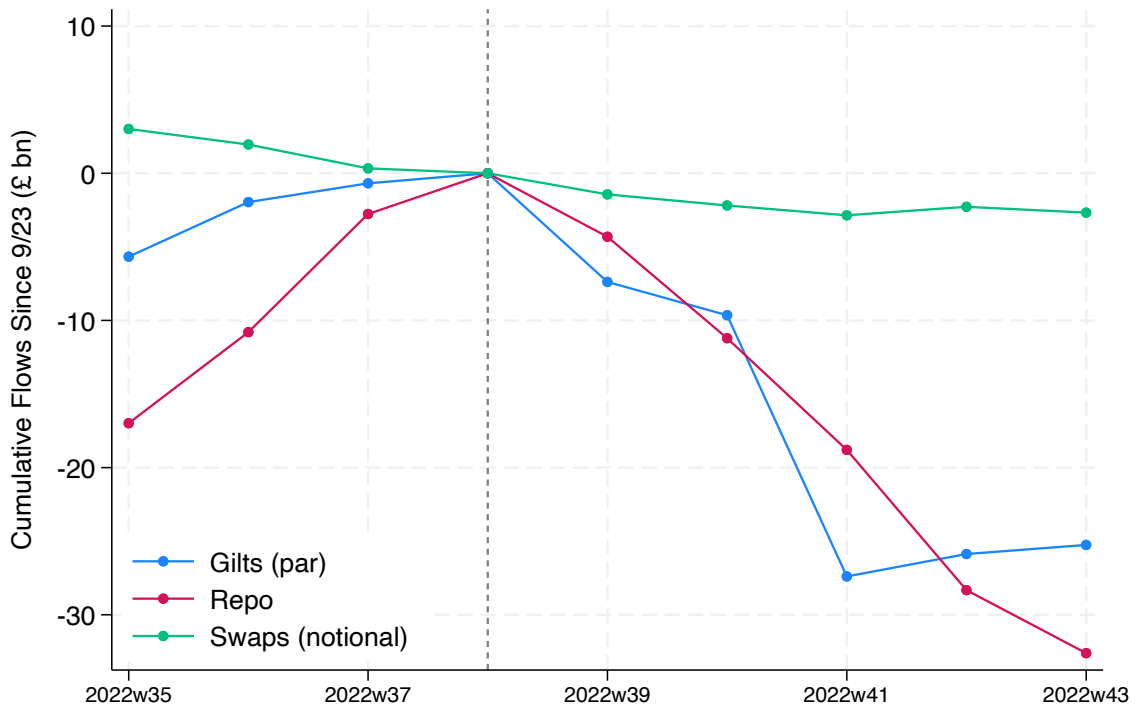


(b) Time-Series of LDI Leverage (Assets-to-Equity)



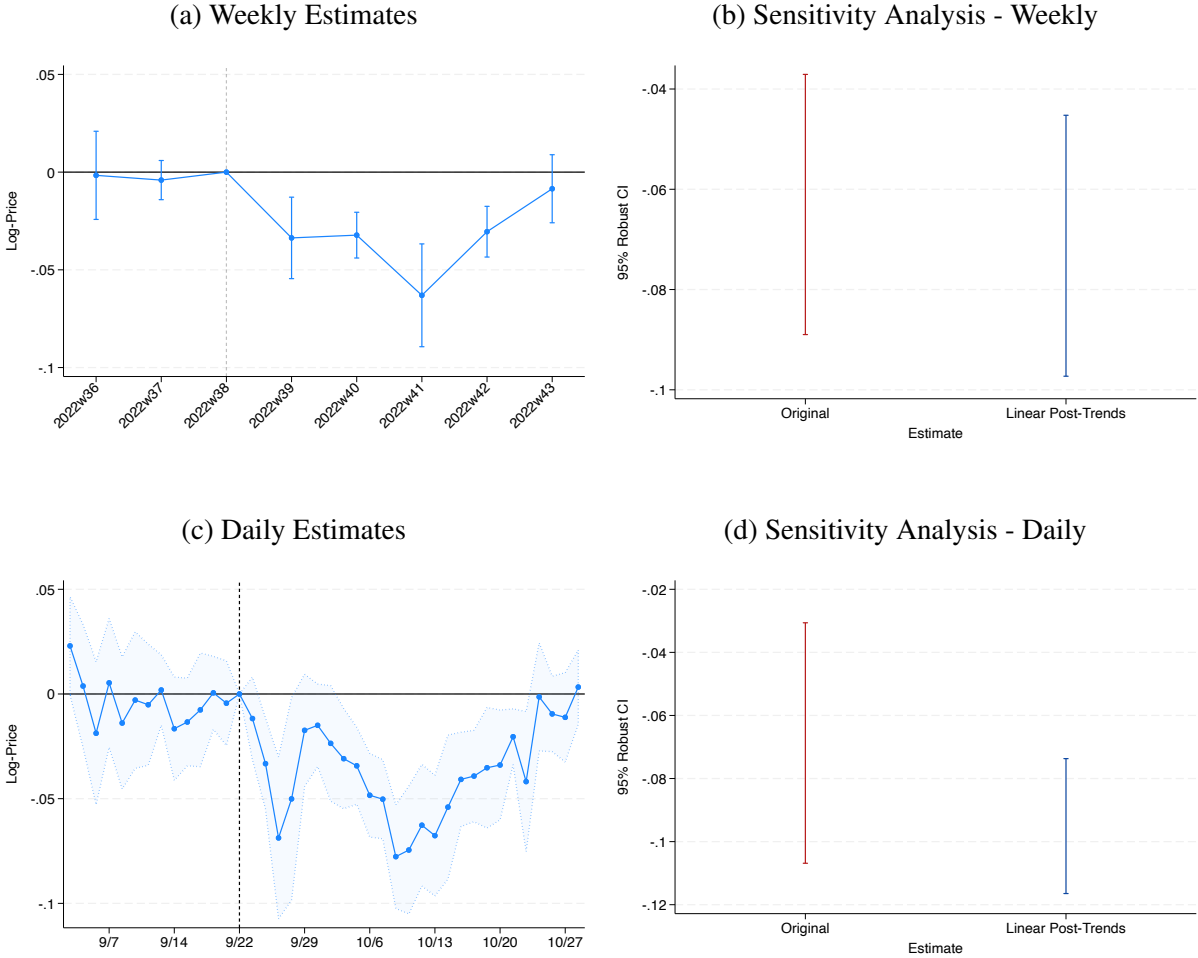
Notes: Panel (a) shows the aggregate balance sheet of the LDI sector as of 9/2/2022. The stock of Gilts is based on cumulative flows since 2017 from the MIFID II database. It is reported in market value terms. OIS + inflation swap assets equal the market value of the receive-fixed leg of OIS swaps plus the receive-floating (inflation) leg of inflation swaps, both from the EMIR TR. OIS + inflation swap liabilities equal the notional value of the pay-floating leg of OIS swaps plus the notional value of the pay-fixed leg of inflation swaps. Repo positions equal par borrowing amounts from the SMMD database. Aggregate LDI equity is the difference between assets and liabilities and is left blank in the plot. Panel (b) shows how the asset-to-equity ratio of the sector evolved from September through October 2022. It is constructed in the same manner as panel (a) and is at the weekly frequency. See Section 2.3 for complete details.

Figure 3: LDI Trading Behavior



Notes: This figure shows the trading behavior of LDIs from September through October 2022. Gilt flows are based on the change in the par amount of bonds held, swap flows are based on the change in net notional receiving fixed OIS swaps and paying-fixed inflation swaps, and repo flows are based on the change in outstanding repo borrowing. All flows are indexed to 9/23 (vertical dashed line), the day of the mini-budget.

Figure 4: Price Pressure from LDIs - Difference-in-Difference Estimates

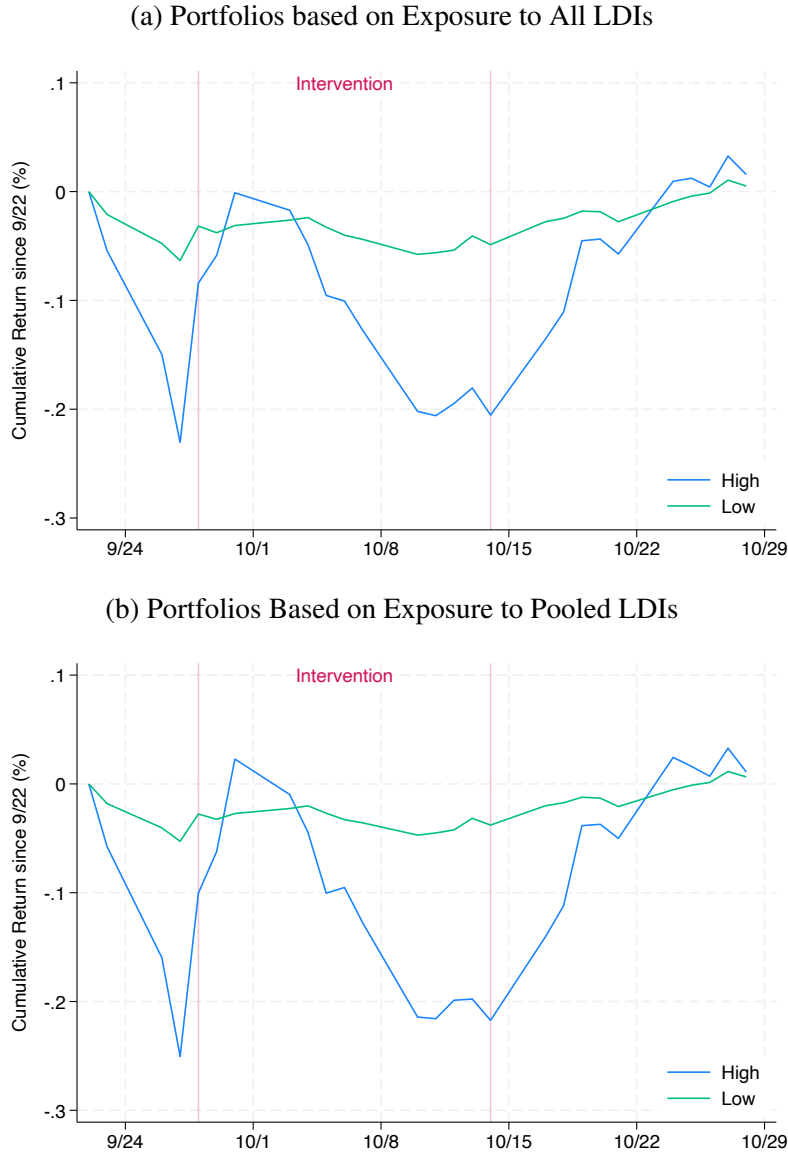


Notes: This figure shows the estimated β 's from the following difference-in-differences regression:

