Asset Purchase Rules: How QE Transformed the Bond Market

Valentin Haddad, Alan Moreira, and Tyler Muir*

January 8, 2025

Abstract

We argue that quantitative easing (QE) and tightening policies constitute a dynamic state-contingent plan instead of a succession of independent interventions. This view changes the main reason QE is effective by adding an insurance channel to the static effect of absorbing bond supply in a given period. QE purchases occur in bad economic states (e.g., 2008-2009 or 2020) when the supply of government debt increases. Increasing long-term bond prices in bad economic states increases their safety, driving up their value and thus lowering ex-ante yields. We estimate that this insurance channel alone lowers long-term bond yields by 75-100 bps. This channel explains the prevalence of low long-term yields, low term premia, and low yield volatility since the introduction of QE, despite the sharp increase in net government debt supply. Consistent with a state-contingent channel, implied volatilities of long-duration risk-free securities fall substantially on QE announcements, even for options with maturities out to 10 years. We calibrate a policy rule for asset purchases to their historical path and include it in a quantitative term structure model. In the model, state-contingent QE offsets term premia fluctuations in long-term bonds. The insurance effect from this channel lowers long-term Treasury yields by 75bps ex-ante, which explains about 75% of the total effect of QE on yields. The calibrated model matches both broad patterns in bond yields and the response to QE announcements.

*Haddad: UCLA & NBER, valentin.haddad@anderson.ucla.edu; Moreira: University of Rochester & NBER, alan.moreira@simon.rochester.edu; Muir: UCLA & NBER, tyler.muir@anderson.ucla.edu. We thank participants at Berkeley, UCLA, Wharton, Chicago, NYU, Harvard, Columbia, Stamford, University of Rochester, Federal Reserve Board, BIS the NBER Summer Institute, UGA Finance conference, Bocconi Asset Pricing conference, and Viral Acharya, Ryan Banerjee, John Cochrane Will Diamond (discussant), Darrel Duffie, Itamar Drechsler, Andrea Eisfeldt, Christian Jensen(discussant), Anil Kashyap, Hanno Lustig, Narayana Kocherlakota, Arvind Krishnamurthy, Francis Longstaff, Thomas Mertens, Mehrdad Samadi (discussant), and Annette Vissing-Jorgensen, for helpful conversations.

1. Introduction

What happens when policymakers systematically purchase long-term assets in downturns with elevated term premia and market participants understand this form of policy rule? Natural examples of such states of the world are 2008-2009 or 2020 when the economy is in turmoil and the supply of government debt sharply increases. The mechanisms and implications of QE when it is part of a state-contingent plan differ sharply from those in the common approach considering QE as surprise one-off interventions. By leaning against the term premium, state-dependent interventions increase the safety of long-term bonds ex-ante. This additional safety leads to unconditionally lower long rates for two main reasons. First, countercyclical interventions have substantial ex-ante impact because they provide support in highly valued states of the world. Second, because the policy rule leads to relatively more stable bond prices, it crowds in investors who seek safety as well as arbitrageurs who absorb temporary supply imbalances, further lowering yields.

In this paper, we flesh out this argument, and provide a framework to quantify the impact of various forms of state-contingent asset purchase policies. We show that a simple policy rule leaning against the term premium accounts empirically for 1) the realized pattern of asset purchases, 2) persistent changes in the behavior of the yield curve since the initial introduction of QE, and 3) market responses right at announcements and updates of QE policies, including responses for long-dated options on long-term bonds. In particular, we estimate that the policy rule lowers 10-year Treasury yields by around 100 basis points (bps). Importantly, we also show that empirical analyses that only view QE as one-off interventions would miss this sizable impact.

First, the realized use of quantitative easing and tightening does adjust dynamically with the state of the economy – central banks systematically use asset purchases during crises or downturns. After a central bank first uses QE, like the Federal Reserve in 2008, its balance sheet seldom shrinks back to the old status quo of tiny asset holdings. Instead,

long-term asset positions remain large, on the order several tens of percent of GDP, contracting and expanding with economic conditions. For example, the Fed purchased over $1 trillion in long-term Treasuries in March and April of 2020 at a time when the economy was in recession, Treasuries markets faced liquidity issues, and Treasury bond issuance from government spending spiked. More broadly, the Fed's balance sheet tends to expand in periods where the supply of Treasuries increases and these are periods where term premia would normally rise (Greenwood and Vayanos, 2014).

If market participants systematically expect the central bank to lean against the term premium by absorbing bonds, a natural consequence should be more stability in long-term interest rates. Motivated by this prediction, we revisit the evidence of Greenwood and Vayanos (2014), who document that higher bond supply increases the slope of the yield curve by raising the term premium in a sample from 1952-2007. This relationship breaks down in the post QE period: if anything, supply increases become associated with a lower slope and thus term premia appear to move far less to supply shocks.

This ex-post support naturally lowers yields ex-ante. We document that the spread between long-term and short term interest rates is indeed persistently low in the post-QE era. This is the case despite a large increase in the supply of long-duration Treasuries over this period. Given the increase in supply in the last 15 years, the earlier relationship between slope and supply would predict long-term bond yields about 150 bps above current levels. Accounting for Fed purchases in a static way — i.e. netting them out from the supply of Treasuries — only lowers this gap by around 40 bps. The conclusion of abnormally low term premia by historical standards appears robust, as it holds across a variety of benchmarks or specifications.

Still, a number of factors other than QE policies can drive low-frequency shifts in the behavior of long-term yields. We exploit the staggered introduction of QE policies across many countries, allowing us to include time-fixed effects that control for potential global trends in long-term rates. This setting also helps alleviate concerns of other changes spe-

cific to the US that happen around the time QE is introduced. We find that the slope of the yield curve is, on average, about 80 bps lower after central bank asset purchases are introduced. Further, the coefficient on bond supply (i.e., debt to GDP) is strongly positive before the introduction of QE — higher bond supply is associated with a steeper yield curve — but drops to near zero or negative afterwards. Another potential concern is that the introduction of QE often coincides with the short term interest rate hitting the zero lower bound (ZLB). We find that replacing the short-term rate by its shadow value predicted by the Taylor rule has no effect on our main results. Markets also price fast departures from the ZLB and small probabilities of getting back to it at long horizons, making the ZLB an unlikely driver of the large change in behavior of long-term rates that we document.

Another invaluable window into how QE policies impact markets are short intervals around the announcement of these policies. This zooms in on the moment news is communicated about the policy, and leads to sharper identification. However, in the potential presence of a policy rule, particular care has to be given in interpreting this evidence.

Consider first the initial announcements of QE interventions. On those days, all yields drop, with the effect particularly pronounced at longer horizons and peaking around the 10-year maturity. The cumulative decline in long-term yields on just these QE announcement days is in the ballpark of 100-140 bps which is enough to explain the broad time-series evidence above that long-term yields appear too low in the post QE era. Importantly, this response might be driven by the asset purchases announced at that point, but also by market participants updating their view on the persistent use of QE as a policy tool. One way to assess this latter aspect is to use option prices, which reveal perceptions of future price dynamics. Consistent with the view that the policy stabilizes long-term bonds, we observe a stark decline in the implied volatility of long-term Treasuries. This is true using options on 10 year Treasury futures, though the maturity of these options is fairly short in the span of 1-12 months. Using data on swaptions, which are liquid

at much longer exercise maturities out to 5-10 years, we find a similar result. Implied volatilities on swaptions with a 10 year exercise maturity on a 10 year swap decline by over 30% on these days. Thus, the introduction of QE appears to dramatically reduce the volatility of long-duration assets even at very long horizons, supporting the idea that that these policies have an important persistent and state-contingent component.

Later announcements, when market participants are more likely to have internalized the policy rule, exhibit a substantially distinct pattern. Responses to these announcements are much more muted and a roughly near zero on average. One reason for these weaker reactions is that the interventions are more largely anticipated.

But the theory points to another explanation for why, even focusing on the surprise component of the announcements, responses are weaker. Because the presence of the policy rule makes bonds less risky, it crowds in more willingness to absorb shocks by arbitrageurs. This additional elasticity reduces the direct impact of additional purchases, even if they are a surprise relative to the perceived policy rule. Said otherwise, there is an event-study paradox: the fact that the response to QE *policy surprises* is muted reveals that the QE *policy rule* has strongly impacted the bond market.

To further isolate this mechanism from the mechanical role of expectations, we focus on another sources of shocks, where the surprises are better measured: financing news by the US Treasury contained in quarterly refunding announcements. We use the yield response over these announcements relative to the actual amount of issuance over the following quarter(s). Long-term yields increase over the announcement periods when future issuance is high, but this relationship substantially weakens in the post QE era. Thus, here again we find that markets are more willing to absorb an additional unit of treasury supply in the post-QE regime.

Finally, we study the pass through of these effects to other long-term rates which helps gauge the broader economic impact of QE policy. We find similarly low rates for mortgages in the post QE era and, to a slightly lesser extent, long-maturity corporate

bonds. The somewhat weaker effects in corporate bonds translate to relatively high convenience yields of Treasuries vs corporate bonds despite increased Treasury supply (Krishnamurthy and Vissing-Jorgensen, 2012).

To assess the quantitative consistency of our estimates, we use a model of the yield curve, in the style of Greenwood and Vayanos (2014) and Vayanos and Vila (2021). In this model, changes in the supply of Treasury or QE interventions affect bond yields because they shift liabilities from the hands of inelastic investors (the broad population) to risk averse arbitrageurs (active market participants), and vice versa. We first calibrate the model on the pre-QE behavior of the yield curve and Treasury supply, matching the evidence in Greenwood and Vayanos (2014) that supply shocks affect term premia. Then, we feed in a QE policy rule consistent with the empirical relation of between asset purchases and the supply of government debt.

In line with our empirical estimates, the model predicts that the presence of the Fed in the Treasury market as a conditional buyer (and seller) has reduced 10 year Treasury yields by around 115 basis points. Of this amount, only 40 basis points can be attributed to the direct effect of what the Fed has purchased to date. The other 75 basis points come from state-contingent expectations of future purchases and the insurance effect they provide. The model also produces a realistic decline in the sensitivity of bond prices to Treasury supply as well as a drop in volatility of long-term yields.

The model also allows us to quantify the impact of alternative QE policy rules. For example, synchronizing QE with conventional monetary policy in following the short term interest rate is ineffective in lowering yields because it misses the long-term effect of stabilizing the term premium.

Taken together, this analysis demonstrates how to get an empirical handle on the quantitative impact of asset purchase policy rules, and highlight the relevance of such rules. These results demonstrate the need to revisit evaluations of the propagation of asset purchases to the real economy. The strong impact we find on Treasuries is relevant

for fiscal sustainability, but one might wonder about transmission to other asset classes. By comparing Treasuries to safe corporate bonds and mortgages, we find strong evidence of transmission, albeit imperfect, to these other assets.

While our main goal is to establish the potency of state-contingent QE, our findings have implications for work on optimal QE policy. Gertler and Karadi (2011) and Sims et al. (2023) study optimal QE policy in New Keynsian models with financial frictions and we revisit their conclusions in light of our evidence. Our finding that QE policy targets variations in term premia, particularly those related to supply and demand imbalances, accords fairly closely to the optimal policy in their work. Similarly Duffie (2023) argues for asset purchases to alleviate severe disruptions and liquidity issues in Treasury markets, emphasizing that investors highly value the safety in price stability of Treasuries even in rare states of the world. Our analysis suggests that intervening in these states indeed has significant effects in lowering unconditional yields and boosting demand for Treasuries as safe assets.

The literature on why QE works has emphasized financial frictions of various sorts. Curdia and Woodford (2011), Gertler and Karadi (2011), and Gertler and Karadi (2018) emphasize the importance of frictions on the balance sheets of banks and how they introduce substitution between treasuries and other private assets. Vayanos and Vila (2021) assumes arbitrageurs are risk-averse, which is also implicitly a type of financial friction. A commonality among these papers is that QE works if it has effects on asset prices.[1]

Multiple papers have carefully investigated how these asset market effects pass-trough to the real economy. Di Maggio et al. (2020) shows increases in MBS bonds pass-trough to mortgage rates, which in turn pass-through to consumption through households refinancing decisions. Foley-Fisher et al. (2016) shows how QE impacted yields of corporate bonds which induced more issuance of long-term bonds, and an increase in investment and employment by corporations. Rodnyansky and Darmouni (2017) and Luck and Zim-

[1]See Bianchi and Bigio (2022) for a liquidity view of QE.

mermann (2020) find that banks that benefited more from the price increases induced by QE end up lending more which translated into more employment and investment.

There is of course a very large literature studying announcement effects of QE-type programs, including Joyce and Tong (2012), Vissing-Jorgensen and Krishnamurthy (2011), Meaning and Zhu (2011), Gagnon et al. (2018), Vissing-Jorgensen (2021), and Krishnamurthy et al. (2018). Relative to this work, we emphasize interpreting the evidence through a state-contingent rule for QE.

2. The yield curve after QE

2.1 Central bank balance sheets and asset purchase patterns

We begin by plotting the size of central bank balance sheets as a fraction of GDP for the US, UK, Euro area, Japan, and Canada in Figure 1.[2] We highlight two patterns. First, central bank balance sheets have grown to increasingly large levels, ranging between 20-90% of GDP, with an average near 40%.[3] This is primarily the result of asset purchase programs, e.g., QE in the US, over the past 15 years. Second, and importantly for our story, central bank balance sheets expand through new asset purchases during bad economic times. For the United States (blue dashed line) we see pronounced increases after the financial crisis of 2008-2009 and again in the COVID episode of 2020. We can see an increase for the Euro area around 2012 surrounding the European sovereign debt crisis. Thus, purchases by and large increase during bad states of the world.

If government bond purchases in a given period raise bond prices, then this dynamic policy will make bonds safer to investors ex-ante because prices will be relatively higher in bad economic states where purchases are concentrated. This improves the insurance

[2] All data are from Fred.

[3] Scaling by the stock of public debt, rather than GDP, produces an average closer to 30-35% but still shows an upward trend.