$$\log(p_{b,h,t}) = \gamma_b + \alpha_{h,t,m} + \sum_w \beta_w \times 1(t \in w) \times D_b + \varepsilon_{b,t},$$


where $p_{b,h,t}$ is the median-transaction price for Gilt b in hour h of day t . γ_b is a Gilt fixed effect and $\alpha_{h,t,m}$ is an date-by-hour-by-maturity fixed effect. $1(t \in w)$ is an indicator variable for whether date t is in week w and D_b is an indicator for whether Gilt b was held heavily by LDIs of September 2, 2022. D_b is defined based on whether LDI holdings of b were above median. Panel (a) shows estimates of the regression described above. Following Rambachan and Roth (2023), Panel (b) shows how most-negative estimate of β_w would change if the trend between Gilt with $D_b = 1$ and $D_b = 0$ had continued linearly through 2022w38, the week of the mini-budget. Panel (c) shows estimates of β when they are allowed to vary by day, not week. Panel (d) repeats the robustness check in panel (b) for the daily estimates of β . Standard errors in the plot are double-clustered by Gilt b and day-by-hour.

Figure 5: Price Pressure from LDIs - Portfolio Sorts



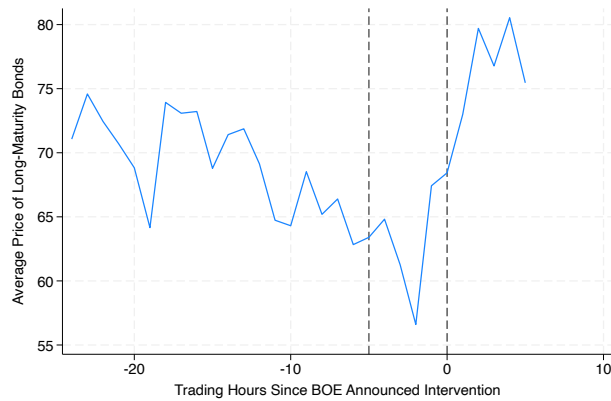
Notes: This figure shows the cumulative average performance of Gilt portfolios constructed based on exposure to LDIs (panel a) and pooled LDIs (panel b). In panel (a), we first sort gilts into terciles based on H_b , the fraction of gilt b held by LDIs as of 9/1/2022. We then plot the performance of the low and high-tercile portfolios relative to their closing prices on 9/22/2022, the day before the mini-budget. In panel (b), we sort gilts into terciles based on P_b , the fraction of b held by pooled LDIs as of 9/2/2022. We then plot the performance of the low and tercile portfolios relative to their closing prices on 9/22/2022. The red vertical lines in the plot correspond to the start and end dates of the Bank's intervention. Data is daily and prices are based on end-of-day quotes from the Thompson Reuters Eikon database.

Figure 6: Timing the Bottom



(a) Cumulative Gilt Purchases

(b) Concentrated Buying



(c) Price Reaction of Long-Term Gilts

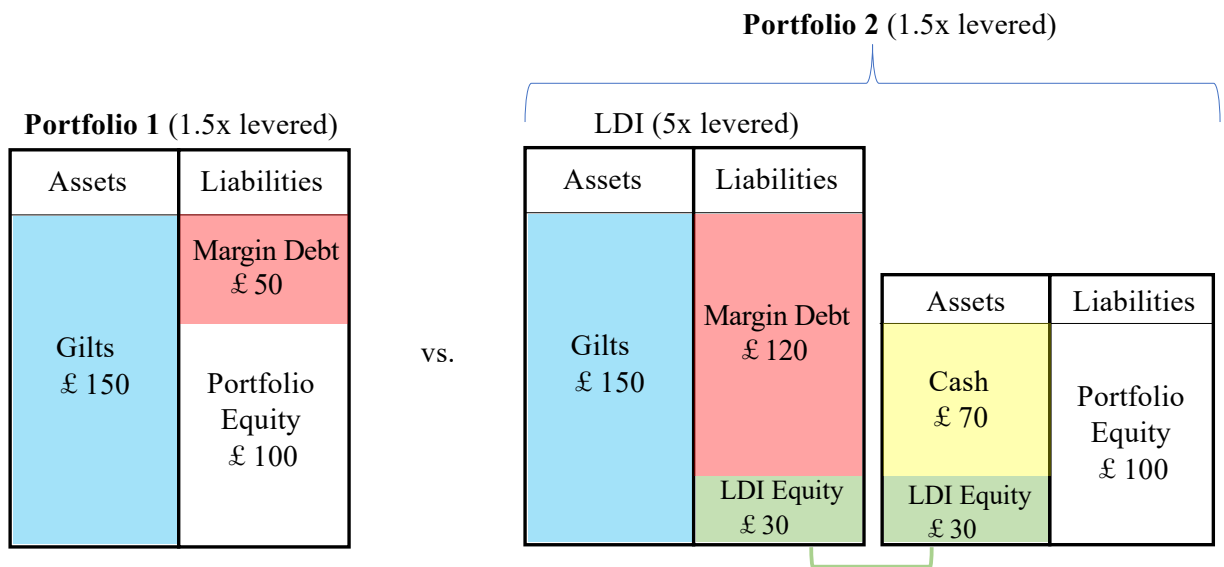
Notes: Panel (a) of this figure shows cumulative Gilt purchases (or sales) for different investor types in the hours around when the BOE first announced it would intervene in the Gilt market at 11 am GMT on 9/28/2022. Panel (b) shows the share of purchases by the top k hedge funds as a function of k , for different windows during the crisis. The line labeled “Morning of 9/28” is based on all trades made overnight and before 11 am on 9/28. The line labeled “Week after Mini-Budget” excludes trades made during this window. Panel (c) shows the average price quote of gilts with at least 20-years to maturity over the same period.

Figure 7: Balance Sheet of LDI-Pension Sector

Assets (£1,581 bn)	Liabilities
Gov't Debt (£659)	Present Value of Benefits (£1,184)
Corporate Debt (£377)	
Public Equities (£191)	Net Financial Debt (£77)
Alternatives (£354)	Funding Balance (£320)

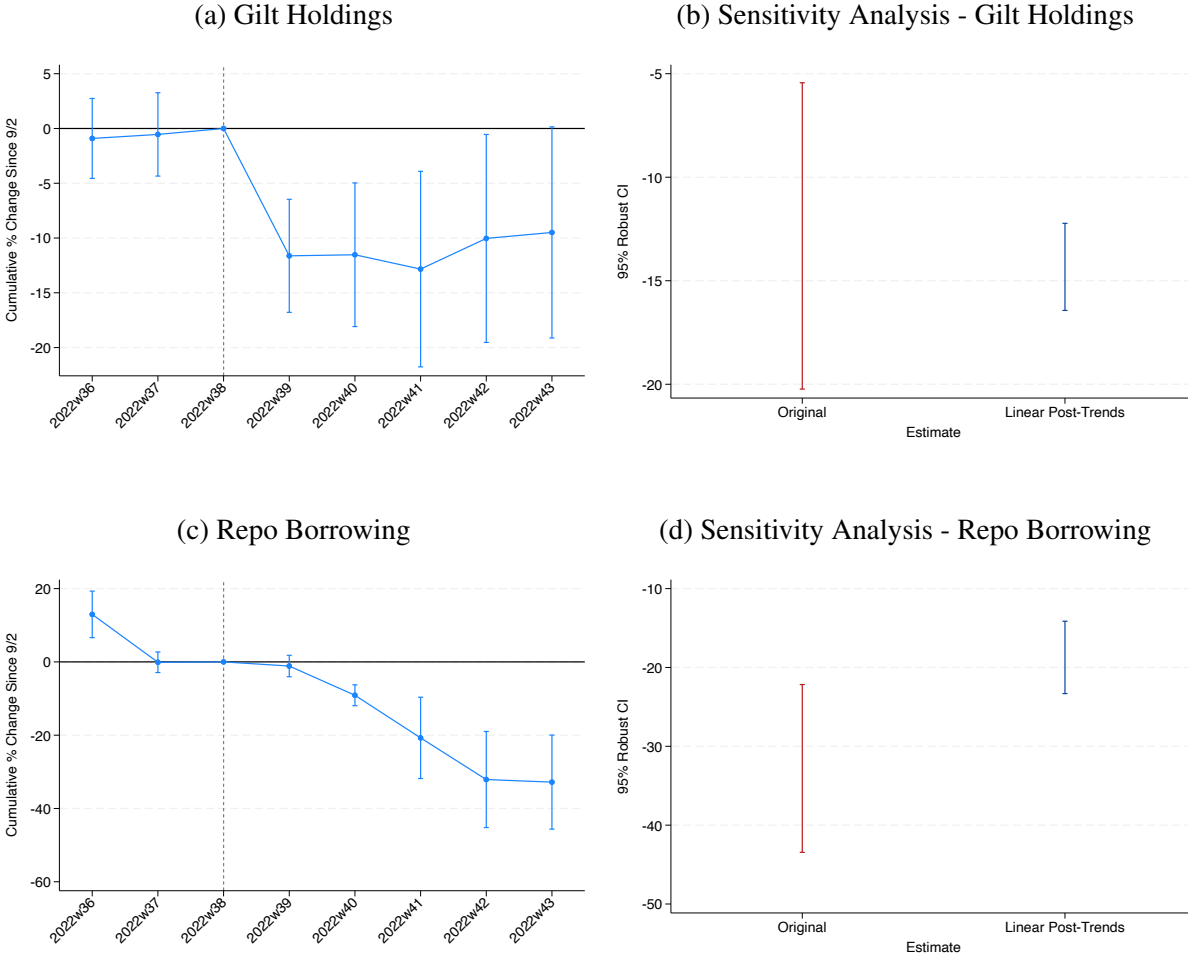
Notes: This figure visualizes the combined balance sheets of the LDI and U.K. pension sectors as of August 31, 2022. The consolidated balance sheet is based on data from the U.K Pension Protection Fund. Net financial debt includes any leverage used by LDIs in which pensions invest and is net of all cash held by both entities. See Section 4.1 for complete details.

Figure 8: Example of Siloed Leverage



Notes: This plot depicts the balance sheets for two hypothetical portfolios. Both are identical on a consolidated basis. In portfolio 2, the investor (e.g., a pension) invests in an LDI fund that uses leverage to purchase Gilts. The item SPV equity represents equity in the the “Special Purpose Vehicle”, namely the LDI.

Figure 9: Trading Behavior of Pooled vs Single LDIs

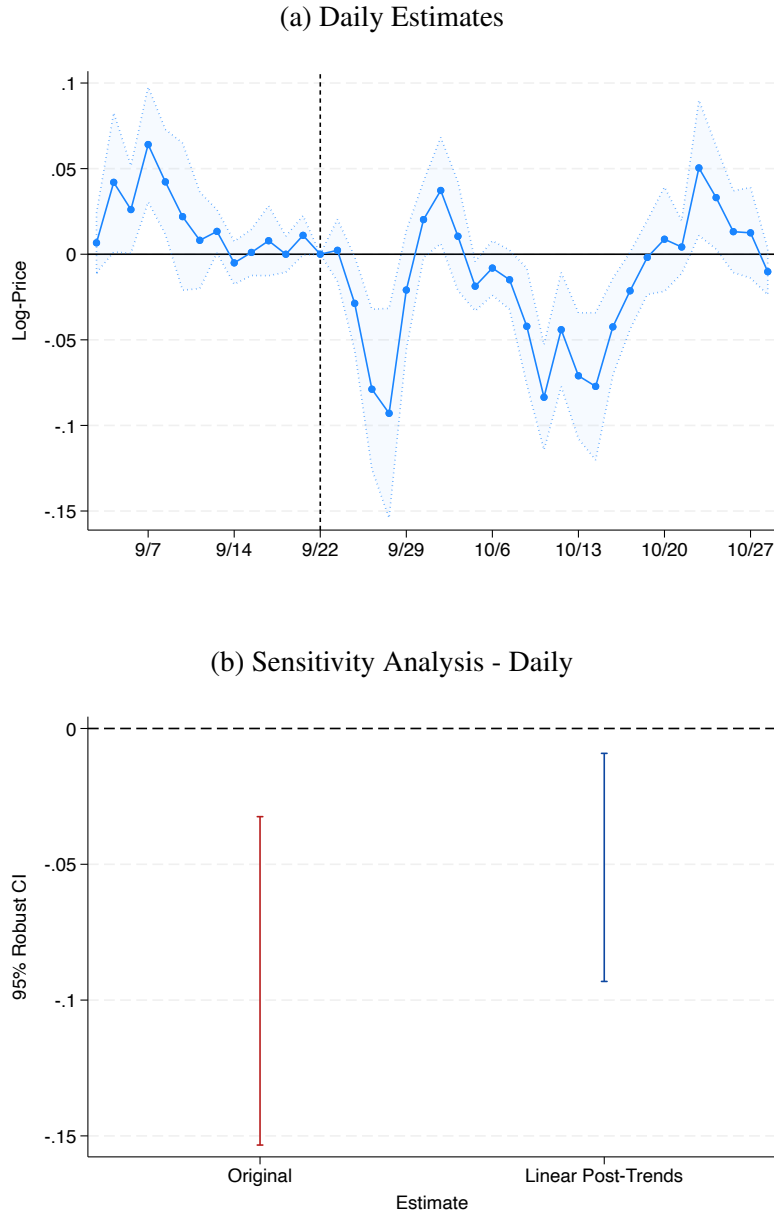


Notes: This figure shows estimates of the following difference-in-differences regression:

$$y_{it} = \lambda_i + \alpha_{m(i),d,t} + \Gamma X_{it} + \sum_w \theta_w \times 1(t = w) \times Pool_i + \varepsilon_{it}$$

where y_{it} is one of two outcome variables for LDI i managed by $m(i)$ on date t . The first outcome (Panel a) is the percentage change in Gilt holdings (par values) and the second (Panel c) is the percentage change in repo outstanding, both relative to their initial stocks on 9/2/2022. λ_i is an LDI-level fixed effect and $\alpha_{m(i),d,t}$ is a manager-by-duration-by-time fixed effect, where duration is defined based on quartiles of the maturity of i 's gilt holdings at the beginning of September. $1(t = w)$ is an indicator variable for whether week t equals w and $Pool_i$ is an indicator for whether LDI i is a pooled vehicle. X_{it} is a vector of controls for i that contains its lagged debt-to-asset ratio (three polynomials) and repo-to-asset ratio. Following Rambachan and Roth (2023), Panel (b) shows how most-negative estimate of θ_w for Gilt holdings would change if the trend between pooled and non-pooled LDIs had continued linearly through 2022w38, the week of the mini-budget. Panel (d) repeats the robustness test when the outcome variable is repo borrowing. Standard errors are clustered by account and week. Data is weekly.

Figure 10: Price Pressure from Pooled LDIs - Difference-in-Difference Estimates



Notes: This figure shows the estimated β 's from the following difference-in-differences regression:

$$\log(p_{b,h,t}) = \gamma_b + \alpha_{h,t,m} + \sum_d \beta_d \times 1(t=d) \times D_b^p + \varepsilon_{b,t},$$

where $p_{b,h,t}$ is the median-transaction price for Gilt b in hour h of day t . γ_b is a Gilt fixed effect and $\alpha_{h,t,m,s}$ is a date-by-hour-by-maturity-by-LDI exposure fixed effect. $1(t=d)$ is an indicator variable for whether date t equals d and D_b^p is an indicator for whether Gilt b was held heavily by pooled LDIs of September 2, 2022. D_b^p is defined based on whether pooled LDI holdings of b were above median. Panel (a) shows estimates of the regression described above. Following Rambachan and Roth (2023), Panel (b) shows how most-negative estimate of β_d would change if the trend between Gilts with $D_b^p = 1$ and $D_b^p = 0$ had continued linearly through announcement of the mini-budget on 9/23/2022. Standard errors in the plot are double-clustered by Gilt b and day-by-hour.

Figure 11: Relative Value Spreads vs the Level of Gilt Yields



Notes: This figure shows daily estimates of LDI selling on gilt prices based on gilts of nearby maturities (relative-value spreads, left axis) against the cumulative average return of all gilts from September 22, 2022 onward (right axis). Price pressure estimates are taken directly from Figure 4c. Cumulative returns on gilts are based on end-of-day quotes from the Thompson Reuters Eikon database.

Table 1: Summary Statistics of the LDI Sector

	Aggregate	Mean	Std. Dev.	p25	p50	p75	<i>N</i>
Assets (£bn)	297.04	1.25	5.12	0.11	0.32	0.82	238
Gilt-to-Assets	0.88	0.90	0.97	0.55	1.00	1.00	238
Fixed-OIS Swaps-to-Assets	0.06	0.00	1.12	0.00	0.00	0.27	238
Floating-Inf. Swaps-to-Assets	0.05	0.09	0.35	0.00	0.00	0.11	238
Debt-to-Assets	0.49	0.55	0.53	0.26	0.61	0.89	238
Repo-to-Assets	0.35	0.44	0.89	0.00	0.25	0.59	238
Floating-OIS Swaps-to-Assets	0.09	0.03	1.17	0.00	0.00	0.34	238
Fixed-Inf. Swaps-to-Assets	0.05	0.08	0.30	0.00	0.00	0.09	238
% Pooled (Count)	35						
% Pooled (Assets)	19						
Number of Managers	18						

Notes: This table shows summary statistics on the LDI sector as of September 2. Gilt-to-assets measures the fraction of assets held in cash Gilt positions, including linkers. Fixed swaps-to-assets measures the fraction of assets that are receiving fixed in an interest swap (net). Debt-to-assets measures the fraction of assets financed through repo or by paying floating in an interest rate swap. Repo-to-assets is the fraction of assets financed through repo. Data is at the LEI level. The number of managers counts the unique number of LDI managers across all LEIs. % pooled is the fraction of accounts that are designated as pooled. % Domiciled in the UK is the fraction of accounts that are domiciled in the UK.