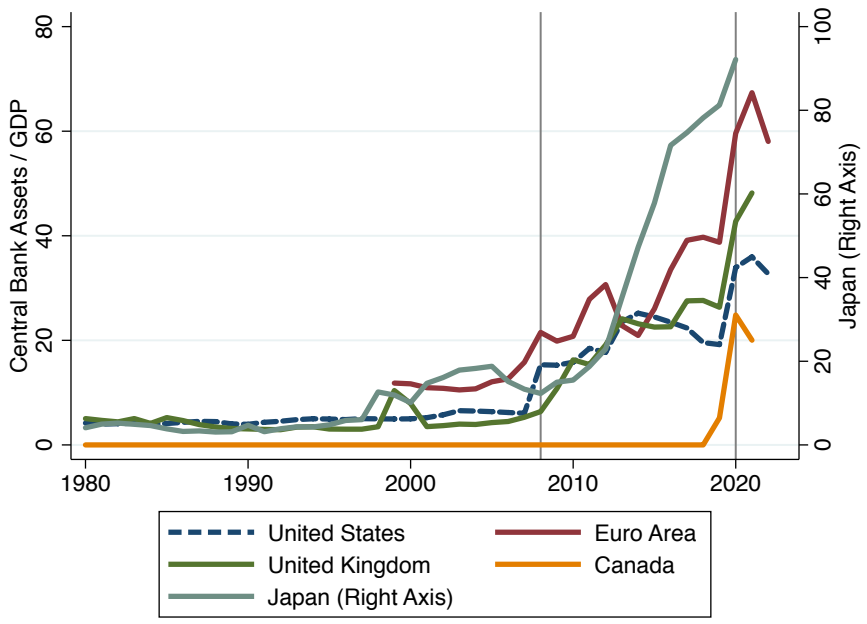


Figure 1: Central Bank Balance Sheets.
All data series from Fred. Gray dashed lines indicate the introduction of QE in the US in 2008 and the COVID period which say an expansion of central bank purchases for all countries.

properties of long-term bonds and, in equilibrium, should reduce yields. Thus, forward-looking measures of risk, and also bond yields, should be lower if investors see these policies as recurring. Further, the dynamics of bond prices (or yields) will also change. While the broad patterns of central bank balance sheet behavior helps provide intuition for these mechanisms, we now turn to direct empirical evidence on the behavior of yields.

2.2 Long-term Treasury Yields are low in the Post-QE era

We show that the slope of the yield curve – long-term Treasury yields relative to short-term interest rates – appears "too low" since QE was introduced by around 100-150 basis points (bps). This result holds both in US data as well as internationally using staggered introduction of QE.

2.2.1 The Slope of the Yield Curve: US Evidence

We show long-term Treasury yields appear low relative to the level of short term interest rates and the supply of duration risk. Our regressions follow Greenwood and Vayanos (2014) who regress the slope of the yield curve on maturity-weighted debt to GDP (MWDGDP) as a measure of the total supply of duration risk. Specifically, this measure computes all future cash flows of the outstanding publicly available Treasury securities (both principal and coupons), weights them by maturity, and normalizes the sum by nominal GDP. We reconstruct and extend the maturity-weighted debt to GDP measure through the end of our sample. We take the long-term Treasury yield, which has a maturity of roughly 20 years, and subtract off the 3-month Treasury-bill yield as our measure of slope.[4]

Table 1 regresses the slope of the yield curve on debt supply (MWDGDP). We control for the short-rate as well as the unemployment rate which picks up the state of the macroeconomy, as the slope of the yield curve is well-known to depend on macroeconomic conditions. In the pre-2008 data, higher debt supply is associated with relatively higher yields on long-term Treasuries (Greenwood and Vayanos, 2014). A 10% increase in maturity weighted debt to GDP is associated with a 12 basis point rise in long-term Treasury yields, relative to the short-rate. Including more recent data since 2008, this relationship substantially weakens, with the coefficient falling by more than half.

Importantly, column (3) of Table 1 adds a post 2008 dummy, and the coefficient indicates long-term Treasury yields are 170 basis points lower then one would expect based on debt supply and the state of the macroeconomy. This primarily occurs both because the slope of the yield curve has been relatively low and because debt supply has expanded dramatically over the past 15 years. The dummy coefficient allows for this lower slope in

[4]Data sources are discussed in Appendix A and all yield data are from Fred. Using the 10 year Treasury instead as the long-rate or 2 year as the shorter term rate, or various other maturities, produces similar results, as we show in Table 2. We chose to use the same yield as in Section 2.4.2 for consistency.

Table 1: Slope of the Yield Curve. We regress the slope of the yield curve on bond supply measured as log maturity-weighted debt to GDP (MWDGDP). Column (1) uses the pre-QE sample, column (2) the full sample, column (3) adds a dummy for the post QE period, and column (4) interacts this dummy with bond supply. Columns (5)-(8) repeat this analysis using one-year ahead excess bond returns in place of the slope of the yield curve to separate term premia from the expectations hypothesis channel. Controls include the 3-month Tbill and the unemployment rate. Newey-West standard errors with 20 quarterly lags in parentheses.

Notes.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Slope of the Yield Curve				Excess Bond Returns			
	Pre-2007	Full	Full	Full	Pre-2007	Full	Full	Full
ln(MWDGDP)	1.24***	0.53*	1.10***	1.25***	0.11***	0.07**	0.09**	0.11***
	(0.26)	(0.29)	(0.26)	(0.26)	(0.04)	(0.03)	(0.04)	(0.03)
Post 2008 Dummy			-1.74***	1.25			-0.07	0.34
			(0.41)	(1.22)			(0.06)	(0.22)
ln(MWDGDP) × Post				-1.95**				-0.28**
				(0.79)				(0.14)
TBill	-0.32***	-0.22***	-0.30***	-0.28***	-0.00	0.00	0.00	0.00
	(0.04)	(0.05)	(0.04)	(0.04)	(0.01)	(0.01)	(0.01)	(0.01)
Unemp	0.60***	0.41***	0.48***	0.42***	0.01	0.01	0.01	0.00
	(0.08)	(0.08)	(0.06)	(0.06)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	227	288	288	288	227	284	284	284
R-squared	0.73	0.54	0.66	0.69	0.09	0.05	0.06	0.09

the post QE era while keeping the positive relation between slope and debt supply. Column (4) instead interacts debt supply with the post 2008 dummy and finds a substantial *negative* coefficient, meaning higher debt supply is associated with a much flatter yield curve since 2008. The sum of the coefficients on debt supply indicate that, if anything, higher supply is associated with a lower slope in the post QE era. This lower sensitivity of yields to debt supply is consistent with King (2019) who finds a lower sensitivity in the 2008-2015 period.

Columns (5)-(7) repeat the same analysis but use one-year ahead excess returns on the long-term Treasury portfolio in predictive regressions. The coefficients confirm that

the positive association between debt supply and slope stems from the term premium channel: higher debt supply is associated with higher bond returns, rather than primarily stemming from an expectations hypothesis channel of higher expected short-rates. These findings suggest that the term premium is abnormally low given debt supply and the state of the macroeconomy. In contrast, the other control variables mainly correlate with slope in columns (1)-(4) because they are informative about expected future short rates rather than term premia, as is confirmed in columns (5)-(8). These findings are in line with Greenwood and Vayanos (2014), among others, who find that the short rate (Tbill) does covary with slope but only through the expectations hypothesis channel and mean-reversion in future short rates – it does not predict returns and thus plays a minimal role for term premia.

Figure 2 plots the predicted path of the slope based on pre-QE data and tells a similar story as our regressions. We use the pre-QE data, as in the regression from column (1) of Table 1, to construct the predicted slope of the yield curve out of sample in the post QE era. We show the actual slope in blue. Again, the actual slope is too low – long-term Treasuries appear expensive – by about 150 basis points on average. The figure makes clear this is not driven by one specific episode in the post QE-era as the predicted is above the actual for the entirety of the post QE period. In contrast, there is a tight fit between the two in the pre-QE era. Importantly, accounting for realized purchases by the Fed (subtracting them from maturity-weighted debt to GDP) does not change these patterns substantially.

The above findings that the slope of the yield curve appears "too low" in the post QE era is robust to various specifications and to measurement issues. Table 2 shows robustness of this result using various measures of the slope of the yield curve, including or excluding controls, or dropping the COVID period.

Column (9) computes "net" debt by subtracting out holdings by the Federal Reserve, including through QE purchases. Specifically, we recompute the maturity weighted debt

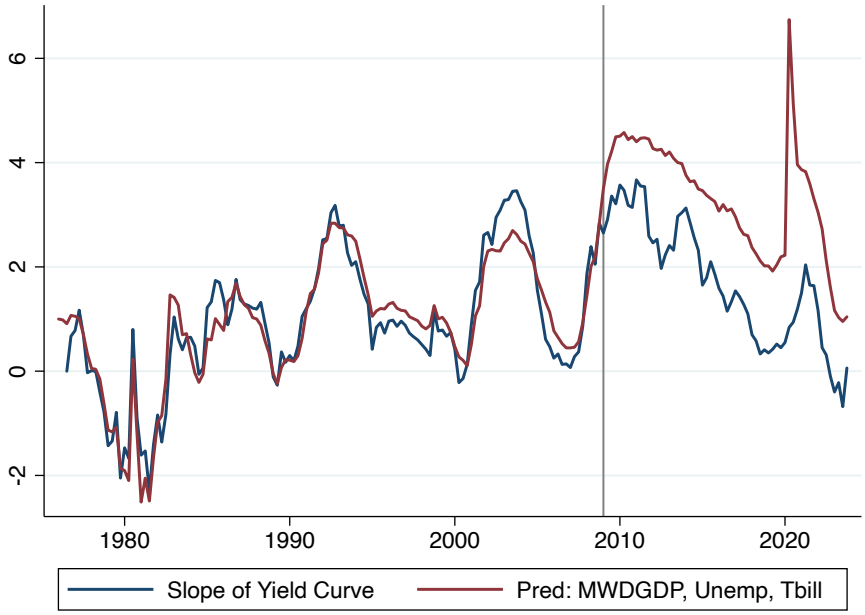


Figure 2: Predicted vs Actual Slope.
We regress the slope of the yield curve on bond supply in the pre-QE data following Table 1 with controls for the Tbill rate and unemployment rate. We plot the predicted value from this regression (red line) and compare to the actual slope of the yield curve (blue line). The gray dashed line indicates Q4 2008 when QE is first introduced.

to GDP measure after subtracting off all cusips held by the Fed in each period. Importantly, long-term Treasury yields remain low in the post QE era even when accounting for holdings by the Fed. Finally, while our main specification follows Greenwood and Vayanos (2014) in using maturity weighted debt to GDP, the results are similar when using total debt to GDP (using either book or market values of debt) or when using long-term debt to GDP.

A reasonable concern is that because the zero lower bound was binding during part of the QE period, the slope of the term structure might have been artificially low, i.e. it not that long bond yields were too low, but short yields that were too high. We follow Hanson and Stein (2015) and use yields on longer maturity bonds as a proxy for the monetary

Table 2: Slope of the Yield Curve: Robustness. We show robustness to various measures of slope as well as alternative regressions specifications of slope on bond supply. Post indicates the dummy for the post-QE period beginning at the end of 2008. Column (8) drops the COVID period. Column (9) uses "net" bond supply which subtracts out holdings by the Federal Reserve.

Panel A

	(1) 20yr-2yr	(2) 20yr-3m	(3) 10yr-3m	(4) 10yr-1yr	(5) 20yr-5yr	(6) No Tbill
Post_t	-1.33***	-1.74***	-1.34***	-1.14***	-0.54**	-1.37***
	(0.20)	(0.38)	(0.34)	(0.30)	(0.22)	(0.27)
N	190	288	283	283	283	190
R^2	0.86	0.66	0.53	0.65	0.75	0.81

Panel B

	(7) No Unrate	(8) No COVID	(9) Net Debt	(10) LT Debt	(11) Book Debt	(12) Net LT Debt
Post_t	-1.01	-1.40***	-1.22***	-1.27***	-1.08***	-1.31***
	(0.62)	(0.19)	(0.35)	(0.39)	(0.38)	(0.29)
N	288	186	185	190	190	190
R^2	0.34	0.89	0.79	0.80	0.80	0.80

policy instance. Our slope results hold using 10 year minus 1 year, 20 year minus 2 year, or even 20 year minus 5 year. Subtracting longer maturity yields further supports the notion that the medium term path of rates or monetary policy cycle is an unlikely driver of our results and further mitigates concerns about a temporarily binding zero-lower bound (ZLB). For example, for the 20yr minus 5yr results, the entire path of rates over the next 5 years is subtracted out.

In the appendix, we also include the "shadow" short rate (Wu and Xia (2016)), defined to be equal to implied short rate from the Taylor rule if the Taylor rule is zero or negative, and defined to be equal to the actual Tbill rate otherwise. Including this shadow rate does not affect the regression coefficients of interest (see Table 10 in the Appendix). If the ZLB was the driver of our results we would expect the shadow rate to predict the slope and

drive out our post 2008 dummy and the interaction with the debt supply variable.

This argument implicitly relies on the assumption that the ZLB was not perceived to be too persistent. Intuitively, if investors expected lift-off within a year, then using a 20 year minus 2 year slope is immune to ZLB effects. While the US economy stayed at the ZLB for a long period, what matters for bond yields is how long investors expected the ZLB to last in real time. Option data and Fed Funds futures data strongly indicate that investors expected a quick lift-off. Mertens and Williams (2021) use interest rate cap contracts to construct the risk-neutral distribution of short rates as far as seven years ahead and find that during this period the probability of short rates being at the zero lower bound hovered between 2.5% and 5%. Appendix B.1 discusses the ZLB in more detail. In particular, we plot the term structure of Fed Funds futures – contracts on where the Fed funds rate will be at various horizons in the future. The Fed Funds futures curve suggests investors consistently expect interest rates to return to normal levels of around 2% at 1-2 year horizons; they do not believe the Fed will leave rates at or near 0 for very long, so that the ZLB is highly unlikely to drive the behavior of long-term Treasury yields.

2.2.2 International Evidence: Exploiting the Staggered Introduction of QE policies

We've shown that long-term Treasury yields appear too low in the post QE era relative to the expansion in government debt supply. The evidence strongly suggests that the result is driven by low term premia in the post QE period. However, one may worry about broader trends affecting long-term Treasuries not captured in our analysis.

We now turn to international data, where we take advantage of staggered introduction to QE or asset purchase policies across countries. We construct the slope of the yield curve in 16 developed countries along with debt to GDP, unemployment, and inflation.[5]

[5]Our slope measure is generally the 10 year government bond yield minus a 3-month short rate due to data availability for international data. Further, our international data only has total debt to GDP rather than maturity weighted debt to GDP. We find this distinction makes little difference in the US sample and total debt to GDP is still a good proxy for the total supply of duration risk.

We then construct dummies for the post QE (or post asset purchase policy) period in each country. For example, we use 2008 for the US, 2009 for the UK, 2001 for Japan, and 2020 for Canada and Australia. For the Euro area, asset purchase policies targeted GIIPS countries from 2010-2012 and the broader Euro area starting in 2015, so we use 2012 for Italy, Ireland, Portugal and Spain (we don't have data for Greece) and we use 2015 for the rest of the Euro area. We note that while it is a misnomer to use the term "QE" for asset purchase policies in 2010-2012 for the GIIPS countries, the same logic of our argument applies: purchase announcements lowered yields substantially, suggesting supply effects, and the market likely interpreted additional future purchases if the state of the economy in these countries worsened and fewer purchases if the economic state improved. Appendix A provides details on the exact construction of these series, details the data sources, and discusses the dating of QE periods in more detail.

We run a panel regression of the slope of the yield curve on debt supply, a dummy for the post QE period, and controls which include the unemployment rate, rate of inflation, and country fixed effects. Table 3 shows a coefficient on the post QE dummy of -87 basis points in column (1), confirming that long-term government bond yields are particularly low following QE. The coefficient on debt to GDP is positive, indicating higher term premia when the supply of government debt expands. Column (2) interacts debt supply with the post QE dummy and finds a strong negative coefficient on the interaction term – indicating that the relation between debt supply and slope becomes much weaker in the post QE era across countries.

Columns (4)-(6) include time fixed effects which removes secular global trends in government bond yields and in debt supply and instead uses cross-country variation in the staggered implementation of QE. We see remarkably similar patterns. The coefficient on the post QE dummy indicates a 73 basis point decline in slope, and the interaction of the dummy with debt supply indicates a far weaker relation between supply and slope after QE. These results show that the differences we see in the post QE era are not driven by

Table 3: International Data. We regress the slope of the yield curve on log debt to GDP. $Post_{i,t}$ is a dummy equal to 1 after the first round of asset purchases occurs in country i in the sample. Data are annual. Columns (4)-(6) include time fixed effects which control for global trends. We include inflation and the unemployment rate as controls. See text for more details. Standard errors computed using Driscoll Kraay with 10 annual lags.

	(1)	(2)	(3)	(4)	(5)	(6)
$post_{i,t}$	-0.87***	1.07*	0.66	-0.73**	1.44***	1.22
	(0.19)	(0.58)	(0.76)	(0.32)	(0.46)	(0.79)
$ln(Debt/GDP)$	1.72***	2.90***	2.87***	1.00	2.43***	2.44***
	(0.54)	(0.71)	(0.68)	(0.61)	(0.59)	(0.58)
$ln(Debt/GDP) \times post_{i,t}$		-2.32***	-2.10***		-2.36***	-2.26***
		(0.79)	(0.76)		(0.73)	(0.72)
Inflation	-0.01	0.01	0.00	0.08	0.07	0.05
	(0.04)	(0.04)	(0.04)	(0.09)	(0.09)	(0.09)
Inflation $\times post_{i,t}$			0.38***			0.18*
			(0.08)			(0.11)
Unemp	0.18***	0.19***	0.20***	0.16**	0.15**	0.15*
	(0.04)	(0.03)	(0.04)	(0.07)	(0.06)	(0.08)
Unemp $\times post_{i,t}$			-0.03			-0.01
			(0.04)			(0.06)
N	515	515	515	515	515	515
Time FE	N	N	N	Y	Y	Y
Groups	16	16	16	16	16	16
R^2_w	0.198	0.248	0.261	0.583	0.622	0.624

the US or global factors and occur at the introduction of QE policies.

One potential concern with using debt to GDP is that it is highly persistent and near unit root. Table 4 runs the same regressions in differences where we use 5-year changes. Using 5-year changes is sensible given the potential mismatch in timing between yield changes and debt supply that higher frequencies would potentially miss. For example, asset prices are forward looking and will respond today to an expected increase in the supply of debt over the next year. Using 5-year changes still captures the broad relationships between these variables but helps avoid potential measurement error that higher frequency changes could introduce. We find similar results of a much weaker relationship

Table 4: International Data: 5 year changes instead of levels. We regress the slope of the yield curve on log debt to GDP. $Post_{i,t}$ is a dummy equal to 1 after the first round of asset purchases occurs in country i in the sample. Data are annual. Columns (4)-(6) include time fixed effects which control for global trends. We include inflation and the unemployment rate as controls. See text for more details. Standard errors computed using Driscoll Kraay with 10 annual lags.

	(1)	(2)	(3)	(4)	(5)	(6)
$post_{i,t}$	-0.69**	-0.18	-0.21	-0.74**	0.13	0.37
	(0.32)	(0.24)	(0.17)	(0.32)	(0.22)	(0.24)
$\Delta_5 ln(Debt/GDP)$	2.99***	4.24***	4.20***	2.18**	4.03**	4.15**
	(0.92)	(1.15)	(1.09)	(1.06)	(1.64)	(1.83)
$\Delta_5 ln(Debt/GDP) \times post_{i,t}$		-4.31***	-3.59**		-4.21***	-4.89**
		(1.57)	(1.54)		(1.28)	(1.92)
$\Delta_5 Inflation$	0.01	0.03	0.02	0.02	0.04	-0.02
	(0.05)	(0.05)	(0.06)	(0.09)	(0.09)	(0.09)
$\Delta_5 Inflation \times post_{i,t}$			0.17*			0.39***
			(0.10)			(0.14)
$\Delta_5 Unemp$	0.26***	0.29***	0.30***	0.17***	0.17***	0.14**
	(0.05)	(0.05)	(0.06)	(0.05)	(0.04)	(0.07)
$\Delta_5 Unemp \times post_{i,t}$			-0.05			0.09
			(0.08)			(0.08)
N	435	435	435	435	435	435
Time FE	N	N	N	Y	Y	Y
Groups	16	16	16	16	16	16
R^2_w	0.294	0.318	0.322	0.679	0.691	0.699

between debt supply and long-term government bond yields after QE or asset purchases are introduced.

In this last section we have seem in a variety of ways that long-term Treasury yields are low given the supply of Treasuries issued by the government or in the hands of private investors. The timing of this drop in yields coincides exactly with the start of the large-scale asset purchase programs. We now inspect these large-scale asset purchase announcements more closely.

2.3 Zooming in on QE News: Announcement Effects

We turn to an event-study analysis of the announcements of the large-scale asset purchase programs. These events are useful for identification as they zoom in on specific days where there is news about QE. These events have of course been extensively studied elsewhere in the literature (see for example Vissing-Jorgensen and Krishnamurthy (2011) for one of the first papers to document these facts). Our contribution here is to link these events back with the secular movements and provide a quantitative explanation for them. The announcement day evidence is consistent with the magnitudes shown in the broad time-series in the prior sections. Finally, we provide new evidence from options that these announcements affected very long-term yield dynamics, which points strongly towards our rule based interpretation.

Table 5 looks at Treasury yield changes by maturity on QE event days. The effects are small at the short end, gradually grow by maturity peaking at 140 bps for the 10 year maturity, and slowly decline to 80 bps in the 30 year bond. These magnitudes – while only coming from few announcement days – align closely with the magnitudes we've seen in the broad time-series evidence. We showed that long-term Treasury yields appeared "too low" by around 100-170 bps depending on which maturity is used. The QE announcement evidence shows that the broad low yield patterns in the time-series can be largely explained by very few days on which QE news was revealed.

Panel B of Table 5 recasts the raw yield changes by maturity in Panel A to the effect on forward rates at various horizons. We see substantial declines in forward rates, for example the 5 year 5 year forward rate falls by 170 bps on these announcement days. This emphasizes that news about the path of interest rates alone are unlikely explanations of the announcement effects – it is unlikely that investors learned that the average short rate would be 170bps lower beyond 5 years from now over these announcements. Instead, these results again point to the explanation that these announcements primarily affected

Table 5: Treasury Yield changes on QE Announcement Dates by maturity. We regress two-day change in Treasury yields on QE event dummies. We scale the event dummies so that the coefficient represents the cumulative change in yield over the announcements using a two-day event window. Newey-West standard errors with 4 lags in parentheses. Panel B repeats the same analysis using forward rates implied by the Treasury curve.

Panel A: Yields by Maturity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
VARIABLES	1M	3M	6M	1Y	2Y	3Y	5Y	7Y	10Y	20Y	30Y
QE Event	-0.10	-0.06	-0.32**	-0.37**	-0.51**	-0.76***	-1.10***	-1.35***	-1.39***	-0.90***	-0.76***
	(0.26)	(0.19)	(0.15)	(0.17)	(0.21)	(0.24)	(0.25)	(0.26)	(0.25)	(0.24)	(0.24)
Observations	4,037	4,037	4,037	4,037	4,037	4,037	4,037	4,037	4,037	4,037	4,037
R-squared	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.00	0.00

Panel B: Forward Rates

	(1)	(2)	(3)
VARIABLES	3yr 7yr Forward	5yr 5yr Forward	7yr 3yr Forward
QE Event	-1.51***	-1.70***	-1.66***
	(0.28)	(0.28)	(0.26)
Observations	4,037	4,037	4,037
R-squared	0.01	0.01	0.01

the term premium.

In Panel A of Figure 3, reproduced from Haddad et al. (2023), we also see that the effect of these announcements is highly front-loaded with the bulk of the effects on the 10 year yield coming in the first few announcements. While Table 5 emphasizes cumulative yield responses to QE news days, the figure highlights that movements in yields tend to be concentrated in initial announcements and fades with later ones. This is not due to the size of announced purchases – when available we show the "multiplier" of the yield change normalized by the purchase amount. This is consistent with earlier announcements having large effects through a "regime change" of a new dynamic policy rule rather than just reflecting news of the specific quantity of bonds purchased at

this announcement. In Haddad et al. (2023) we show that these early announcement responses are puzzlingly large given the relatively modest quantity of bonds that were actually bought by the Fed. That is, they would need abnormally large price impact to explain the announcement effects. In contrast, later announcements imply very weak or zero price impact (if they are assumed to be a surprise). Under the rule view, early announcements are much more powerful given they embed the present value of an entire state-contingent plan, while later announcements will be zero on average provided they do not (on average) represent deviations from the policy rule. Panels B and C indicate that this weakening announcement effect is a broader pattern that holds for both the UK and Euro area.

A natural possibility is that these events were associated with a new policy regime for asset purchases. Such a transition has the potential to amplify the asset pricing effects of these announcements because they lead investors to price not only the current announced quantities.

2.3.1 QE changed *dynamics*: Evidence from option prices around announcements

Finally, our story predicts that dynamic QE policies made bonds *safer*. That is, because investors expect the Fed to do purchases in bad times through a state-contingent plan, and thus long-term bond prices will be relatively higher in bad times, this would dampen risk from investors' perspective and lower bond yields ex-ante. We aim at this question by using option prices that reveal perceived changes in bond price dynamics around QE announcements.

We report these results in Table 6. Column (1) shows the cumulative implied volatility on the 10 year Treasury falling substantially, by 43%, over QE1 announcements. A drawback of this result is that the implied volatility is from options on the 10 year Treasury with a 1-month expiration. Ideally, we would measure implied-volatility using a much

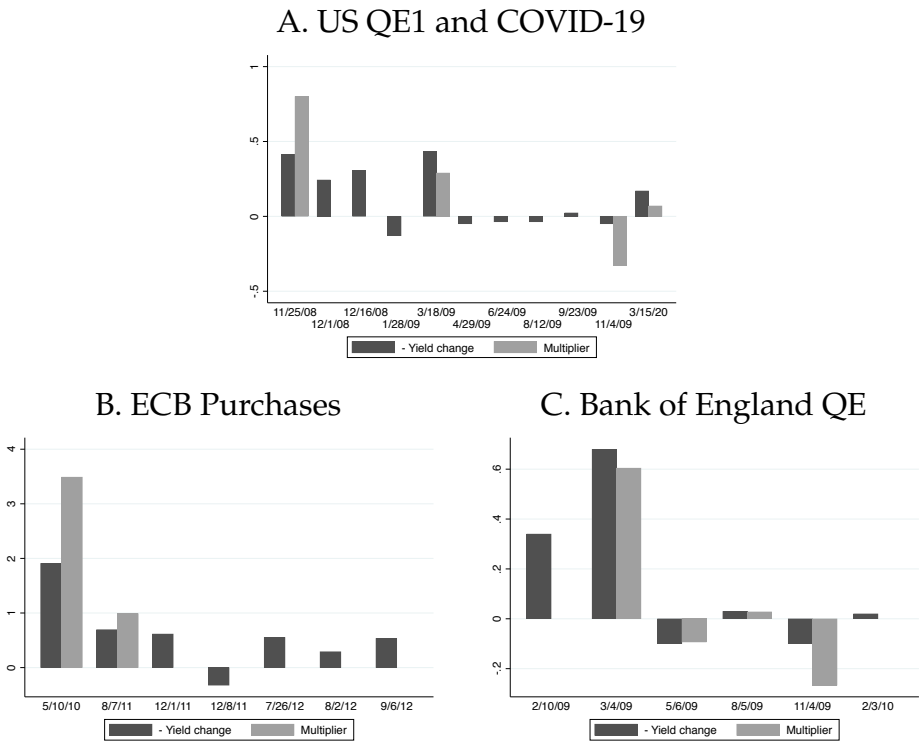


Figure 3: Weakening Announcement Effects of Asset Purchases

This figure is reproduced from Haddad et al. (2023) and plots announcement effects to asset purchase announcements made by the US Federal Reserve, Bank of England, and the ECB. See appendix for details, numbers come from various studies including Joyce and Tong (2012), Vissing-Jorgensen and Krishnamurthy (2011), Meaning and Zhu (2011), Gagnon et al. (2018), Vissing-Jorgensen (2021), and Krishnamurthy et al. (2018).

longer maturity for the options, but options on Treasury futures do not have significantly longer maturities. Columns (2)-(7) overcome this issue by using swaptions, which have much longer maturities. The drawback is that they are options on swaps and not Treasuries directly. Under no-arbitrage conditions the swap-Treasury spread should be small, but this doesn't have to be the case when there are convenience yields. [6] The advantage is that the long-maturity of the swap options allows us to assess long-term effects on volatility that our story predicts. We see large declines in implied volatilities in swaptions, with

[6]We note that Treasury yields and swaps respond quite similarly to QE announcements, meaning announcements do not result in large moves in the swap-Treasury spread.

Table 6: Option prices and implied volatility responses on QE dates. We compute log changes in implied volatilities and report numbers in percentage points, so that the QE event dummy coefficients represent percentage declines in volatility from the announcements.

	Implied Volatilities						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Mat 3M	Mat 3M	Mat 1Y	Mat 1Y	Mat 10Y	Mat 10Y
	$\Delta ln(\sigma_{10yr})$	Tenor 10yr	Tenor 30yr	Tenor 10yr	Tenor 30yr	Tenor 10yr	Tenor 30yr
QE1 Event	-42.98***	-24.71*	-39.66***	-35.83***	-39.66***	-37.94***	-42.16***
	(15.78)	(14.25)	(8.57)	(11.63)	(8.57)	(11.15)	(10.98)
N	3,092	3,960	3,695	3,960	3,695	3,960	3,695
R^2	0.00	0.00	0.01	0.00	0.01	0.00	0.00

magnitudes similar to those in Treasury futures. In particular, we see a decline of 38% for an option to enter a swap in 10 years that will last 10 years. This very long horizon suggests a large decline in very long-run volatility coming from QE announcements in line with our story. Specifically, QE announcements made long-duration assets much less risky, even at very long horizons.

2.4 Policy rule pass-through: Corporate yields and mortgage rates

While we've shown the effect of QE as a state-contingent plan is important for Treasury yields, we also evaluate the effects on other rates relevant for economic activity and investment. Specifically, we focus on corporate bond yields, mortgage rates, and MBS yields. A pass through of QE to these rates directly affects the cost of capital for firms and can spur housing demand and house price increases. The effect of corporate yields primarily reflects a spillover or indirect effect because long-term corporate bonds are substitutes to long-term debt securities targeted by QE. The effect on mortgage rates and MBS is both direct and indirect, as QE purchases have targeted both MBS and Treasuries.