Table 2: Weekly Difference-in-Difference Estimates of Price Pressure

	Log(Price)		
	(1)	(2)	(3)
2022w36	-0.00 (-0.15)	-0.01 (-0.61)	-0.01 (-0.55)
2022w37	-0.00 (-0.81)	-0.01 (-1.19)	-0.00 (-0.85)
2022w39	-0.03** (-3.21)	-0.03** (-2.74)	-0.03** (-2.89)
2022w40	-0.03** (-5.47)	-0.03** (-5.42)	-0.03** (-5.36)
2022w41	-0.06** (-4.76)	-0.06** (-4.61)	-0.06** (-4.83)
2022w42	-0.03** (-4.67)	-0.03** (-4.51)	-0.03** (-4.73)
2022w43	-0.01 (-0.97)	-0.01 (-1.53)	-0.01 (-1.61)
Maturity Bins	5	10	15
$p: \beta_{2022w36} = \beta_{2022w37}$	0.83	0.91	0.89
Adj- R^2	0.97	0.97	0.97
Total R^2	22,220	22,094	21,814

Notes: This figure shows the estimated β 's from the following difference-in-differences regression:

$$\log(p_{b,h,t}) = \gamma_b + \alpha_{h,t,m} + \sum_w \beta_w \times 1(t \in w) \times D_b + \varepsilon_{b,t},$$

where $p_{b,h,t}$ is the median-transaction price for Gilt b in hour h of day t . γ_b is a Gilt fixed effect and $\alpha_{h,t,m}$ is a date-by-hour-by-maturity fixed effect. $1(t \in w)$ is an indicator variable for whether date t is in week w and D_b is an indicator for whether Gilt b was held heavily by LDIs of September 2, 2022. D_b is defined based on whether LDI holdings of b were above median. The omitted week in the regression is 2022w38, the week of the mini-budget announcement. Each column of the table uses a different number of maturity bins to construct the fixed effect $\alpha_{h,t,m}$. The row $p: \beta_{2022w36} = \beta_{2022w37}$ shows the p -value from testing the null of no parallel pre-trends. Standard errors in the plot are double-clustered by Gilt b and day-by-hour.

Table 3: Gilt and OIS Positions by Institution Type

	9/1		
	Gilts	OIS Swaps	Total
Hedge Funds	23.5	-227.6	-204.1
Dealers	58.7	-69.8	-11.1
Insurance Companies	101.5	33.5	135.0
Asset Managers (non-LDI)	86.8	46.7	133.5

Notes: This table shows the stock of gilt holdings and OIS swaps (indexed to SONIA) held by different institution types as of 9/1/2022. Data for the gilt holdings of dealers/banks is taken from Form BT (item 32D), collected by the Bank of England and are reported in market values. All other positions are based on the MIFID II and EMIR databases. Gilt holdings are reported in par terms for all institutions except dealers/banks. OIS swaps are reported as the net amount of notional in which the institution is receiving fixed and paying floating.

Table 4: Consolidated Balance Sheet of LDI-Pension Sector

Assets		Liabilities	
Fixed Income	1,036	Pension Obligations	1,184
<i>Sovereign Debt</i>	659		
<i>Corporate</i>	377		
<i>IG</i>	257		
<i>HY/Private Credit</i>	120		
Public Equities	191	Financial Debt	77
<i>UK</i>	21		
<i>Developed</i>	149		
<i>Emerging</i>	21		
Alternatives	354	Funding Balance	320
Total	1,581	Totals	1,581

Notes: This tables presents granular asset positions of the combined balance sheets of the LDI and U.K. pension sectors as of August 31, 2022. The consolidated balance sheet is based on data from the U.K Pension Protection Fund. Net financial debt includes any leverage used by LDIs in which pensions invest and is net of all cash held by both entities. See Section 4.1 for complete details.

Table 5: Weekly Difference-in-Difference Estimates of Pooled Structure

	% Change Since 9/2	
	(1) Gilts	(2) Repo
2022w36	-0.91 (-0.59)	12.97** (4.83)
2022w37	-0.54 (-0.34)	-0.10 (-0.09)
2022w39	-11.63** (-5.33)	-1.13 (-0.91)
2022w40	-11.53** (-4.16)	-9.10** (-7.56)
2022w41	-12.84** (-3.40)	-20.73** (-4.42)
2022w42	-10.04** (-2.50)	-32.10** (-5.79)
2022w43	-9.49* (-2.33)	-32.81** (-6.04)
$p: \beta_{2022w36} = \beta_{2022w37}$	0.01	0.01
Adj- R^2	0.64	0.80
N	1,797	1,336

Notes: This table shows the estimated θ_w from the following difference-in-differences regression:

$$y_{it} = \lambda_i + \alpha_{m(i),t} + \Gamma X_{it} + \sum_w \theta_w \times 1(t = w) \times Pool_i + \varepsilon_{it}$$

where y_{it} is one of two outcome variables for LDI i managed by $m(i)$ on date t . The outcome variables in the first and second columns are, respectively, the percentage change: (i) in Gilt holdings (par values) and (ii) repo outstanding, both relative to their initial stocks on 9/2/2022. λ_i is an LDI-level fixed effect and $\alpha_{m(i),t}$ is a manager-by-time fixed effect. $1(t = w)$ is an indicator variable for whether week t equals w and $Pool_i$ is an indicator for whether LDI i is a pooled vehicle. X_{it} is a vector of controls for i that contains its lagged debt-to-asset ratio (three polynomials) and repo-to-asset ratio. The row $p: \beta_{2022w36} = \beta_{2022w37}$ shows the p -value from testing the null of no parallel pre-trends. Standard errors are clustered by account and week. Data is weekly.

Bank of England

Appendix to Staff Working Paper No. 1,089

Fire sales of safe assets

Gabor Pinter,⁽¹⁾ Emil Siriwardane⁽²⁾ and Danny Walker⁽³⁾

(1) Bank for International Settlements. Email: gabor.pinter@bis.org

(2) Harvard Business School and NBER. Email: esiriwardane@hbs.edu

(3) Bank of England. Email: daniel.walker@bankofengland.co.uk

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH

Email: enquiries@bankofengland.co.uk

©2024 Bank of England

ISSN 1749-9135 (on-line)

A.1 Additional Results

A.1.1 Survey Estimates of LDI Size

As described in Section 2.3, our construction of the LDI sector's balance sheet relies on: (i) identifying LEIs associated with LDIs; and (ii) cumulating transactions in the gilt market to arrive at a stock of gilt holdings. Because both steps introduce some measurement error, we validate our balance sheets against a regular survey containing information on more than 500 leveraged LDI funds gathered from the largest UK LDI managers. This survey was started after the 2022 crisis to monitor the resilience of the LDI sector. It is administered by the Financial Conduct Authority (FCA) and shared with the Bank of England and the Pensions Regulator (tPR). Here we focus on information dated 10/31/2022 because this is the earliest survey data available, which allows us to compare LDI balance sheet data as close as possible to the crisis in late September to mid October 2022.

In the survey the LDI managers reported leverage (an asset-to-equity ratio) of 1.7 on average for their non-pooled LDI funds, specifically single LDI funds and segregated mandates, and 2.7 on average for their pooled LDI funds. By comparison, we estimate average leverage of 2.2 for non-pooled LDI funds and 2.9 for pooled LDI funds.

For our comparison of total gilt holdings we focus on survey responses from non-pooled LDI funds, as the information gathered on pooled LDI funds in early survey waves was more limited than for non-pooled LDIs. LDI managers reported holding £357 bn of gilts in their non-pooled LDI funds. To our knowledge, this number represents the book value for gilts, marked at their purchase price, rather than their current market value. By comparison, we estimate the book value of the gilt stock at the end of October to be £331 bn or £26 bn less than the survey results.

Overall, the survey responses provide some comfort that our estimates confirm the finding in the survey data that pooled LDI funds were more leveraged than non-pooled LDI funds and that our measures of gilt holdings are relatively accurate, as the survey data and our estimates are within 7% of each other. Our estimates may diverge from the survey results for number of reasons. For one,

we our list of LDIs associated with LDIs may not be fully exhaustive. Second, the survey responses are likely lagged whereas our estimates are based on flows up to 10/28/2022. It is also important to note that the amount by which we underestimate gilt positions is quite small compared to the combined LDI-pension sector balance sheet. For example, suppose the £26 bn of gilts that we miss is fully financed by repo borrowing. Even in this extreme case, the amount of repo borrowing by the LDI sector would still be less than 10% of the combined sector's assets. Consequently, any measurement error in our estimates does not change our basic argument that the pension sector had sufficient assets to easily collateralize the debt held on LDI balance sheets (see Section 4.1)

A.1.2 LDI Trading Behavior Over the Long Run

Figure A1 shows cumulative Gilt purchases by LDIs over a long horizon. The main point of the plot is that LDI selling behavior during September 2022 was clearly abnormal for historical standards, as they rarely were net sellers of Gilts in the preceding four years.

A.1.3 Portfolio Sorts

In Section 3.2.3 we sort Gilts into portfolios based on H_b , the fraction of Gilt b held by LDIs as of 9/2/2022. Figure A2a shows the average H_b in each tercile. The plot indicates that sorting based on H_b generates meaningful variation in exposure of Gilts to LDIs. This variation is central to our identification strategy throughout the paper.

Figure A2b shows the average maturity of Gilts in each tercile as well. The plot clearly reveals that LDIs held longer-duration Gilts. Gilts in the low- H_b tercile had an average maturity of 7.7 years whereas those in the top tercile had an average maturity of 24 years. We are careful to control for these duration differences when identifying the price impact of LDI selling in the main text.

A.1.4 Price Pressure

In Section 3.2.2, we use a difference-in-difference strategy to estimate the size of the fire sale discount caused by LDI selling. Our estimates in that subsection are based on hourly (median) transaction prices. While transactions offer several advantages when it comes to measuring fire sale discounts, they might also create selection issues for our empirical strategy. For example, if LDIs selectively sold certain Gilts, perhaps those for which they received private information about fundamentals, then our estimates of the fire sale discount caused by LDI selling could be biased. To alleviate these types of concerns, Figure A3 shows estimates of regression 5 using hourly price quotes from the Thompson Reuters Eikon database. The plot shows the average weekly fire sale discounts based on price quotes are comparable to those apparent from transactions (Figure 4a).

Figure A4 shows difference-in-difference at the daily level, after interacting the date-hour-maturity bin fixed effect with an additional indicator for whether gilt b is a nominal instrument. The purpose of this fixed effect is to control for any differential exposure between nominal and inflation-linked gilts to fundamental shocks (e.g., the mini-budget). The price pressure estimates are comparable to those found in the main text (Figure 4c). Moreover, the point estimates do not materially change under the counterfactual where the pre-trend continues linearly through the mini-budget (Rambachan and Roth, 2023).