The literature has already shown substantial pass through from QE announcements, so we focus our attention more on the long time-series patterns in long-term yields.[7]

2.4.1 The response of broader assets

Table 7 repeats our main regressions in Table 1 where we regressed the spread between long-term Treasuries and the short rate on debt supply and a dummy for the post QE era. We assess whether the result that long-term rates appear too low in the post QE era is also true for corporate bonds and mortgage rates. We find negative coefficients on the post QE dummy that are economically large. Corporate yields are lower by 1.2% for Aaa corporate debt and 0.9% for Baa rated debt (though the estimate for Baa is noisy). Mortgage rates are lower by 1.3%. While large in magnitude, these numbers are a bit smaller than what we find for Treasuries, reproduced in column (1). This suggests that some of the effect on Treasury yields from dynamic QE policy comes from affecting Treasury "specialness" or convenience yields. The next subsection explores this convenience yield effect in more detail for corporate bonds. Still, the effect on long-term debt yields is economically substantial, suggesting that the real effects are likely large.

Table 8 finds similar results when zooming in only on QE announcement days. We report cumulative yield changes on QE announcements for corporate bonds and mortgage backed securities. Yield changes for MBS are relatively large (1.2% and 1.6%) and in line with the response of Treasuries. This makes sense because QE purchases have targeted both MBS and long-term Treasuries. Yield changes for corporate bonds are smaller but still economically large at around 0.5-0.6%. The corporate bond series have a maturity of around 20 years, so these numbers are smaller than those for the 20 year Treasury yield at

[7]For example, Gilchrist and Zakrajsek (2013) and Gilchrist et al. (2015) found large and significant effects of QE on corporate bond yields and credit risk in the United States, D'Amico and Kaminska (2019) found significant and persistent effects of various rounds of QE on corporate bond yields for the UK, Rosa (2012) and Mamaysky (2018) found significant effects of QE on equities and equity-implied volatility for several countries, and Di Maggio et al. (2020) showed that the Fed's QE programs that did not include MBS purchases reduced mortgage rates.

Table 7: Pass Through: Mortgage Rates and Corporate Bonds. We regress the yield spread of long-term Treasuries, long-term corporate bonds, and mortgage rates on bond supply measured as log maturity-weighted debt to GDP (MWDGDP) and a dummy for the post QE era. Spreads are taken relative to the 3-month Treasury bill rate. Controls include the 3-month Tbill and the unemployment rate. For corporate bond spreads, we additionally control for log stock market volatility (computed using the sum of squared weekly stock market returns in a given quarter) as a measure of default risk. Newey-West standard errors with 20 quarterly lags in parentheses.

	(1)	(2)	(3)	(4)
VARIABLES	Treas	Corp: Aaa	Corp: Baa	Mortgage
ln(MWDGDP)	1.10***	0.76**	0.65*	0.57**
	(0.26)	(0.33)	(0.35)	(0.24)
Post 2008 Dummy	-1.74***	-1.17**	-0.87	-1.31***
	(0.41)	(0.54)	(0.58)	(0.37)
TBill	-0.30***	-0.26***	-0.21***	-0.19***
	(0.04)	(0.06)	(0.06)	(0.04)
Unemp	0.48***	0.51***	0.63***	0.42***
	(0.06)	(0.07)	(0.08)	(0.06)
$ln(\sigma_t)$		0.60**	0.90***	
		(0.25)	(0.28)	
Observations	288	285	285	211
R-squared	0.66	0.59	0.59	0.51

-0.9%. These relative magnitudes – stronger effects on mortgages and somewhat weaker effects on corporate bonds, are consistent with the broad time-series results from Table 7 and again point towards positive, but imperfect, pass through to other assets, particularly corporate bonds. Thus, the well-identified event study analysis of QE announcements again accords with the broad time-series patterns of lower yields.

2.4.2 Convenience Yield Effects

The behavior of convenience yields has been the focus of a large literature after the 2008 financial crisis. While in our context they are a direct reflection of the positive but imperfect pass through we just documented, looking at convenience yields directly is useful to em-

Table 8: Corporate and MBS yield responses on QE dates. We compute two day changes in yields and report numbers in percentage points. We scale the event dummies so that the coefficient represents the cumulative change in yield over the announcements using a two-day event window. Newey-West standard errors with 4 lags in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Treas (10 yr)	Treas (20 yr)	Aaa	Baa	MBS	MBS: Ginnie Mae
QE Event	-1.39***	-0.90***	-0.49**	-0.62**	-1.62***	-1.17***
	(0.25)	(0.24)	(0.25)	(0.24)	(0.30)	(0.40)
Observations	4,046	4,046	4,046	4,046	4,040	3,536
R-squared	0.01	0.00	0.00	0.00	0.01	0.00

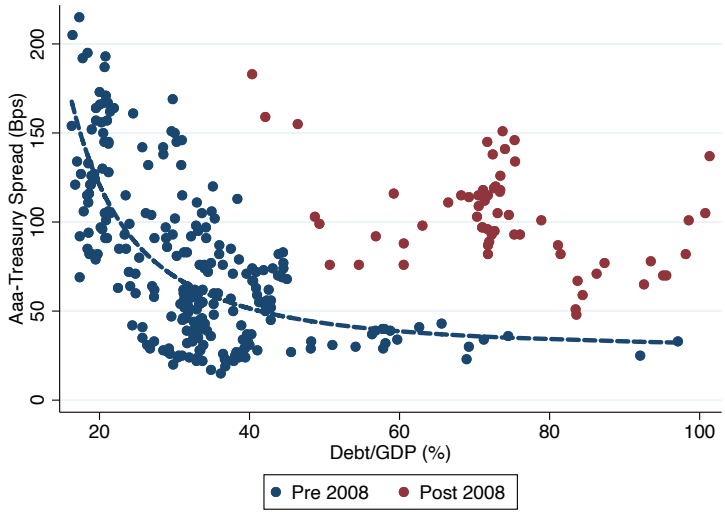
phasize two points. First, consistent with a supply explanation, long-term Treasuries have been abnormally expensive since 2008 relative to Aaa corporate bonds of similar maturities. The fact that Treasuries in particular were too expensive lends credence to the view that low long-term bond yields were driven not by broad economic forces (e.g., secular stagnation or the ZLB), but instead by forces originated in the Treasury market. Second, it gives us an alternative, more conservative, benchmark for what long-rates would have been if we assume there is no effect of QE on corporate yields.

We start from the original evidence in Krishnamurthy and Vissing-Jorgensen (2012) in the blue dots in Figure 4 using quarterly data from 1947-2008.[8] The figure plots the Aaa-Treasury spread against the supply of US government debt relative to GDP. Treasuries and Aaa bonds have very similar cash-flows, so the difference in yield measures Treasury convenience yields stemming from safety, liquidity, or "preferred habitats." As shown in the blue dots, when the supply of Treasuries is scarce, the premium on Treasuries goes up, consistent with a downward sloping demand curve for Treasuries.

The red dots in Figure 4 show this relationship breaks down since 2008. Comparing

[8]Krishnamurthy and Vissing-Jorgensen (2012) use a longer sample back to 1919 but focus on annual data, but show very similar patterns. We start in 1947 since we focus on quarterly data. See Appendix A for details on the exact construction of these series.

Panel A: Total Supply



Panel B: Net Supply

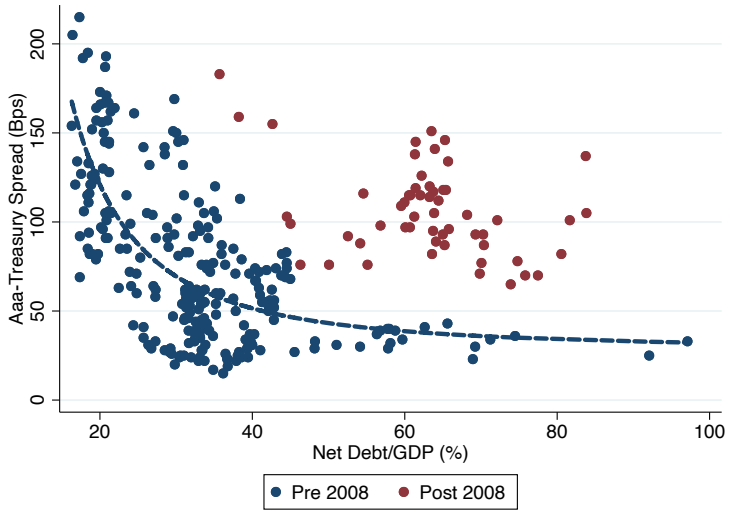


Figure 4: KVJ Updated.
This figure updates Krishnamurthy and Vissing-Jorgensen (2012). Panel B uses net supply, where net supply subtracts off holdings of Treasuries by the Federal Reserve.

the blue dashed-line and the red dots we see that the convenience yield has been substantially higher than the pre-2008 price-supply relationship would have predicted. Roughly,

the gap between the red dots and the blue dashed curve suggests that Treasury yields, relative to Aaa corporate bonds, are too low by around 75 basis points. In Panel B of Figure 4 we net out from the supply numbers the treasuries held by the Fed. As a result the red dots shift only slightly to the left. While convenience yields are a bit less excessive, they are still substantially so.

Appendix B.2 provides formal regression analysis to show that convenience yields for long-term Treasuries are meaningfully larger in the post QE period. Importantly, this result is present only for long-term bonds, the spread between commercial paper and Treasury bills shows no such break. This is informative because it tell us this is something about long-term bonds in particular and therefore points to the large-scale asset purchase programs which were initiated in 2008 and targeted long-term bonds. It cuts against "financial repression" and regulation driven demand shifts for government bonds since if anything Treasury bills have been treated more favorably in the post-2008 bank regulatory framework. As before, we find that QE announcement effects are largely in line with these same patterns.

Acharya and Laarits (2023) emphasize the relation between convenience yields and the stock-bond covariance, which could be another driver of the effects we show here. Figure 9 analyzes the relation between Treasuries and the stock market by computing the stock-bond covariance. We don't find a positive relation between this covariance and the slope of the yield curve. Moreover, we don't find a significant break in the stock-bond covariance around QE – while the beta is negative during the QE period, the trend in covariance occurs earlier (e.g., in the late 90's), consistent with other findings in the literature. On QE announcement days, there is no significant reaction of the stock market. Thus, we find no significant evidence of a spillover to the stock market from QE or a change in the stock-bond relationship due to QE.

2.5 Crowding In and the reduced sensitivity to supply shocks after QE

An important implication of the rules-view of how QE announcements impact asset prices is that we should see a "crowding in" effect of a greater willingness to absorb risk. The presence of the rule makes markets more elastic as arbitrageurs expect that eventually the central bank will absorb supply-demand imbalances, and therefore are more willing trade against these imbalances.

To assess this, we look at the bond market response to the Treasury quarterly refunding announcements. In the first week of second month of each calendar quarter (i.e., February, May, August, and November), the US Treasury announces the planned borrowing for the next calendar quarter, how it will distribute the borrowings across maturities, and how it is borrowing in the current quarter. Specifically, on Monday Treasury puts out a statement disclosing how much they will be borrowing in the current quarter, and how much they plan to borrow in the next calendar quarter. Then on Wednesday, the US Treasury puts out a statement and holds a press conference where they disclose the planned size of the auctions of each bond for the quarter ahead and discuss how they expect these auction sizes to evolve afterwards.

We use the bond market response to these announcements to measure the market sensitivity to news about the supply of duration in these auctions. We construct a proxy for this news by looking at the realized issuance behavior of the Treasury over the following quarters. To stay consistent with the rest of our analysis, we construct the maturity-weighted-issuance over GDP (*mwi*). We use the CRSP monthly return data and assign the quarter for the first time a particular bond appears in the data set as the quarter the bond was issued and we use the amount outstanding in the data set as the amount issued. We then sum across all the amounts outstanding of all bonds that appear in the CRSP data set for the first time in a given quarter and multiply by the outstanding maturity as of that

date. We normalize this by nominal GDP.

$$mwi_t = \frac{\sum_i B_{i,t} M_{i,t}}{GDP_t}.$$

We then look at future changes in this variable as our proxy for news. To understand why this is a valid proxy consider the following decomposition for innovations,

$$\Delta mwi_t = mwi_{t+q_+} - mwi_{t-q_-} = E_{t_-}[\Delta mwi_t] + (E_t - E_{t_-})[\Delta mwi_t] + (1 - E_t)[\Delta mwi_t]),$$

where the first term captures expectations before the refunding announcement, the second term captures the expectation revision during the announcement, i.e. the funding news shock we are interested in, and finally the third term captures the wedge between the realization and what was expected. Both the first and the third term are sources of measurement error, but we look at yield changes around the announcements which are uncorrelated with the first and third terms.

Thus a regression of yield changes around refunding announcements on our proxy for re-funding news Δmwi_t recovers the elasticity of yields to funding shocks up to a downward bias due to the presence of first and third terms. We present results for a few different specifications in the Table 9. We first see that it is indeed the case that yields respond to this news with the expected sign: an increase in amount of duration the government will place in the market leads to an increase in the slope of the yield curve. Importantly, this affect is attenuated after 2008. Consistent with our low frequency results, bonds markets are empirically less sensitive to supply news after 2008.

Table 9: Supply Shocks

The left hand-side variable are changes in the slope of the term structure around Treasury Quarterly Refunding Announcements. We compute slope by using the 20 year and 2 year maturities. Our event window goes from Friday close to Thursday close in weeks with a refunding announcement. This choice is due to the fact the refunding announcement occurs partially on Monday before market opening and partially during Wednesday, so we follow the rest of our analysis and consider an event window of one day after the announcement. The different columns use different proxies for the maturity news. In columns 1 and 2 we use the change in issuance between the current quarter and four quarter ahead. Columns 3 and 4 look at changes between the current quarter and two quarters ahead. Columns 2 and 4 control for seasonality effects of issuance behavior with quarter dummies.

	(1)	(2)	(3)	(4)
ΔMWI	0.69***	0.69***	1.03***	1.04***
	(0.24)	(0.24)	(0.35)	(0.35)
$\Delta MWI \times$ Post	-0.59**	-0.59**	-0.74*	-0.76*
	(0.27)	(0.27)	(0.40)	(0.40)
Post	-0.00	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)
R^2	0.0009	0.0012	0.0010	0.0013
N	7606	7606	7606	7606
Quarter FE	No	Yes	No	Yes
Horizon (qtrs)	4	4	2	2

3. A Simple Model with an Asset Purchase Rule

We turn to a quantitative model of the yield curve which incorporates an asset purchase rule. Such a model allows us to investigate the quantitative plausibility of our findings and the importance of various mechanisms. Furthermore, it provides a laboratory to assess the impact of alternative purchase rules.

We follow a structure similar to Greenwood and Vayanos (2014), Haddad and Sraer (2020), Vayanos and Vila (2021), Greenwood et al. (2023), with the addition of an asset purchase rule. At a high level, the bond market is populated by two types of investors: inelastic investors and arbitrageurs. The latter are marginal investors in bond market, and must absorb variation in the supply of bonds. Similarly, by shifting their portfolios,

quantitative easing impacts the term premium.

3.1 Setting

State of the economy. Time is discrete, and there are default-free zero-coupon bonds maturing at dates 1 up to τ^+. We assume the vector Z_t summarizes the state of the economy and follows a VAR(1) process that we can write as:

$$Z_{t+1} = \Gamma Z_t + \epsilon^z_{t+1}, \tag{1}$$

In our main analysis we focus on the case that the economic fundamentals Z_t consists of the risk-free rate $r_t = \bar{r} + \delta_r Z_t$ (the yield on the one-period bond) and the total supply of government bonds as a share of GDP $s_t = \bar{s} + \delta_s Z_t$. We later augment Z_t to study demand shocks, market functioning shocks, and crisis-driven supply shocks.

The mix of government bonds of different maturities is given by vector w_s, with $1'w_s < 1$, where $1 - 1'w_s$ is the relative supply of one-period bonds, and the length of w is $\tau^+ - 1$, where τ^+ is the longest maturity issued.

Investor preferences. Arbitrageurs and inelastic investors hold this bond supply. Inelastic investors have demand for individual bonds given by $w_\zeta \bar{\zeta}$, where the vector w_ζ denotes the mix of bonds that they buy and $\bar{\zeta}$ the overall size of their bond portfolio as a share of GDP.

Arbitrageurs close the model by buying and selling bonds with the objective of maximizing their mean-variance objective as follows:

$$\max_{X_t} E_t[W_{t+1}] - \frac{\gamma}{2}\text{Var}_t[W_{t+1}] \tag{2}$$

The vector X_t is of dimension $(\tau^+ - 1)$ where $X_t(\tau)$ is the arbitrageur's holding of bonds

of maturity $\tau > 1$. Their wealth evolves according to their portfolio choice and the returns $R_{\tau,t+1}$ of the different bonds:

$$W_{t+1} = (W_t - 1'X_t)(1 + R_t) + \sum_{\tau=2}^{\tau^+} X_t(\tau)(1 + R_{\tau,t+1}). \tag{3}$$

Following Greenwood et al. (2023), we log-linearize returns $r_{\tau,t+1} = \log(1 + R_{\tau,t+1})$:

$$R_{\tau,t+1} \approx r_{\tau,t+1} + \frac{1}{2}\text{Var}_t(r_{\tau,t+1}).$$

This allows us to write the arbitrageur's problem as:

$$\max_X W_t(1 + r_t) + \sum_{\tau=2}^{\tau^+} X(\tau)E_t\left[r_{\tau,t+1} + \frac{1}{2}\text{Var}_t(r_{\tau,t+1}) - r_t\right] - \frac{\gamma}{2}\text{Var}\left[\sum_{\tau=2}^{\tau^+} X(\tau)r_{\tau,t+1}\right].$$

Equilibrium yield curve. In equilibrium, the arbitrageur must hold the entire net supply of bonds at each maturity:

$$X_t = ws_t - w_\zeta\bar{\zeta} = w\bar{s} + \delta_s Z_t - w_\zeta\bar{\zeta}, \tag{4}$$

The first-order condition of the arbitrageur's problem is then

$$E_t[r_{\tau,t+1}] + \frac{1}{2}\text{Var}_t(r_{\tau,t+1}) = r_t + \gamma\text{Cov}_t\left(r_{\tau,t+1}, \sum_{\tau=2}^{\tau^+} X(\tau)r_{\tau,t+1}\right), \tag{5}$$

which states that each bond should earn the risk-free rate plus a premium that depends on the arbitrageur's risk aversion γ and the covariance of bond τ with the arbitrageur's portfolio. Supply or demand shocks will feed into risk-premia as they change the composition of the arbitrageur portfolio.

We conjecture that log bond prices are given by

$$p_{\tau,t} = A(\tau) + B(\tau)Z_t. \tag{6}$$

Plugging this conjecture into the arbitrageurs' FOC, we obtain[9]:

$$B(\tau) = -\delta_r + B(\tau-1)(\Gamma - \gamma\Sigma B' w_s \delta_s), \tag{7}$$

$$A(\tau) = -\bar{r} + A(\tau-1) - \gamma B(\tau-1)\Sigma B'(w_s\bar{s} - w_\zeta\bar{\zeta}) + \frac{1}{2}B(\tau-1)\Sigma B(\tau-1)', \tag{8}$$

with boundary conditions $B(1) = -\delta_r$ and $A(1) = -\bar{r}$. Here B is a matrix with all the maturity specific bond state-variable sensitivities ($B = [B(2), ..., B(\tau^*)]$). Thus this is a fixed point problem. Note that what matters are the supply-weighted sensitivities to shocks $B_w = B'(w_s\delta_s)$ in the arbitrageur's portfolio. To solve it, we guess an initial sensitivity of the investors portfolio to the state variables B_w, we then solve for the coefficients, and iterate until the equilibrium sensitivity of the underlying bonds is consistent with the sensitivity of the investors portfolio, i.e. $B_w = B' w_s \delta_s$ holds.

3.2 Asset Purchase Rules

We consider a simple family of purchase rules, where the central bank's balance sheet is a function of the state Z_t:

$$qe_t^d = \overline{qe} + \delta_{qe} Z_t, \tag{9}$$

which indicates that purchases are affine with respect to the state variables Z_t (for example, the Fed buys based on the output gap, inflation, and debt supply). Note that since bond prices are a function only of Z_t this rule also allows for purchases that depend on term premia or price dislocations, consistent with interventions based on market func-

[9] See Appendix B.3 for details on the model solution.

tioning. The term qe_t represents the quantity of bonds purchased, normalized by GDP. The superscript d denotes the dynamic nature of this rule. We denote by w_{qe} the portfolio corresponding to one unit of purchases, and assume it is constant over time.

A simplifying assumption of this setting is that "policy shocks" are absent and that the purchases depend solely on the state of the economy. Note also that, because we take short-rate dynamics as exogenous, this model gives no role for the "signaling channel" that QE announcements affect long-term bond yields by changing the expectation of future short rates and thus lowering yields through the expectations hypothesis. We make this assumption because there is substantial evidence that the signaling channel plays a fairly minor role in QE announcement effects, particularly for longer term yields.[10]

Finally, note also that the model focuses only on the yield curve. As such the risk premium in our affine model applies specifically to Treasuries, so that the pricing can be driven by duration risk Greenwood and Vayanos (2014) or by a pricing kernel that embeds convenience yields in Treasuries as in Krishnamurthy and Vissing-Jorgensen (2012). While our empirical work suggest an imperfect degree of substitution between Treasuries and corporate bonds, we leave such modelling aside in the interest of simplicity.

In equilibrium, arbitrageurs are on the other side of the central bank's asset purchases. Their holdings must then satisfy:

$$X_t = w_s \bar{s} + \delta_s Z_t - w_\zeta \bar{\bar{\zeta}} - w_{qe} qe_t \tag{10}$$

$$= (w_s \delta_s - w_{qe} \delta_{qe}) Z_t + (w_s \bar{s} - w_\zeta \bar{\bar{\zeta}} - w_{qe} \overline{qe}). \tag{11}$$

[10]Eser and Schwaab (2016) find "signaling of future low interest rates did not play a role" in the yield impact of asset purchases for the ECB.[11] Vissing-Jorgensen and Krishnamurthy (2011) do find a role for the signaling channel by looking at the response of Fed funds futures, but quantitatively the effects from this channel are relatively small particularly for longer maturity bond yields. Vissing-Jorgensen and Krishnamurthy (2011) estimate upper bounds of 40bps from this channel. Joyce et al. (2010) find that the signaling channel played a similarly small role in the UK.

This leads to an altered recursion for the the coefficient of bond prices:

$$B(\tau) = -\delta_r + B(\tau-1)(\Gamma - \gamma\Sigma B'(w_s\delta_s - w_{qe}\delta_{qe})), \tag{12}$$

$$A(\tau) = -\bar{r} + A(\tau-1) - \gamma B(\tau-1)\Sigma B'(w_s\bar{s} - w_\zeta\bar{\zeta} - w_{qe}\overline{qe}) + \frac{1}{2}B(\tau-1)\Sigma B(\tau-1)'. \tag{13}$$

Both the unconditional level of bond purchases $\overline{qe}$ and the conditional aspect of the policy rule δ_{qe} shape yields. From the above expressions, we observe that a purchase rule has the static effect of altering current risk premia—terms like $\gamma B(\tau-1)\Sigma B'(w_{qe}\overline{qe})$ for the state-contingent rule. Moreover, it affects the unconditional risk premia through the term $B(\tau-1)\gamma\Sigma B'(w_{qe}\delta_{qe})$, which directly impacts the bond prices' sensitivity to the state variables.

This unconditional risk-premium effect should be evident in data if the economy undergoes a sharp transition across policy regimes where the purchase rule impacts both $B(\tau)$ and $A(\tau)$. Therefore, this model suggests that QE announcements associated with regime transitions can be correlated with an unusually large effect on yields and implied volatilities of yields.

4. Quantitative implications of the model

4.1 Calibration

Simplifying the Debt Maturity Structure To take the model to the data we make the additional simplifying assumption that the government only supplies a single maturity $\tau*$ every period. Thus $w(\tau = \tau^*) = w_{qe}(\tau = \tau^*) > 0$ and $w(\tau \neq \tau^*) = w_{qe}(\tau \neq \tau*) = 0$. One can think of this maturity as the average duration of government bonds. In this interpretation, the model approximates an economy where the government issues many

maturities but keeps the overall maturity structure of long duration bonds constant. The duration of the net-supply of government debt is free to vary as the government implicitly increases the issuance of short-term bonds as it buys more long-term bonds.

Yield curve without purchases We calibrate the model without purchases using the pre-QE data used in Section 2.2 from 1952-2007. To estimate the state variable dynamics, we run a first-order VAR using the short-rate (3 month T-bill) and debt supply, where supply is measured as maturity weighted debt to GDP as in Section 2.2. Relative to the estimated VAR, we change the mean of the short rate to be 2%. The mean in the historical sample is higher but there are a few good reasons to go with the lower number. First, 2% is much closer to the average short-rate in the last 30 years. Second, 2% is typically thought of as the Fed's target rate in the modern era. This choice is irrelevant for any of our conclusions because the mean short rate only determines the average level of rates but does not enter in to prices of risk that control the average slope of the yield curve or the dynamics of term premia.

We then choose the parameters driving the price of risk γ and $\bar{\zeta}$ to match the behavior of the yield curve coefficients A and B. Specifically, we target the sensitivity of yields to supply and the short rate, focusing on the 10-year maturity as the long-term yield. We will show that this simple target leads to matching the pattern of coefficients across maturities in the data quite well. We also target the unconditional slope of the term structure in our data, which produces an average slope between the 10-year Treasury yield and the 3-month Treasury yield of about 200 bps.

Purchase rule We calibrate the purchase rule to the empirical behavior of purchases. Figure 5 plots empirical purchases against the supply of Treasuries in the data. We see a clear positive correlation that times when supply expands, purchases also expand significantly. A regression coefficient suggests a loading of 0.5 on supply. A regression in levels

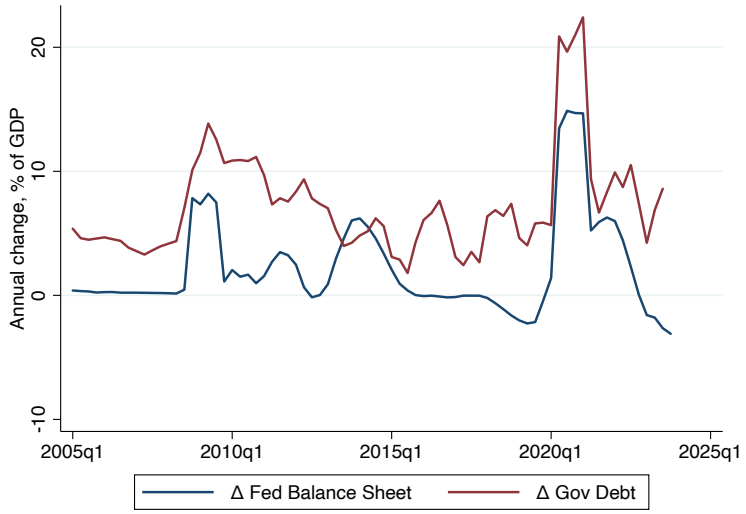


Figure 5: Purchases and Supply Empirically

Plot of purchases against supply. Annual changes.

produces a lower number at around 0.35. In the model we have outlined, this correlation makes sense if the Fed is targeting term premia, since term premia increase with supply both in the model and in the data we have presented earlier. Thus, we do not interpret this relation in the sense that supply is being directly targeted, but that it provides a proxy for term premia. Motivated by this evidence, we consider a rule with a loading on supply of $b_s = 1/3$, which is at the low end of what we find empirically.

As we will show in the next section, using a value around $1/3$ is a compromise between the empirical purchases, which tend to favor a higher number, and the yield dynamics in the post QE era, which tend to favor a slightly lower number. For now, we set the coefficient of purchases on the short-rate b_r to zero so that we only focus of supply, but we will return to rules that depend on short rates in Section 4.3.

4.2 How effective is the asset purchase rule?

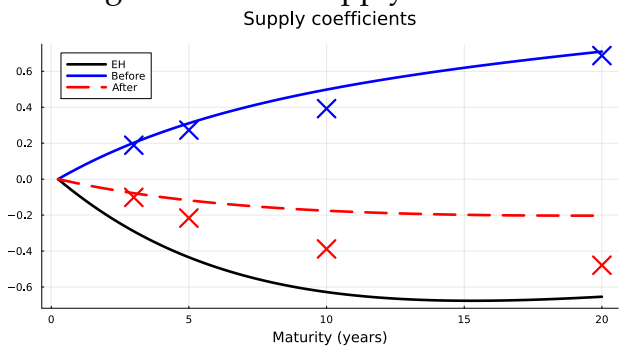
Figure 6 reports the behavior of the yield curve after inclusion of the policy rule. Note that by setting the constant a_{qe} equal to 0, we are assessing the effects of purchases that are, on average, equal to 0 – the rule is buying bonds as often as it sells them. This is a useful benchmark because it isolates the pure "insurance channel" of QE on the yield curve, without mixing static effects coming from constant purchases.

In Panel A, the blue curve plots supply coefficients by maturity in the model without purchases (the loadings $B(n)$[12]). We see an increasing pattern. Longer duration assets have a larger loading on supply because they are more affected by changes in the term premium. The blue X's represent the same coefficients in the data, estimated using the 1952-2007 sample, before QE was implemented. The red dashed curve plots supply coefficients in the model once we introduce the asset purchase rule. As expected, purchases that are positively correlated with supply will lower the coefficient or loadings of yields on supply. The coefficients in the model go from positive to negative, even though the rule only purchases 1/3 of supply deviations. This is because of the correlation between supply and interest rates in the dynamics of the state variables X_t: positive supply shocks are negatively correlated with short term interest rates, and so all else equal a positive shock to debt supply lowers long-term yields absent any changes in term premia. The black line makes this clear – it represents the expectations hypothesis benchmark in the model, which turns off all risk premia effects by settings $\lambda_t = 0$. The red line thus ends up between the black and blue lines because the rule only partially offsets shocks to debt supply. The red "X"'s show the coefficients in the data of yields on debt supply in the post QE era, and are in line with the models' predictions. We are thus able to match the low sensitivity to debt supply after the introduction of QE.