A.1.5 The Behavior of Pooled LDIs

A.1.5.1 Alternative Specification

In Section 4.3, we use the following difference-in-difference regression to estimate the impact of the pooled LDI structure on trading behavior through the gilt crisis:

$$y_{it} = \lambda_i + \alpha_{m(i),d,t} + \Gamma X_{it} + \sum_w \theta_w \times 1(t = w) \times Pool_i + \varepsilon_{it}.$$

One potential issue with this specification is the inclusion of controls X_{it} , which account for lagged capital structure and balance sheet composition. As Caetano et al. (2022) point out, the inclusion of

controls in difference-in-difference setups can confound the causal interpretation of the coefficients θ_w , like for instance if X_{it} is affected by the treatment (in this case, being a pooled fund after the mini-budget).

The intent of the controls is to account for any impact that balance sheet composition may have had on trading behavior. For example, if pooled LDIs had higher leverage entering the crisis, then their selling may be driven by this capital structure choice and not their pooled structure. As a different way to account for this possibility that avoids including covariates, we instead modify the fixed effect $\alpha_{m(i),d,t}$ so that it is based on the intersection of manager identity, weighted-average maturity of assets (binned into quartiles), week, and debt-to-assets as of the week before the mini-budget (also binned into quartiles). This means that θ_w is identified by comparing LDIs within the same manager, with comparable asset durations, and initial balance sheet composition. Figures [A6a](#) and [A6b](#) show estimates of θ_w for gilt and repo flows, respectively, using this alternative fixed effect and no covariates. Reassuringly, the figure continues to show that pooled LDIs sold more gilts and closed more repo after the mini-budget.

A.1.5.2 Sophistication

In Section [4.3](#), we demonstrate that pooled LDIs liquidated gilts and reduced repo positions more aggressively than non-pooled counterparts. This behavior suggests that pooled LDIs' balance sheets were more segmented from their pension owners due to factors like coordination frictions. An alternate theory could be that pension owners of pooled LDIs, typically smaller, were less adept at managing collateral. To test this, we reanalyze the data in Section [4.3](#) using only large pooled LDIs and small non-pooled LDIs. Large pooled LDIs are classified as those in the top tercile by total assets at the start of September, while small non-pooled LDIs are defined similarly. The rationale behind this analysis is that larger pensions likely own the larger pooled LDIs, whereas smaller pensions own smaller non-pooled LDIs. If sophistication, as proxied by size, is driving our results, this sample should mitigate the likelihood of observing pooled LDIs deleveraging more aggressively than non-pooled LDIs.

Figure A7 shows the resulting difference-in-difference estimates for this subsample of LDIs. The coefficients in the plot are based on estimates of regression (4) with one modification: we now define the fixed effect $\alpha_{m(i),d,t}$ so that is based on the intersection of LDI manager $m(i)$, two bins of the initial duration of i 's gilt holdings, and time. We use two bins for duration instead of four since the sample is reduced relative to our baseline analysis. The main takeaway from the figure is that large pooled LDIs still sold gilts and reduced repo more than small non-pooled LDIs. This finding cuts against the idea that pension sophistication is the only reason why pooled LDIs seemingly struggled to recapitalize.

A.1.6 Swap Spreads

In Section 5, we used our difference-in-difference estimates of the impact of LDI selling to develop a sense of how much their activity affected the level of the yields. A complimentary way to answer this question is to study the joint evolution of gilt yields and OIS swap rates. The basic idea is that gilts were more impacted by LDI selling whereas OIS swaps were not. This is of course a very strong assumption. To the extent that the two markets share a common set of arbitragers, any selling pressure in the gilt market should spillover into the OIS market (Vayanos and Vila, 2021). Thus, any estimates of how much LDI selling impacted the level of yields based on such a comparison should be viewed as a lower bound on the true effect.

With those caveats in mind, Figure A9 plots the spread between 25-year gilts and OIS swap rates (left) and the 25-year gilt yields (right) from September 2022 onward. Under the (strong) assumption that OIS rates were unaffected by LDI selling, the gilt-OIS spread captures the impact of LDI selling. In the window after the mini-budget but before the BOE first intervened, the gilt-OIS spread and gilt yields increased by 13 bps and 108 bps, respectively. This implies that LDI selling was responsible for at least 12% of the price drop during this period. From September 29 to October 14, when the fire sale reopened (see Section 3.2), the gilt-OIS spread and gilt yields rose 40 bps and 100 bps, respectively, implying that LDI selling was responsible for at least 40% of the price drop during this period. Again, the actual effect of LDI selling is likely much higher given

LDI selling in the gilt market should impact OIS rates if there is common arbitrage across the two markets. Moreover, a regression of the 25-year gilt yield on the associated gilt-OIS spread during this period yields an R^2 of 52%, reinforcing the notion that much of the gilt variation during this time was driven by LDI selling.

A.2 UK Life Insurers

In this section, we compare and contrast the balance sheet composition of UK life insurers to the LDI-Pension sector.

Asset Mix and Duration To start, we analyze the portfolio composition of life insurers as of the end of 2022Q2. Our entire analysis of insurers is based on the Bank of England’s aggregate insurance quarterly report.¹ At the end of 2022Q2, life insurers held the majority (57%) of their assets in fixed-income instruments, with roughly 15% of their assets held in government bonds, 30% held in corporate bonds, and 12% held in mortgages and loans. For comparison, Figure 7 of the main text indicates that the LDI-pension sector held about 65% of its assets in fixed income securities going into the gilt-market crisis. Moreover, virtually all bond holdings of life insurers were investment-grade, with 70% rated A or higher. In terms of duration, roughly 60% of bonds held by life insurers had a maturity of over 10 years and over 30% had a maturity over 30 years. Together, the plots clearly indicate that the assets held by UK life insurers were heavily exposed to interest rate risk entering the crisis.

Net Financial Leverage Next, we calculate the net amount of financial leverage (i.e., excluding policy-related liabilities) on life insurer balance sheets. According to the Bank of England’s regulatory data, as of 2022Q2, life insurers held about £95 bn of derivative liabilities and £57 bn of other liabilities (e.g., including repo), totaling £152 bn of what we label as financial debt. Because they held £82 bn of cash and deposits, life insurers therefore had roughly £70 bn of net

¹Available [here](#).

financial debt on their balance sheets. As with the pension-LDI sector, this financial debt was quite small (~3%) compared to insurers' total assets.² By comparison, the net financial debt of the LDI-pension sector represented roughly 5% of their total assets entering the crisis (Section 4).

Solvency and Liquidity Finally, we analyze the evolution of liquidity and solvency in the insurance sector alongside the LDI-pension sector. As many others have noted (e.g., Chen and Kemp, 2023; Alfaro et al., 2023), the solvency and liquidity needs of the LDI-pension sector diverged as interest rates increased in 2022, especially after the mini-budget. On one hand, the solvency of the LDI-pension sector improved as rates rose, reaching a two-decade high (see Figure A5), because the duration of pension obligations exceeded that of their assets. On the other hand, the sector's liquidity needs increased with interest rates, due to margin calls stemming from financial debt used to finance long-duration assets (e.g., via LDIs). Alfaro et al. (2023) label this negative correlation between solvency and liquidity needs as “liquidity after solvency hedging” or “LASH” risk.

Figure A10 reveals that the life insurance sector encountered a very similar situation. Panel (a) of this figure displays the solvency of the insurance sector, as measured by the Solvency Capital Requirement (SCR) ratio, from Q1 2021 through Q3 2022, alongside the solvency ratio of the LDI-pension sector. The solvency ratios of both sectors closely align, rising in response to the increase in global interest rates throughout 2022. The solvency of both sectors reached a peak in Q3 2022, coinciding with the onset of the gilt market crisis.

Similarly, Figure A10b shows that, like the LDI-Pension sector, the liquidity needs of the life insurance sector increased as interest rates rose. The figure uses variation margin payments on derivative positions—equivalently, changes in the mark-to-market value of liabilities—to proxy for the sector's liquidity needs. This proxy likely understates the true liquidity needs of life insurers, as it does not account for any repo financing used for purchasing gilts or corporate bonds. In the first half of 2022, life insurers made approximately £30 billion in variation margin payments on derivatives, with these liquidity needs surging to nearly £70 billion in Q3 2022 alone.³

²This number is based on all assets of life insurers, including unit-linked assets. Relative to the pure investment portfolio of life insurers, net financial debt was still somewhat small (~10%).