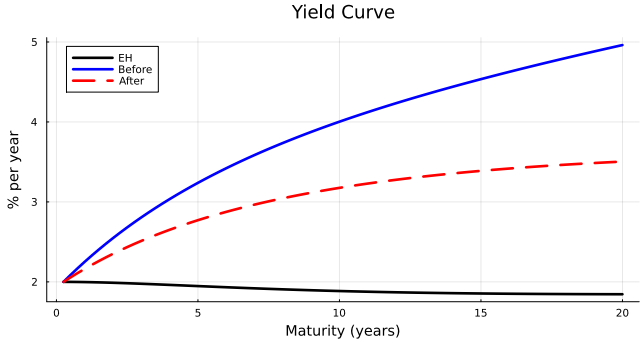
It is useful to compare the blue (pre QE), red (post QE rule), and black (expectations

[12] $B(n)$ gives the loadings for log bond prices, which we convert to loadings for yields as $-B(n)/n$.

Panel A: Loading of Yields on Supply Before and After Rule



Panel B: Unconditional Yield Curve Before and After Rule



Panel C: Yield Curve Conditional on Elevated Supply

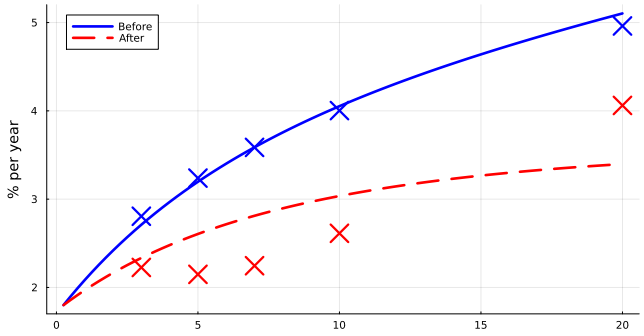


Figure 6: Implementing the purchase rule in the model

Panel A shows regression coefficients (loadings) of bond yields on supply in the data and in the model, both before (blue) and after (red) the purchase rule is implemented. Loadings in the data are indicated by "X" markers in the pre-QE and post-QE samples. The black line, labeled "EH" gives loadings in the model under the expectations hypothesis benchmark where all risk premia are zero. Panel B shows the unconditional yield curve before and after the rule in the model. Panel C shows the effect of the rule on the yield curve when the rule is implemented during a period of elevated supply, calibrated to the level of supply when QE was implemented in 2008-09.

hypothesis) lines in Panel A. Surprisingly, the introduction of the purchase rule moves the supply coefficients the majority of the way to the black line (no risk premia), despite only absorbing 1/3 of supply shocks. This occurs because of the "crowding in" effect in the model which amplifies the effect of the purchase rule on risk premia. By making supply shocks less volatile through purchases, there is a direct reduction in the sensitivity of yields to supply. However, because this reduces the overall risk investors must bear, they also become more willing to absorb shocks.

Panel B shows the effects of the rule on the unconditional yield curve in the model, found by computing the coefficients $A(n)$ before and after the rule. The red curve shows a substantial reduction in yields, particularly at longer maturities. For a 10 year-bond, the reduction in yield coming from the rule is 75 bps. Why are these effects so large? In the model, the purchase rule makes long-term bonds safer, which has an immediate direct effect. However, this safety also affects the volatility of the pricing kernel, which depends on the long term bond return as well. In this way, the end result is a substantial reduction in yields.

Finally, Panel C reports the impact of starting the rule in a period with elevated supply, combining a conditional and unconditional effect. This is relevant empirically since QE was introduced in the US in 2008-2009 as debt supply was large. We symmetrically introduce the QE rule in the model when supply was equal to debt supply in 2009 and compare the model response to the announcement effects from event studies presented earlier. We see that the magnitude of the response is very close to what is observed in event studies around this period, as shown in the red vs blue "X"'s.

Figure 7 plots the volatility of bond returns before and after the rule in the model. Because the rule dampens fluctuations in term premia, it lowers the volatility of longer-maturity bonds but has virtually no effect on bonds at shorter horizons. The rule lowers the volatility of the return on a 10-year Treasury by 10% and lowers the 20-year by 20%. In contrast, the evidence presented empirically suggests larger drops: around 38% using

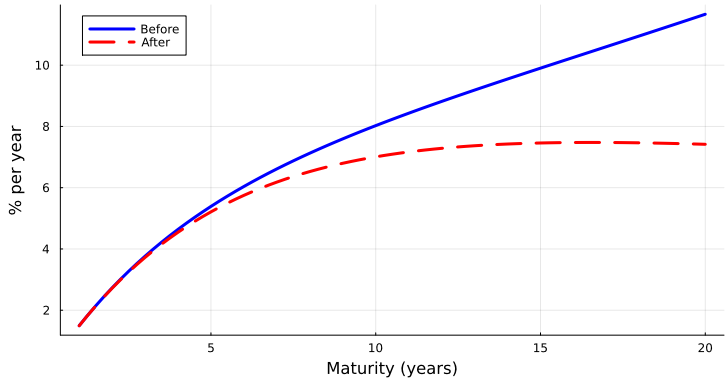


Figure 7: Volatility of Bond Returns

Volatility of long-term bonds goes down after purchase rule. The blue line plots the volatility of bond returns in the model before the QE rule is implemented, the red dashed line plots the volatility afterwards. The y-axis is in percent per year.

the event study approach and around 17% by taking a post-QE average. Thus, while the purchase rule in the model does lower volatility, as in the data, it is not as large as what we see empirically.

4.3 The impact of counterfactual rules: QE and short-rates

So far, we have only focused on the effects of a rule depending on supply through the coefficient b_s. In reality, purchases also depend on the short-term interest rate r_t with purchases occurring in periods where the short-term interest rate is low or near 0. We set $b_r = -10$. Since the unconditional short-rate in our calibration is set at 2%, this implies that a drop of interest rates to 0 is associated with purchases of 20% of GDP. This is a bit larger than the empirical variation and thus likely an upper bound on effects coming from this channel. Figure 8 plots the results for when we add a rule that depends on the short-rate, on top of the supply purchases of $1/3$, shown in green.

We observe the loading of the rule on the short rate has little effect on the term structure of yields and if anything raises yields slightly. The increase in yields comes from the fact that long-term bond prices fall when short rates rise (all else equal) and vice versa.

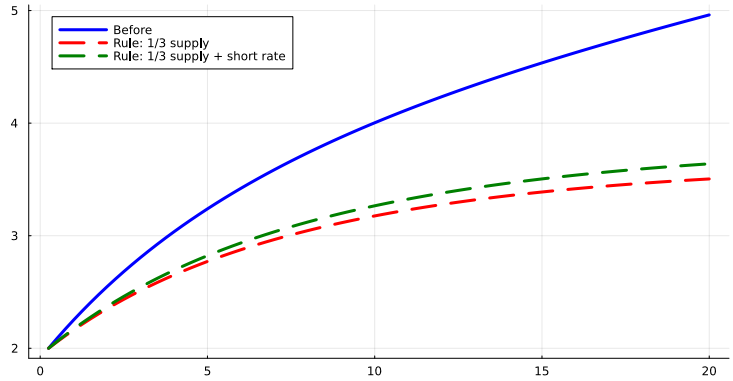


Figure 8: Yield Curve Effects with Alternative Rules

We plot the effects on the yield curve for alternative bond purchase rules. Specifically, the green line shows the effects of a rule that depends both on supply and interest rates, while the red line shows the effects of a rule that only depends on supply as shown earlier. The green line assumes purchases of 10% of supply for a 1% decline in the short rate (for example, if the short rate falls from 2% to 0%, the Fed purchases 20% of supply).

Purchases that are large when interest rates are low add volatility to long-term bonds (and similarly for sales or quantitative tightening policy that acts when interest rates are high). This increased volatility results in a slightly higher term premium in the model, though its effects are quantitatively small.

This observation implies that even if purchases respond to the short rate in practice, this dependency is unlikely to explain the large impact of QE that we measure empirically. Second, it also implies that asset purchase rules moving only with the short rate, the target of conventional monetary policy, are ineffective at altering long-term interest rates.

5. Model Extensions

Our baseline model is intentionally simple to transparently understand and quantify the impact of asset purchase rules. However, it also leaves out realistic factors in yield dynamics as well as additional factors such as market functioning that the central bank might want to respond to through purchases. We outline how one can incorporate sev-

eral extensions of the model with this in mind.

5.1 Demand Factors

There is growing evidence that demand shocks play a critically important role in treasury markets. To allow for such shocks we augment the inelastic investors demand for government bonds to $\zeta_t = \bar{\zeta} + \delta_\zeta Z_t$, where fundamentals Z_t now augmented with a demand state follows a VAR(1) process. The arbitrageur now absorb the demand-induced imbalances with the resulting variation showing up in term premia (for maturity τ):

$$term_t^\tau = \gamma B(\tau)\Sigma B'(w_s s_t - w_\zeta \zeta_t - w_{qe} qe_t) - r_t.$$

Term premia therefore reveals the demand-supply imbalances. A policy maker that wants to neutralize such demand effects from bond prices might choose to exactly offset demand-induced quantity variation, potentially having to not only trade at quite high frequency, but also needing to have direct knowledge about these demand shocks.[13]

5.2 A Term Premium Based Rule

In this section, we show how a rule that responds only to term-premia, and therefore requires knowledge only of term premia and not of its specific drivers, can stabilize bond prices with respect to these demand imbalances. Further, in our baseline model a natural interpretation of the rule loading on supply is simply that supply drives term premia variation, and policy makers are responding to the tightening of financial conditions that

[13] Expressed in terms of the state variable we have $\overline{term} + \delta_{term} Z_t$ where

$$\begin{aligned} \overline{term} &= \gamma B(\tau)\Sigma B'(w_s \bar{s} - w_\zeta \bar{\zeta} - w_{qe}\overline{qe}) - \bar{r}, \\ \delta_{term} &= \gamma B(\tau)\Sigma B'((w_s \delta_s - w_\zeta \delta_\zeta - w_{qe}\delta_{qe}) - \delta_r \end{aligned}$$

result from an increase in term premia, rather than supply per se. Consider a policy rule of the following form

$$qe_t^d = \overline{qe} + \delta_{qe} term_t,$$

If we assume the mix of long-term bonds the government issues, investors demand, and the fed targets are the same, then $w_s = w_{qe} = w_\zeta$ the equilibrium term premium follows a dampened version of the laissez faire term premium dynamics, with it's variance being a fraction of the original dynamics, and a policy maker can choose δ_{qe} to make sure the term-premium does not deviate too much from it's desired level. For example, if the policy maker aims at reducing the term premia variance from V^0 to V^* he would set

$$\delta_{term} = \frac{1}{\gamma B(\tau)\Sigma B' w}\left(\sqrt{\frac{V^0}{V^*}} - 1\right).$$

Note here that the bond sensitivity to state variables B is endogenous to the choice of rule. A reduction in term premia volatility feeds-back into lower sensitivities, which in turn means that the average positions $\overline{qe}$ the fed must hold to hit a desired level of risk-premia is lower

$$a_{qe} = \bar{s} - \bar{\zeta} - \frac{E[term_t]}{\gamma B(\tau)\Sigma B' w},$$

Or conversely, if the policy maker would like to have a zero balance-sheet on average, such conditionality will imply a lower average term-premium.

5.3 Crisis-driven Supply Shocks

A possible interpretation of the asset purchase interventions since 2008 is that the Fed and central banks around the world accommodated the large increase in bond issuance because they were driven by large macroeconomic shocks (e.g. 2008 financial crisis, 2012

European debt crisis, Covid-19), but would have made different choices if the expansion was driven by an active expansion of health and aging related entitlements. Under this view, it is important to distinguish these crisis-driven supply surges from general increases in borrowing. To capture this intuition we augment our state vector to now include two supply states, where aggregate supply is $s_t = \bar{s} + \delta_g Z_t + \delta_c Z_t$.

A policy maker concerned with the political economy of accommodating general increases in government borrowing might only respond to the term-premia variation cleaned up of such general supply effects,

$$qe_t^d = \overline{qe} + \delta_{qe}(term_t - \gamma B(\tau)\Sigma B' w \delta_g Z_t).$$

One can imagine how such a rule could be more desirable once political economy considerations are taken into account. Specifically, such rule implies that the government would internalize any increase in term-premia due to a politically-driven increase in borrowing. Understanding the particulars of the rule would also be important to forecast how the bond market will respond to these different types of fiscal shocks going forward.

5.4 Market Malfunction Shocks

More recently, there has been a focus on disruption in government bond markets (Duffie, 2023) following the extreme volatility in US Treasuries, and international bond markets more broadly, at the start of of COVID-19 crisis, and also in Gilts following the 2022 UK mini-budget debacle. In both cases, the respective central banks responded by quickly purchasing large quantity of long-term bonds. The UK case is particularly noteworthy as the Bank of England decided to purchase Gilts aggressively even when this action was in direct tension with their objective of tightening policy to lean against inflation. Market malfunction has been pointed to as a key driver of such episodes and it has been argued

that such shocks add a new rational for asset purchase policies.

To explore these ideas in our setting we augment the model to include a market malfunction shock $\nu_t \in \{0, \nu\}$, which we model following He et al. (2022). The idea is that when dealers cannot intermediate effectively, so that when inelastic investors sell orders, that would otherwise be accommodated by other inelastic investors buy orders, now have to be absorbed by dealers and their equilibrium balance-sheet now includes these bond sales ν_t,

$$X_t = (w_s s_t - w_\zeta \zeta_t + w_\zeta \nu_t - w_{qe} qe_t).$$

We model such malfunction as a low probability event – we assume the market malfunction shock arrives with probability p_ν and recovers with probability $p_{-\nu}$. This sudden, but persistent, inability of intermediaries to match sellers with buyers translates into price effects as a standard, but sudden, demand shock ζ_t.

It is easy to see, as we show in the appendix, that a policy contingent on ν_t can completely neutralize the ex-ante and ex-post effects of this shock which we show that can be meaningful for reasonable calibrations. Consistent with the standard intuition such policy has exactly the same persistence as the market malfunction shocks themselves. Because such shocks are typically described as acute but fairly transitory, they predict a similarly quickly unwinding of central bank purchases. We did not see this happening in the US COVID-19 experience, but is consistent with the UK case when the government bought about 20 billion pounds during the two-week period following the mini-budget announcement, but sold these assets back to investors within 3 months of the intervention.

6. Conclusion

We argue that quantitative easing (QE) and tightening policies constitute a dynamic state-contingent plan instead of a succession of spot decisions. This dynamic view adds an insurance benefit to the static effect of absorbing bond supply. Purchasing long-term bonds when the term premium becomes abnormally large not only achieves ex-post stabilization but also makes bonds safer today, reducing long-term rates. This channel explains the prevalence of low yields, low term premia, and high Treasury convenience yields since the introduction of QE despite the sharp increase in net government debt supply. It also rationalizes the strong market response to initial QE announcements — and weak response to later ones — and why option markets at these announcements price in substantial downside support. We calibrate a policy rule for asset purchases to their historical path and include it in a quantitative term structure model. In the model, state-contingent QE results in substantial stabilization ex-post, and the safety provided by this insurance lowers long-term Treasury yields by an additional 75bps ex-ante.

References

Acharya, Viral V, and Toomas Laarits, 2023, When do treasuries earn the convenience yield?: A hedging perspective, Working paper, National Bureau of Economic Research.

Bianchi, Javier, and Saki Bigio, 2022, Banks, liquidity management, and monetary policy, *Econometrica* 90, 391–454.

Curdia, Vasco, and Michael Woodford, 2011, The central-bank balance sheet as an instrument of monetary policy, *Journal of Monetary Economics* 58, 54–79.

Di Maggio, Marco, Amir Kermani, and Christopher J Palmer, 2020, How quantitative easing works: Evidence on the refinancing channel, *The Review of Economic Studies* 87, 1498–1528.

Diamond, William, and Peter Van Tassel, 2021, Risk-free rates and convenience yields

around the world, *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper* .

Duffie, Darrell, 2023, Resilience redux in the us treasury market, in *Jackson Hole Symposium, Federal Reserve Bank of Kansas City*.

Eser, Fabian, and Bernd Schwaab, 2016, Evaluating the impact of unconventional monetary policy measures: Empirical evidence from the ecb's securities markets programme, *Journal of Financial Economics* 119, 147–167.

Foley-Fisher, Nathan, Rodney Ramcharan, and Edison Yu, 2016, The impact of unconventional monetary policy on firm financing constraints: Evidence from the maturity extension program, *Journal of Financial Economics* 122, 409–429.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack, 2018, The financial market effects of the federal reserve's large-scale asset purchases, *24th Issue (Mar 2011) of the International Journal of Central Banking* .

Gertler, Mark, and Peter Karadi, 2011, A model of unconventional monetary policy, *Journal of monetary Economics* 58, 17–34.

Gertler, Mark, and Peter Karadi, 2018, Qe 1 vs. 2 vs. 3...: A framework for analyzing large-scale asset purchases as a monetary policy tool, *29th issue (January 2013) of the International Journal of Central Banking* .

Greenwood, Robin, Samuel Hanson, and Dimitri Vayanos, 2023, Supply and demand and the term structure of interest rates, Working paper, National Bureau of Economic Research.

Greenwood, Robin, and Dimitri Vayanos, 2014, Bond supply and excess bond returns, *The Review of Financial Studies* 27, 663–713.

Haddad, Valentin, Alan Moreira, and Tyler Muir, 2023, Whatever it takes? the impact of conditional policy promises, Working paper, National Bureau of Economic Research.

Haddad, Valentin, and David Sraer, 2020, The banking view of bond risk premia, *The Journal of Finance* 75, 2465–2502.

Hanson, Samuel G, and Jeremy C Stein, 2015, Monetary policy and long-term real rates, *Journal of Financial Economics* 115, 429–448.

He, Zhiguo, Stefan Nagel, and Zhaogang Song, 2022, Treasury inconvenience yields during the covid-19 crisis, *Journal of Financial Economics* 143, 57–79.

Joyce, Michael, Ana Lasaosa, Ibrahim Stevens, and Matthew Tong, 2010, The financial market impact of quantitative easing .

Joyce, Michael AS, and Matthew Tong, 2012, Qe and the gilt market: a disaggregated analysis, *The Economic Journal* 122, F348–F384.

King, Thomas B, 2019, Expectation and duration at the effective lower bound, *Journal of Financial Economics* 134, 736–760.

Krishnamurthy, Arvind, Stefan Nagel, and Annette Vissing-Jorgensen, 2018, Ecb policies involving government bond purchases: Impact and channels, *Review of Finance* 22, 1–44.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2012, The aggregate demand for treasury debt, *Journal of Political Economy* 120, 233–267.

Luck, Stephan, and Tom Zimmermann, 2020, Employment effects of unconventional monetary policy: Evidence from qe, *Journal of Financial Economics* 135, 678–703.

Meaning, Jack, and Feng Zhu, 2011, The impact of recent central bank asset purchase programmes, *BIS Quarterly Review, December* .

Mertens, Thomas M, and John C Williams, 2021, What to expect from the lower bound on interest rates: Evidence from derivatives prices, *American Economic Review* 111, 2473–2505.

Rodnyansky, Alexander, and Olivier M Darmouni, 2017, The effects of quantitative easing on bank lending behavior, *The Review of Financial Studies* 30, 3858–3887.

Sims, Eric, Jing Cynthia Wu, and Ji Zhang, 2023, The four-equation new keynesian model, *Review of Economics and Statistics* 105, 931–947.

Van Binsbergen, Jules H, William F Diamond, and Marco Grotteria, 2022, Risk-free interest rates, *Journal of Financial Economics* 143, 1–29.

Vayanos, Dimitri, and Jean-Luc Vila, 2021, A preferred-habitat model of the term structure of interest rates, *Econometrica* 89, 77–112.

Vissing-Jorgensen, Annette, 2021, The treasury market in spring 2020 and the response of the federal reserve, *Journal of Monetary Economics* 124, 19–47.

Vissing-Jorgensen, Annette, and Arvind Krishnamurthy, 2011, The effects of quantitative easing on interest rates: Channels and implications for policy, *Brookings Papers on Economic Activity* 43, 215–287.

Wu, Jing Cynthia, and Fan Dora Xia, 2016, Measuring the macroeconomic impact of monetary policy at the zero lower bound, *Journal of Money, Credit and Banking* 48, 253–291.

A. Data Appendix

We detail our main data sources here.

Aaa-Treasury Spread, Commercial Paper T-Bill Spread, and Debt to GDP: Our calculation follows exactly the work of Krishnamurthy and Vissing-Jorgensen (2012) (KVJ) in construction. See their data appendix for additional details. All data are from FRED. Aaa-Treas spread is the series AAA minus LTGOVTBD. Starting in the year 2000 we replace LTGOVTBD with GS20 (the yield on a 20 year Treasury) as the LTGOVTBD series is discontinued. The CP-TBill spread is uses CPN3M when available and CP3M when not available and subtracts TB3MS when available and TB6MS when not available. Debt to GDP is computed using market value of debt. While our construction follows KVJ, our observations are quarterly rather than annual, and the quarterly data series begin in 1947. Net debt to GDP is constructed by subtracting the size of the Fed balance sheet WALCL. This is somewhat of an overestimate of Fed holdings of Treasuries and is therefore a conservative assumption on the role that Fed balance sheet plays in reducing supply, for example since it includes holdings of Agency MBS. Using the series TREAST which only includes Treasury holdings makes the effect of realized purchases slightly weaker. In reality, since there is some substitutability between Agency MBS and Treasuries, we choose to make the assumption more conservative.

Slope of the Yield Curve: We use LTGOVTBD or GS20 based on availability as mentioned above as the long-rate. Results are similar when using the 10-year yield (GS10). We use TB3MS as the short rate as our data is quarterly. This makes the slope variable consistent with the same yields we use above in the Treasury convenience yield from KVJ. We also consider using the 10 year zero coupon bond GS10, as in Greenwood and Vayanos (2014), as the slope measure instead. We proxy for excess bond returns as follows: we compute the 1 year return on the long-term bond using 20*yield-19*(yield in 1 year). Since we don't have the yields of 19 year bonds, we are implicitly assuming the term structure is roughly flat between year 19 and year 20. Alternatively, we can compute the yield of a 19 year bond as a weighted average between the 10 and 20 year bond with weights that assume a linear yield curve between these points, but this has little effect. Annual excess returns are computed relative to the 1 year Treasury GS1. Greenwood and Vayanos (2014) find the same pattern of predictability across the entire maturity structure, suggesting our approximation mentioned above plays little to no role. As in Greenwood and Vayanos (2014), our regressions using these measures start in 1952, see their paper for details.

US QE Announcements: We follow the literature on QE event days for QE1 and QE2 (see, for example, Vissing-Jorgensen and Krishnamurthy (2011) and Gagnon et al. (2018)). First, we note that almost all action in yields occurs on during QE1, when QE is first implemented. The dates includes are as follows: November 25th, 2008; December 1st, 2008; December 16th, 2008; January 28th, 2008; March 18th 2009; August 10th, 2010; Septem-

ber 21st, 2010, November 3rd, 2010. Finally, we include the "QE infinity" date of March 16th, 2020. The inclusion of this date has a relatively minor impact on results. Using QE3 dates also produces similar results as the yield responses to these announcements are low. All event study methodology uses 2 day windows (1 day after the event relative to 1 day before) to allow for potential illiquidity in prices (yields). This follows the work of Vissing-Jorgensen and Krishnamurthy (2011). We use the daily yield series from Fred for these events, e.g., DGS1 DGS2 DGS5 DGS5 DGS7 DGS10 DGS20 DAAA DBAA. Volatility measures include the VIX (VIXCLS) and 10-year Treasury Note volatility (VXTYN). Data on swaptions comes from Bloomberg.

International Evidence: We use data from the Jorda Schularick Taylor Macrohistory database https://www.macrohistory.net/. This data contains annual observations for 18 advanced economies. We use annual data on debt to GDP, unemployment, and inflation. We supplement this with data from Fred on government bond yields when possible. For example, for the UK, the series IRLTLT01GBM156N provides 10 year government bond yields in the UK and the series IR3TTS01GBM156N provides 3-month government bond yields for the UK. Similar series exist for Japan, Canada, Australia, Italy, Portugal, Germany, France, Belgium, Denmark, Spain, Finland, Ireland, the Netherlands, and Sweden. These series are monthly. We use end of year values from Fred to match to the annual data from JST. For the US, we use our earlier measure of slope using end of year values. For Euro area countries, we use the 30 day interbank rate as the short rate (IR3TIB01EZM156N) and subtract this rate from all Euro country government yields to compute the slope.

Next, we need measures of the "post QE" period for each country in our sample. We define the post dummy as equal to 1 the year after the first phase of QE was implemented in a given country. For the US, we use the QE1 period beginning in 2008 through March of 2009. For Japan we use 2001, for the UK 2009, and for Canada and Australia 2020.The trickiest issue is how to handle the Euro zone. We proceed as follows: early asset purchase interventions by the ECB from 2010-2012 are typically not considered "QE" as they were targeted at reducing sovereign yields of specific countries – the GIIPS countries (see Krishnamurthy et al. (2018) for details on these announcements). Nevertheless, these purchase programs fit the asset purchase assumptions in our paper for the set of countries they targeted. There were multiple announcements and programs aimed at specific countries from 2010-2012, with important announcements in 2012 in terms of GIIPS yield responses (e.g., the Draghi speech in August 2012 and the OMT program in September 2012, which were broad based). We therefore use 2012 as the date for all GIIPS countries (Greece, Ireland, Italy, Portugal, and Spain). For the remainder of Euro area countries, which were not targeted by these interventions, we use 2015. This coincides with the ECB's Expanded Asset Purchase Programme which is more commonly known as quantitative easing, or QE, in the Euro zone and expanded purchases to euro-denominated, investment-grade securities issued by euro area governments and institutions.

Table 10: Slope of the Yield Curve controlling for shadow rate. We regress the slope of the yield curve on log debt to GDP, but also control for the shadow rate, the Taylor-rule implied interest rate absent a zero lower bound.

Notes.

	Slope of Yield Curve				Predicting Excess Bond Return			
	(1) Pre	(2) Full	(3) Full	(4) Full	(5) Pre	(6) Full	(7) Full	(8) Ful
log(Debt/GDP)	1.89**	0.03	1.24*	1.93***	0.24***	0.08	0.15*	0.24***
	(0.73)	(0.54)	(0.67)	(0.69)	(0.08)	(0.06)	(0.08)	(0.07)
Post 2008 Dummy			-1.63**	16.32***			-0.10	2.30***
			(0.64)	(5.93)			(0.08)	(0.83)
log(Debt/GDP) × Post				-4.29***				-0.57***
				(1.44)				(0.20)
(Shadow) Short Rate	-0.24***	-0.25***	-0.25***	-0.20***	0.01	0.01	0.01	0.01*
	(0.06)	(0.07)	(0.06)	(0.06)	(0.01)	(0.01)	(0.01)	(0.01)
Unemp	0.60***	0.40***	0.46***	0.41***	0.01	0.01	0.01	0.01
	(0.09)	(0.08)	(0.07)	(0.07)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	227	288	288	288	227	284	284	284
R-squared	0.63	0.49	0.56	0.62	0.09	0.03	0.04	0.11

B. Additional Tables

B.1 Robustness with respect to the Zero Lower bound

We've already highlighted that one would need to believe the ZLB would be binding for very long periods of time to explain the data – up to many years. For example, our results on the slope of the yield curve found a low slope when using the 20 year minus 2 to 5 year yield, so a ZLB story would need a binding ZLB beyond years 2 to 5.

To address further address this, we look at Fed funds futures markets – markets that specifically bet on the path of the short-rate – and then use them to evaluate how much the ZLB mechanism could explain our results. The main result from Fed funds futures markets is that investors persistently expect "lift off" over relatively short horizons. That is, they do not expect the ZLB to bind for very long and expect rates to perpetually return to values within 1-3% over 1-2 year horizons.

Fed funds futures are contracts on what the fed funds rate will be T periods in the future. Below we plot the time-series of the implied term-structure of Fed funds rates. For example, a line that starts in March-2009 tells us the March-2009 Q-expectations of the Fed funds rate for different dates in the future. Thus the plot reveals how quickly markets expected lift-off from the ZLB. For example, in late 2008, when rates first hit the

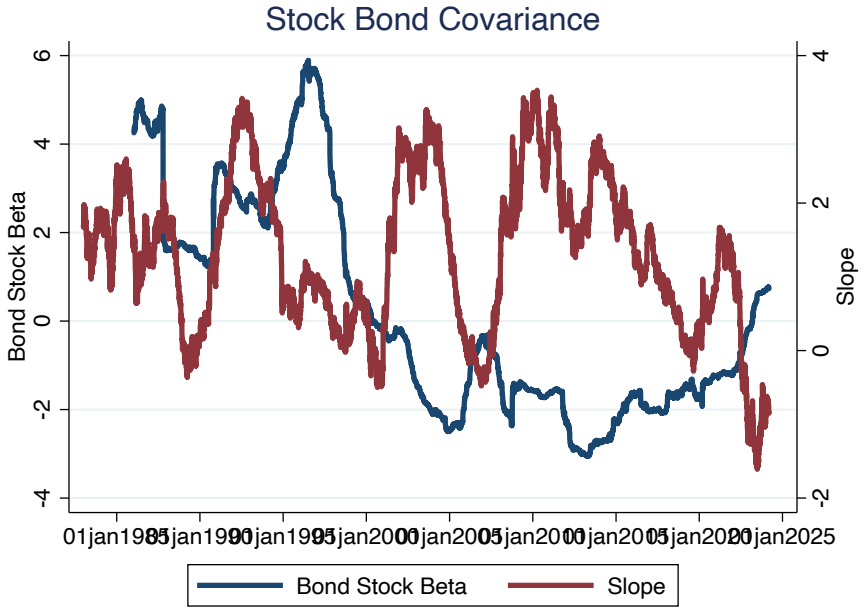


Figure 9: Stock bond covariance and slope of the term structure.
We regress daily bond returns on the stock market and plot the rolling 3-year beta against the slope of the term structure. We use daily data on 10 year Treasury yields and stock market excess returns from Ken French and compute the negative change in 10 year Treasury as a measure of daily bond returns. We use the 10 year yield minus 1 year yield to compute slope, so that the reference 10 year is the same as used in the stock-bond regressions.