³Alfaro et al. (2023) estimate the anticipated liquidity needs of the insurance sector in response to a 100 basis point

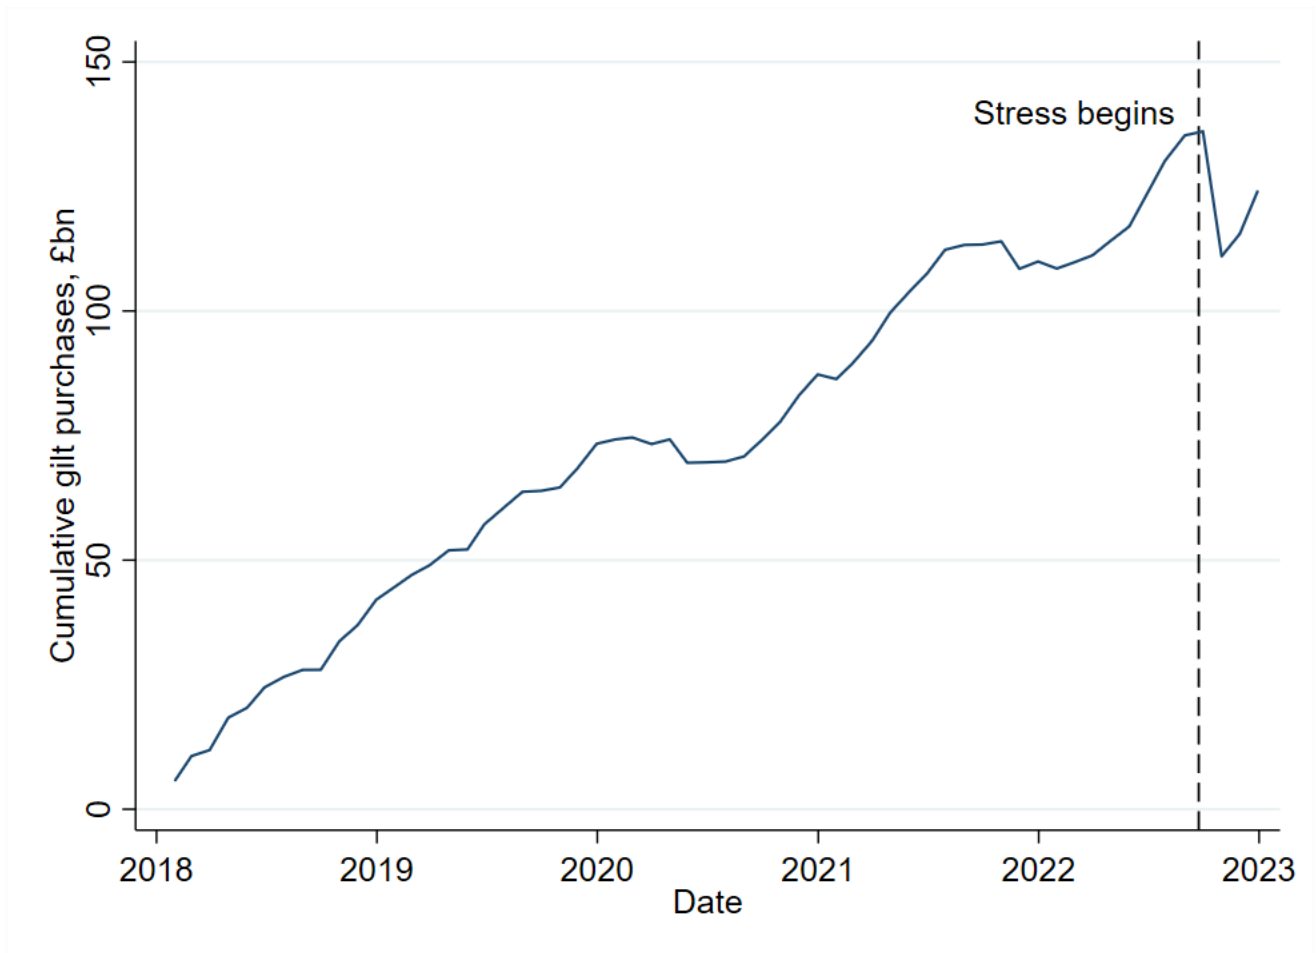
Takeaways Overall, the analysis in this subsection underscore that UK life insurers and the LDI-Pension sector had similar balance sheet compositions as they entered the gilt-market crisis, a somewhat predictable outcome given their analogous economic functions. However, it is noteworthy that life insurers did not experience nearly the same level of stress as the LDI-pension sector following the mini-budget. For example, Figure A11 indicates that life insurers sold almost no gilts, decreased their swap exposure by less than £4 bn, and slightly increased their repo borrowing. We attribute this divergence in behaviors to the fact that life insurers held their financial leverage directly on their balance sheets, in contrast to pensions and their LDI investments.

increase in the level of gilt yields. Their estimates provide a lower bound on the true liquidity needs of the insurance sector because they reflect only gilt positions, which are half the size of corporate bond positions.

References

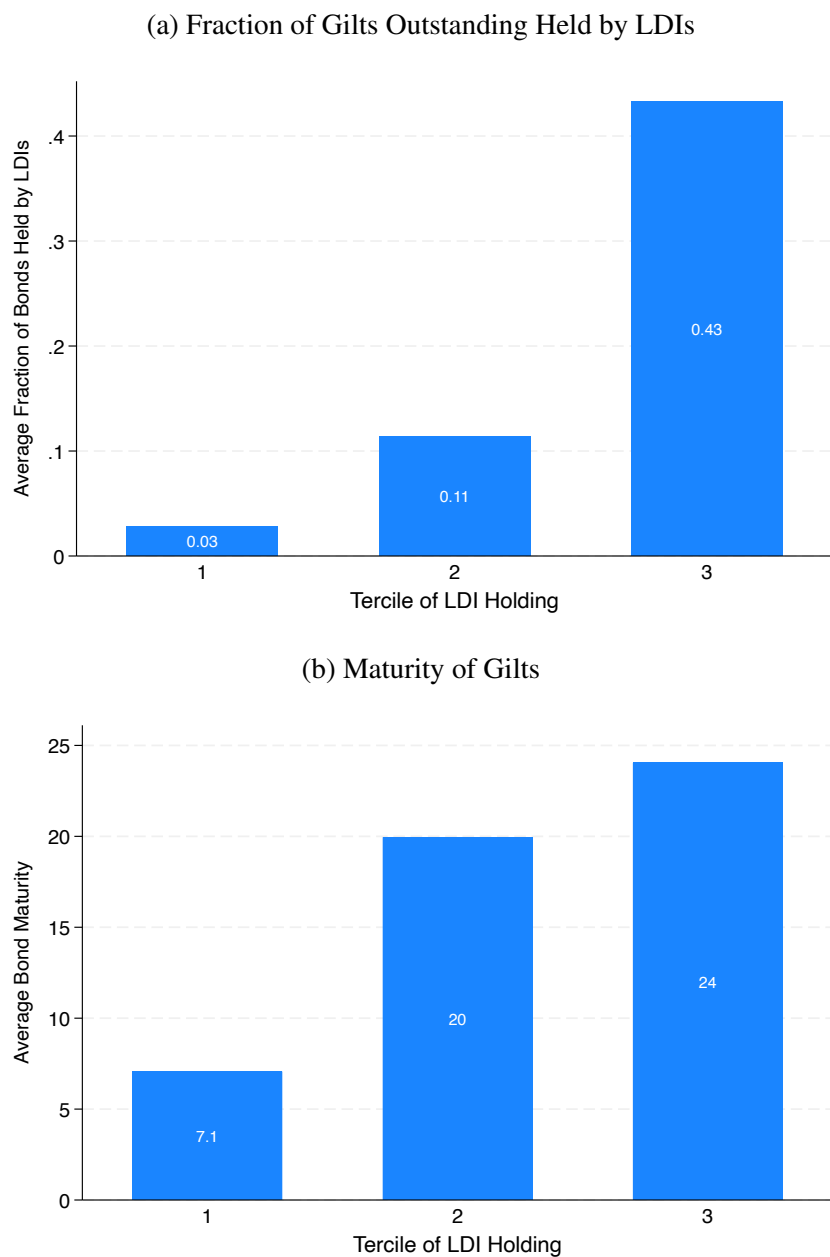
- Alfaro, L., S. Bahaj, R. Czech, J. Hazell, and I. Neamt (2023). Lash risk and interest rates. Technical report, Bank of England.
- Caetano, C., B. Callaway, S. Payne, and H. S. Rodrigues (2022). Difference in differences with time-varying covariates. *arXiv preprint arXiv:2202.02903*.
- Chen, R. and E. Kemp (2023). Putting out the nbfire: Lessons from the uk’s liability-driven investment (ldi) crisis. Technical report, International Monetary Fund.
- Rambachan, A. and J. Roth (2023, 02). A More Credible Approach to Parallel Trends. *The Review of Economic Studies* 90(5), 2555–2591.
- Vayanos, D. and J.-L. Vila (2021). A preferred-habitat model of the term structure of interest rates. *Econometrica* 89(1), 77–112.

Figure A1: Long-Run Trading Behavior of LDIs




Notes: This figure shows cumulative gilt purchases by LDIs since 2018 based on transactions in the MIFID II database.

Figure A2: Fraction of Gilts Held and Maturity by LDI Holding Intensity



Notes: To construct this figure, we sort Gilts into terciles based on H_b , the fraction of Gilt b held by LDIs as of 9/2/2022. Panel (a) then plots the average H_b in each tercile and panel (b) plots the average Gilt maturity in each tercile.

Figure A3: Price Pressure from LDIs using Price Quotes

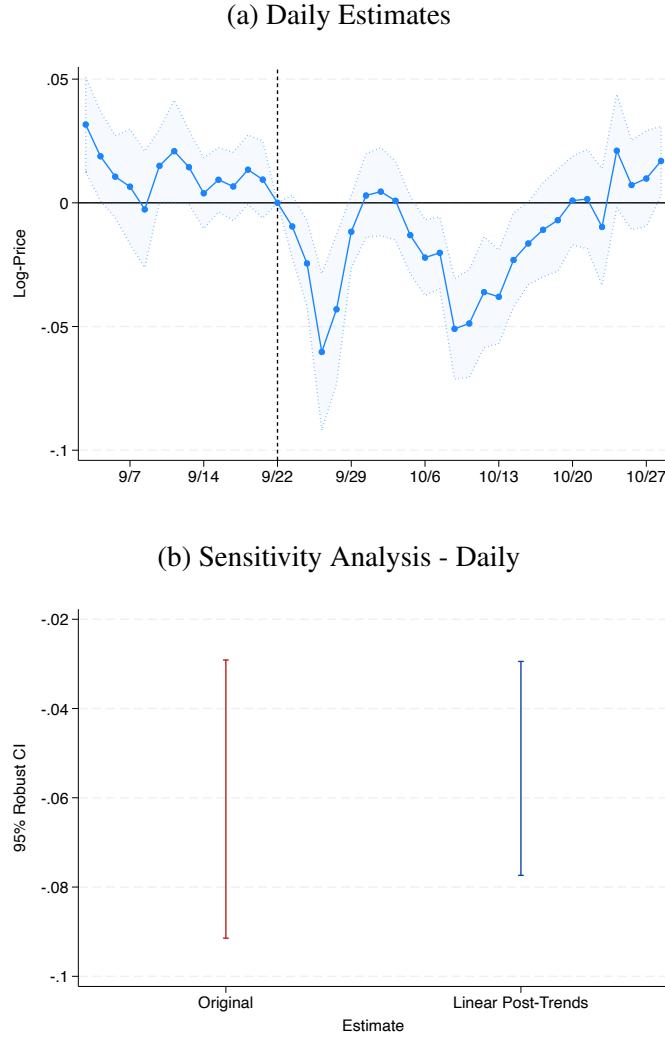


Notes: This figure shows the estimated β 's from the following difference-in-differences regression:

$$\log(p_{b,h,t}) = \gamma_b + \alpha_{h,t,m} + \sum_w \beta_w \times 1(t \in w) \times D_b + \varepsilon_{b,t},$$

where $p_{b,h,t}$ is the price quote for Gilt b in hour h of day t from the Thompson Reuters Eikon database. γ_b is a Gilt fixed effect and $\alpha_{h,t,m}$ is an date-by-hour-by-maturity fixed effect. $1(t \in w)$ is an indicator variable for whether date t is in week w and D_b is an indicator for whether Gilt b was held heavily by LDIs of September 2, 2022. D_b is defined based on whether LDI holdings of b were above median. Standard errors in the plot are double-clustered by Gilt b and day-by-hour.

Figure A4: Daily Price Pressure from LDIs (Robustness)

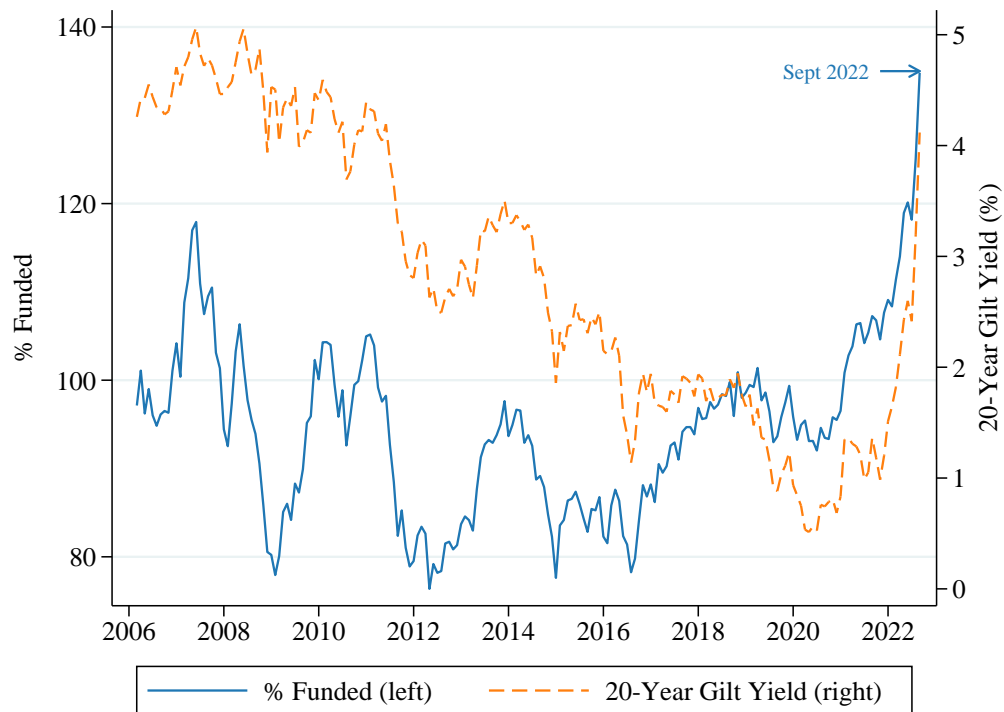


Notes: Panel (a) of this figure shows the estimated β 's from the following difference-in-differences regression:

$$\log(p_{b,h,t}) = \gamma_b + \alpha_{h,t,m,n} + \sum_s \beta_w \times 1(t \in s) \times D_b + \varepsilon_{b,t},$$

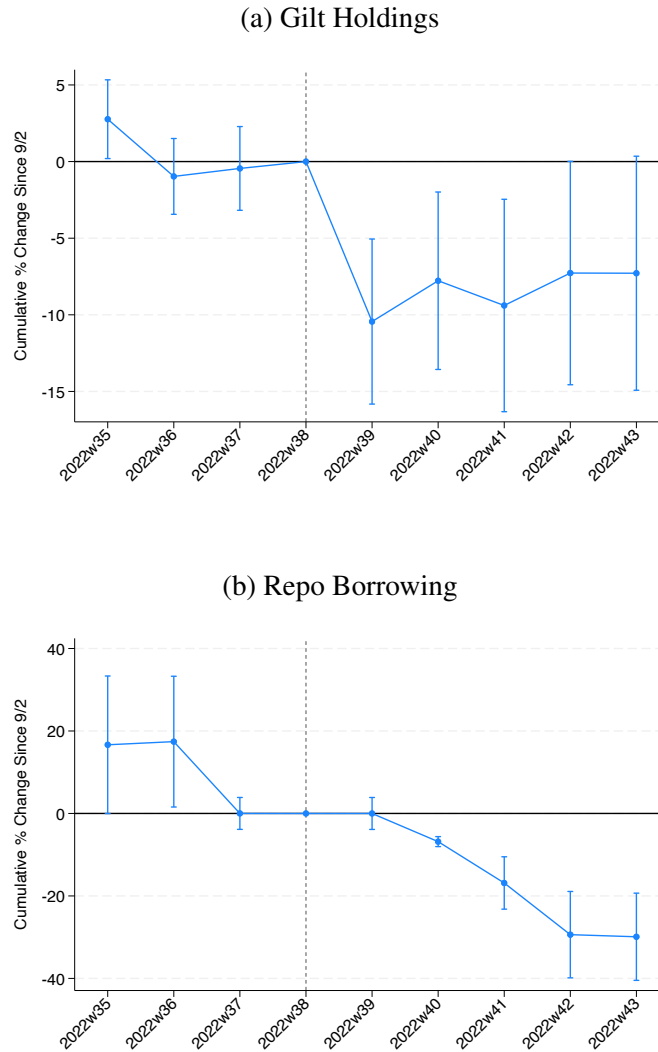
where $p_{b,h,t}$ is the median-transaction price for Gilt b in hour h of day t . γ_b is a Gilt fixed effect and $\alpha_{h,t,m,n}$ is a date-by-hour-by-maturity-by-nominal fixed effect. It is based on the intersection of indicators for date t , hour h , the maturity-quintile of the gilt, and whether the gilt is a nominal instrument. $1(t \in s)$ is an indicator variable for whether date t equals s and D_b is an indicator for whether Gilt b was held heavily by LDIs of September 2, 2022. D_b is defined based on whether LDI holdings of b are in the top tercile. Following Rambachan and Roth (2023), Panel (b) shows how most-negative estimate of β_w would change if the trend between Gilts with $D_b = 1$ and $D_b = 0$ had continued linearly through 2022w38, the week of the mini-budget. Standard errors in the plot are double-clustered by Gilt b and day-by-hour.

Figure A5: Funding Ratio for U.K. Corporate Pensions



Notes: This plot shows how the monthly funding ratio (assets-to-pension liabilities) of U.K. corporate pensions evolved in conjunction with 20-year nominal Gilt yields. Data is monthly and taken from the PPF 7800 index, published by the U.K. Pension Protection Fund.

Figure A6: Trading Behavior of Pooled vs Single LDIs (Robustness)

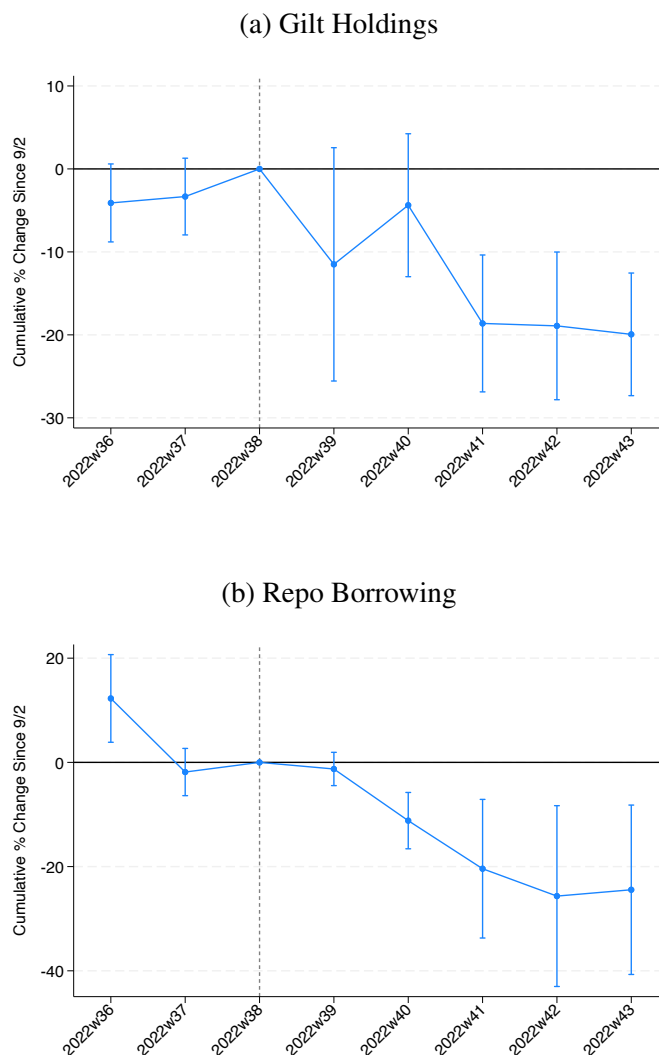


Notes: This figure shows estimates of the following difference-in-differences regression:

$$y_{it} = \lambda_i + \alpha_{m(i),d,l,t} + \sum_w \theta_w \times 1(t = w) \times Pool_i + \varepsilon_{it}$$

where y_{it} is one of two outcome variables for LDI i managed by $m(i)$ on date t . The first outcome (Panel a) is the percentage change in Gilt holdings (par values) and the second (Panel b) is the percentage change in repo outstanding, both relative to their initial stocks on 9/2/2022. λ_i is an LDI-level fixed effect and $\alpha_{m(i),d,l,t}$ is a manager-by-duration-by-leverage-by-time fixed effect, where duration is defined based on quartiles of the maturity of i 's gilt holdings at the beginning of September and leverage is based on quartiles of i 's debt-to-asset ratio in the week before the crisis. $1(t = w)$ is an indicator variable for whether week t equals w and $Pool_i$ is an indicator for whether LDI i is a pooled vehicle.. Standard errors are clustered by account and week. Data is weekly.