ZLB, markets expected the Fed funds rate to be at 1.25% by the end of 2009, just one year ahead. This implies a fairly large probability of lift off. If the two possibilities were 2% or 0.25% by the end of the year, this would imply a liftoff probability of 50%. Going six-months further would imply a probability of close to 90%.

The figure also shows that low rates proved to be much more persistent than initially expected. We also see that by late 2011, when QE1 and QE2 announcements have ceased, investors still expected a brisk increase in the short-rate. Only by early 2012 we see the curve flatten in a way that is more consistent with a persistent low rate environment. Consistent with this pattern, Mertens and Williams (2021) use interest rate cap contracts to construct the risk-neutral distribution of short rates as far as seven years ahead and find that during this period the probability of short rates being at the zero lower bound hovered between 2.5% and 5%. Thus, this channel can not account for the large decline in long-term yields or the declines in long-term volatility we observed around QE1 announcements. Specifically, our results show that the expected volatility for the next ten

years of the 10x10 forward rates drops right at the QE1 announcements by roughly 40%. The zero lower bound mechanism could only account for this fact if the probability of the zero lower bound binding ten to twenty years from now had dramatically increased around these events. Neither the Feds Future data or the interest-rate cap data suggests this was the case.

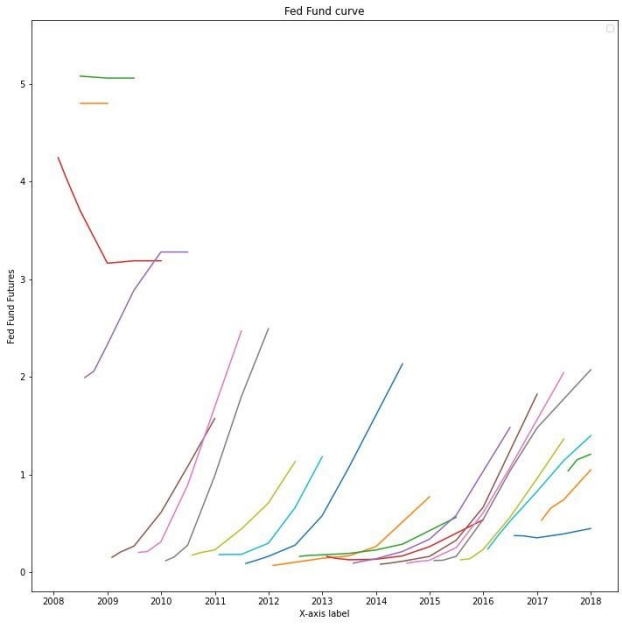


Figure 10: Term Structure of Fed funds rates.
Each line shows the fed fund future implied risk-neutral expectations of the fed fund rate at a particular date in the future from the perspective of the date at the start of the line.

The behavior on bond prices are driven by expectations of future rates and as we see in the figure these expectations were not consistent with the view that the ZLB was a persistent state and therefore it would be challenging for this view to explain the drop in long term bond yields or volatility around early QE announcements.

B.2 Convenience Yields: Additional Evidence

Table 11 shows these results with a regression of the Aaa-Treasury spread, or convenience yield, on the two different proxies for the US debt supply. Column (1) confirms the nega-

Table 11: KVJ Regression Updates

Notes. Quarterly data, sample period 1946-2023. The pre sample uses data before 2008 only. Standard errors are Newey-West with 20 lags in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
VARIABLES	Pre 2008	Full	Full	Full	Full	Full	Full
ln(Debt/GDP)	-0.93***	-0.22	-0.90***	-0.93***	-0.30	-0.91***	-0.93***
	(0.20)	(0.19)	(0.18)	(0.19)	(0.21)	(0.18)	(0.19)
Post 2008 Dummy			1.02***	-0.41		0.93***	-0.56
			(0.20)	(2.26)		(0.18)	(2.46)
ln(Debt/GDP) × Post				0.34			0.36
				(0.54)			(0.61)
Observations	247	308	308	308	303	303	303
R-squared	0.43	0.05	0.44	0.45	0.07	0.46	0.46

tive pattern in the pre-2008 data and matches the results in Krishnamurthy and Vissing-Jorgensen (2012). Economically, a 10% increase in the supply of Treasuries to GDP is associated with a 9 basis point decline in convenience yields. Column (2) shows that this result does not hold when extending the sample to the present. Column (3) includes a post 2008 dummy. The coefficient is about 1 and is highly statistically significant. This accords with a level shift see in the red dots in Figure 4: Treasury yields are around 100 bps too low in the post 2008 data. Columns (5)-(7) repeat this analysis using net supply that subtracts holdings by the Federal Reserve and comes to similar conclusions. These conclusions hold when controlling for a host of additional factors including market volatility or the short term interest rate (see Appendix).

Figure 11 shows visually in the time-series these results: specifically, we use predicted values from the regression in column (1) applied to the post 2008 data. Right in 2008 a large wedge emerged between supply-predicted convenience yields and actual convenience yields. The additional lines in this figure include macroeconomic controls such as unemployment or inflation show that this does little to change the results. As before, accounting for the reduction in supply due to Fed purchases helps very little. A more explicit way to state the result is: given the past relationship between bond supply and convenience yields, and the very large expansion in bond supply, the large-scale asset purchase programs were too small to account for the level of convenience yields. Implicit in these regressions is that the slope of demand curve for Treasuries stayed the same across this period. One concern is that post crisis regulation, or other changes in the economy, may be responsible for a higher demand for Treasuries.

Table 12 looks at the spread in the short-end of the treasury curve, the spread between

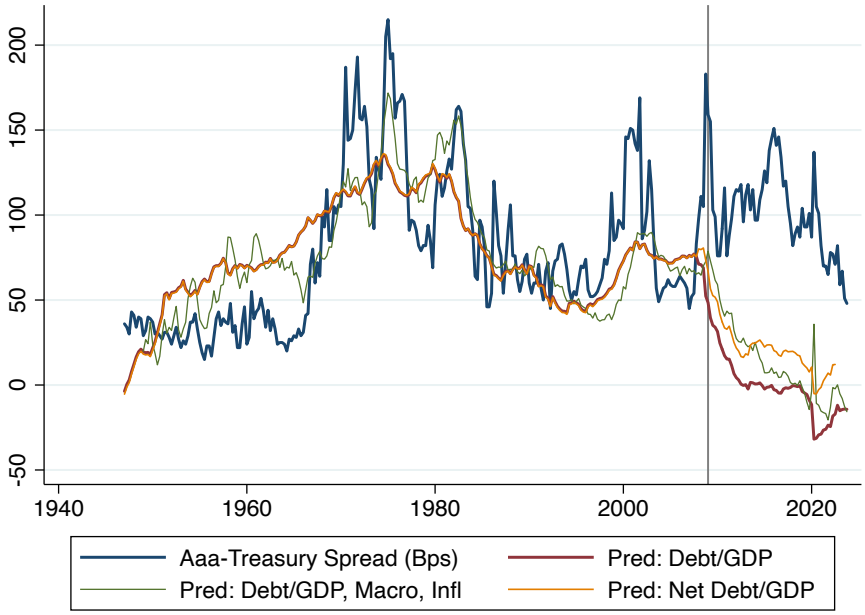


Figure 11: Predicted Spreads.
Predicted Aaa-Treasury spread using pre 2008 regressions, adding controls.

commercial paper and Treasury bills. Focusing again on the post-2008 dummy we see that now there is no puzzle. No excess convenience yield for bills. This is informative because it tell us this is something about long-term bonds in particular and therefore points to the large-scale asset purchase programs which were initiated in 2008 and targeted long-term bonds. It cuts against "financial repression" and regulation driven demand shifts for government bonds since if anything Treasury bills have been treated more favorably in the post-2008 bank regulatory framework. The next subsection explores this difference in maturity by analyzing the slope of the yield curve.

The difference with the short-end of the yield curve makes sense from the perspective of the convenience yield literature if there is some segmentation in convenience yields for very short term vs long-term bonds. Van Binsbergen et al. (2022) find a decline in convenience yields at relatively shorter horizons following QE announcements.[14] This makes sense because QE operations inject reserves which are highly liquid short maturity assets, and the increase in supply should lower convenience yields on safe liquid assets with short maturities. Our finding suggests this is not the case at very long maturities (recall the Aaa-Treasury spread uses a maturity around 20 years), suggesting there is a

[14]See also Diamond and Van Tassel (2021).

Table 12: KVJ Regression Updates: CP Bill Spread.

Notes.

VARIABLES	(1) Aaa-Trs Pre	(2) CPbill Pre	(3) Aaa-Trs	(4) CPbill	(5) Aaa-Trs	(6) CPbill
ln(Debt/GDP)	-0.93***	-0.76***	-0.22	-0.57***	-0.90***	-0.69***
	(0.20)	(0.24)	(0.19)	(0.10)	(0.18)	(0.19)
Post 2008 Dummy					1.02***	0.15
					(0.20)	(0.21)
Observations	247	148	308	208	308	208
R-squared	0.43	0.18	0.05	0.30	0.44	0.31

Table 13: Spreads on QE Announcement Dates

Notes. Units are in basis points.

VARIABLES	(1) Δ(Aaa-Tr)	(2) Δ(Baa-Tr)	(3) Δ CPBill	(4) Δ Treas
QE Event	58.83***	45.98***	34.37	-107.82***
	(13.17)	(13.15)	(31.54)	(23.62)
Observations	4,037	4,037	3,680	4,037
R-squared	0.00	0.00	0.00	0.01

special convenience demand for long-term safe assets.

As before, we show that these broad patterns in safe corporate bond Treasury spreads match evidence from specific QE announcement. Table 13 shows the cumulative effects of these announcements: we divide the dummy for these by the total number of events so that the coefficient represents a cumulative effects across all events. First, we see that the spreads in the long end of the curve are highly significant and quantitatively large. Relative to corporate bonds, long-term treasury yields were lowered by around 60 bps. Also in line with our low frequency analysis we see no effect on the short end of the curve. Column 5 focus on simple changes in treasury yields, where we take an average of maturities between 10 and 30 years.

B.3 Details on Model Solution

We conjecture that log bond prices are given by

$$p_{\tau,t} = A(\tau) + B(\tau)Z_t. \tag{14}$$

It immediately follows that

$$\begin{aligned}
r_{\tau,t+1} &= p_{\tau-1,t+1} - p_{\tau,t} = A(\tau-1) + B(\tau-1)Z_{t+1} - (A(\tau) + B(\tau)Z_t), \\
\text{Var}_t(r_{\tau,t+1}) &= B(\tau-1)\Sigma B(\tau-1)', \\
E_t(r_{\tau,t+1}) &= A(\tau-1) - A(\tau) + (B(\tau-1)\Gamma - B(\tau))Z_t, \\
r_{\tau,t+1} - E_t[r_{\tau,t+1}] &= B(\tau-1)Z_{t+1} - B(\tau-1)\Gamma Z_t = B(\tau-1)\epsilon_{t+1},
\end{aligned}$$

Plugging this conjecture into the arbitrageurs' FOC and substituting the risk-free rate $r_t = \bar{r} + \delta_r Z_t$ and the equilibrium condition on the arbitrageur's portfolio $X_t = w\bar{s} + \delta_s Z_t - w_\zeta(\bar{\zeta} + \delta_\zeta Z_t)$ gives:

$$\begin{aligned}
A(\tau-1) - A(\tau) + (B(\tau-1)\Gamma - B(\tau))Z_t + \frac{1}{2}B(\tau-1)\Sigma B(\tau-1)' \\
= \bar{r} + \delta_r Z_t + \gamma B(\tau-1)\Sigma \sum_{\tau=2}^{\tau^+} \Big(B(\tau-1) \\
\times \left(w(\tau)\bar{s} - w_\zeta(\tau)\bar{\zeta} + \left(w(\tau)\delta_s - w_\zeta(\tau)\delta_\zeta \right) Z_t \right) \Big),
\end{aligned}$$

Because the above expression has to hold for any value of the state variable Z_t, the above equation will be solved if and only if the coefficients A and B satisfy the following difference equations shown in Equation 7 in the main text.

$$B(\tau) = -\delta_r + B(\tau-1)(\Gamma - \gamma\Sigma B' w_s \delta_s), \tag{15}$$

$$A(\tau) = -\bar{r} + A(\tau-1) - \gamma B(\tau-1)\Sigma B'(w_s\bar{s} - w_\zeta\bar{\zeta}) + \frac{1}{2}B(\tau-1)\Sigma B(\tau-1)', \tag{16}$$

A. Model Appendix

A.1 Market Malfunction

We model market malfunction by augmenting our model with a market functioning state ν_t. We model this state as follows. When the market stop functioning a mass of sellers cannot find buyers, due to some breakdown in intermediation, and investors have to absorb an additional amount $\nu_t = \nu$) of government bonds. When markets are functioning well, sellers get matched with buyer and that extra demand imbalance disappears. We

model this transition across market regimes with probability $p(\nu_t)$ of transitioning across states and probability $1 - p(\nu_t)$ that stays in the same state.

We keep the model solution affine by making the simplifying assumption that $(1 - p(\nu))p(\nu) = (1 - p(0))p(0)$. Given this assumption, prices are still affine in the state-variables Z, ν,

$$p(\tau, \nu_t, Z_t) = A(\tau, \nu_t) + B(\tau) Z_t, \tag{17}$$

and the price function coefficients solve the following system of coupled difference equations:

$$\begin{aligned}
B(\tau) &= B(\tau-1)\Gamma - \delta_r - \gamma \left(\sigma_p C(\tau-1) D_W + B(\tau-1)\Sigma B_w\right) \delta_x & (18)\\
A(\tau,\nu) &= A(\tau-1,\nu) - p(\nu)C(\tau-1) + \frac{1}{2}\left(\sigma_p C(\tau-1)^2 + B(\tau-1)\Sigma\Sigma B(\tau-1)\right) - \bar{r} \\
&+ \gamma \left(\sigma_p C(\tau-1) D_W + B(\tau-1)\Sigma B_w\right)(\bar{x} + \nu) & (19)\\
A(\tau,0) &= A(\tau-1,0) + p(0)C(\tau-1) + \frac{1}{2}\left(\sigma_p C(\tau-1)^2 + B(\tau-1)\Sigma\Sigma B(\tau-1)\right) - \bar{r} \\
&+ \gamma \left(\sigma_p C(\tau-1) D_W + B(\tau-1)\Sigma B_w\right)\bar{x}, & (20)
\end{aligned}$$

where $C(\tau) = A(\tau, \nu) - A(\tau, 0)$ and $\sigma_p = p(0)(1 - p(0))$