Figure A7: Trading Behavior of Large Pooled vs Small Non-Pooled LDIs

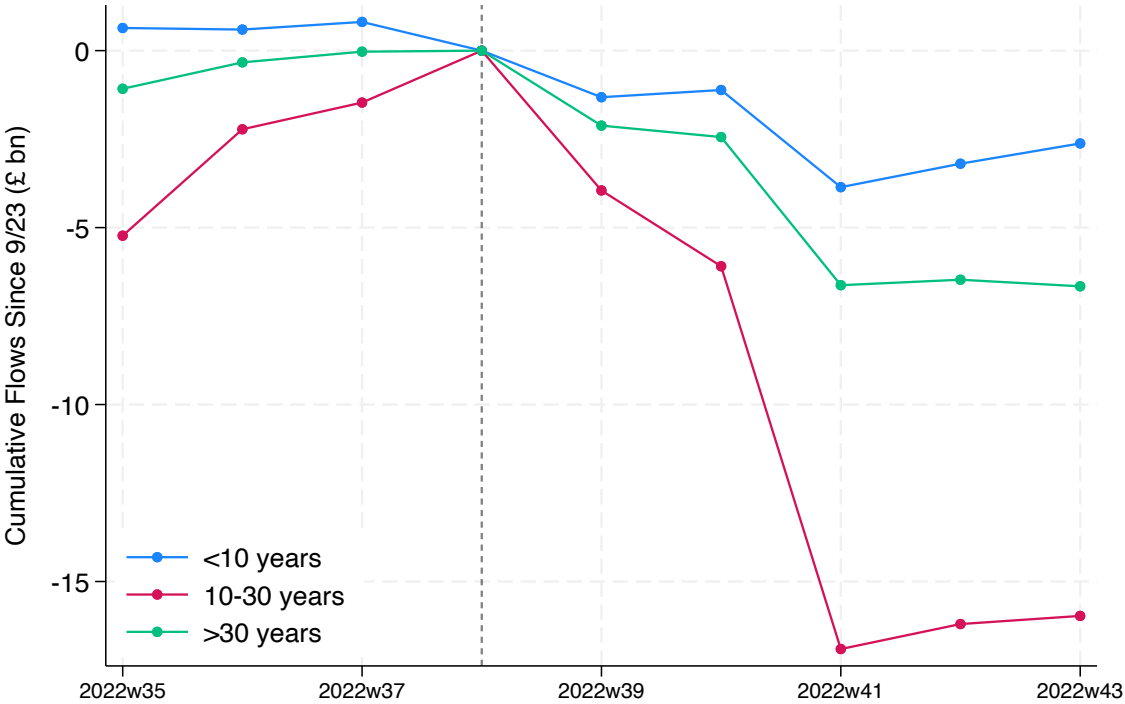


Notes: This figure shows estimates of the following difference-in-differences regression:

$$y_{it} = \lambda_i + \alpha_{m(i),d,t} + \Gamma X_{it} + \sum_w \theta_w \times 1(t = w) \times Pool_i + \varepsilon_{it}$$

where y_{it} is one of two outcome variables for LDI i managed by $m(i)$ on date t . The first outcome (Panel a) is the percentage change in Gilt holdings (par values) and the second (Panel b) is the percentage change in repo outstanding, both relative to their initial stocks on 9/2/2022. λ_i is an LDI-level fixed effect and $\alpha_{m(i),d,t}$ is a manager-by-duration-by-time fixed effect, where duration is defined based on whether the maturity of i 's gilt holdings at the beginning of September is above or below median. $1(t = w)$ is an indicator variable for whether week t equals w and $Pool_i$ is an indicator for whether LDI i is a pooled vehicle. X_{it} is a vector of controls for i that contains its lagged debt-to-asset ratio (three polynomials) and repo-to-asset ratio. Standard errors are clustered by account and week. Data is weekly. The regression is run for the sample of large pooled LDIs and small non-pooled LDIs. Within each group, size is defined based on terciles of assets measured at the beginning of September.

Figure A8: LDI Gilt Flows by Maturity



Notes: This figure shows gilt flows by LDIs from September through October 2022. Gilt flows are based on the change in the par amount of bonds held. All flows are indexed to 9/23 (vertical dashed line), the day of the mini-budget.

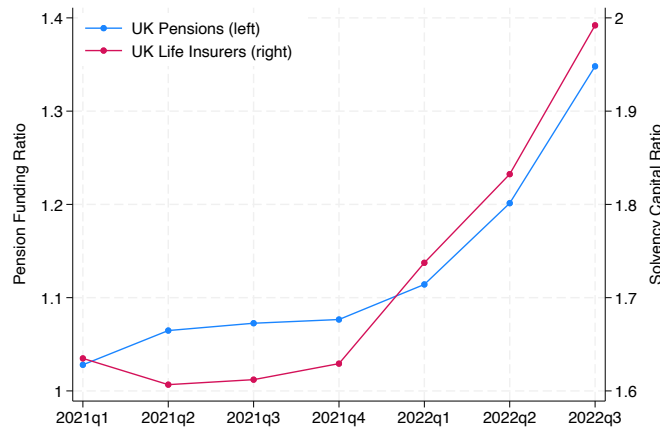
Figure A9: Swap Spreads vs the Level of Gilt Yields




Notes: This figure shows spreads between 25-year gilts and 25-year OIS swap rates (left axis) alongside the cumulative average return of all gilts from September 22, 2022 onward (right axis). Cumulative returns on gilts are based on end-of-day quotes from the Thompson Reuters Eikon database. Gilt yields and OIS rates are from the Bank of England's website.

Figure A10: Solvency and Liquidity of UK Life Insurers

(a) Solvency

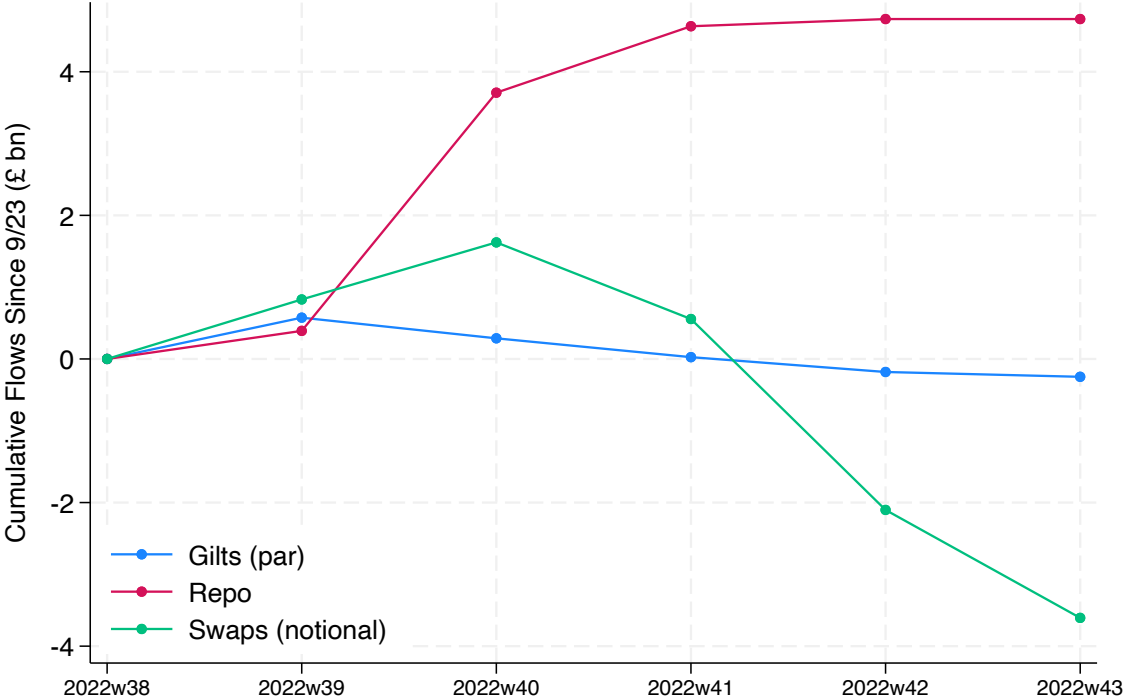


(b) Liquidity Needs



Notes: Panel (a) of this plot shows solvency ratios for the UK pension and life insurance sectors. Pension solvency is measured as asset-to-liabilities and insurance solvency is measured using the Solvency Capital Requirement (SCR) coverage ratio. Panel (b) shows total quarterly variation margin payments made by insurers on derivatives positions, as measured by the change in the mark-to-market value of derivative liabilities. Pension data is from the PPF 7800 index and insurance data is from the Bank of England's aggregate insurance quarterly report.

Figure A11: Trading Behavior of Insurers



Notes: This figure shows the trading behavior of UK insurers in the five weeks following the mini-budget on 9/23. Gilt flows are based on the change in the par amount of bonds held, swap flows are based on the change in net notional receiving fixed OIS swaps and paying-fixed inflation swaps, and repo flows are based on the change in outstanding repo borrowing. All flows are indexed to 9/23.