Unbundling Quantitative Easing: Taking a Cue from Treasury Auctions

Walker Ray

London School of Economics

Michael Droste

Harvard University

Yuriy Gorodnichenko

University of California Berkeley, National Bureau of Economic Research, and Center for Economic Policy Research

We study the role of preferred habitat in understanding the economic effects of the Federal Reserve's quantitative easing (QE). Using high-frequency identification and exploiting the structure of the primary market for US Treasuries, we isolate demand shocks that are transmitted solely through preferred habitat channels but otherwise mimic QE shocks. We document large localized yield curve effects when financial markets are disrupted. Our calibrated model, which embeds preferred habitat in a New Keynesian framework, can largely account for the observed financial effects of QE. QE is modestly stimulative for output and inflation, but alternative policy designs can generate stronger effects.

I. Introduction

While evaluating the first rounds of quantitative easing (QE), then Fed chair Ben Bernanke observed, "The problem with QE is it works in practice but

We thank Harald Uhlig, anonymous referees, Michael Bauer, Rhys Bidder, Michael Fleming, Ed Knotek, Matteo Maggiori, Eric Swanson, Michael Weber, and Milena Wittwer as well

Electronically published July 11, 2024

Journal of Political Economy, volume 132, number 9, September 2024.

https://doi.org/10.1086/729581

it doesn't work in theory." Indeed, QE was successful in reducing short- and long-term interest rates, but the mechanisms behind these effects are still not well understood. Nevertheless, this theoretical ambiguity has not stopped the Fed from continuing to utilize QE programs, including during the onset of COVID-19 and then reversing course by implementing quantitative tightening (QT) in response to recent inflationary pressures.

Although workhorse macroeconomic models imply that Treasury demand is determined solely by economic agents' intertemporal consumption decisions, several explanations have been put forth to rationalize the workings of QE. For instance, QE could have signaled to financial markets a commitment to keep short-term interest rates low for a long time (forward guidance). Or perhaps the Fed exploited financial market frictions (limited arbitrage and market segmentation) by purchasing securities in a particular segment. Alternatively, large-scale asset purchases by the Fed could signal a poor state of the economy, pushing interest rates down (Delphic effect).[1] Given the paucity of QE events, it has been difficult to provide clear empirical evidence for (and assess the relative contributions of) the proposed channels. Moreover, QE policies have been implemented as responses to severe macrofinancial conditions, which further confounds identification and interpretation and raises questions about the effectiveness of QT going forward.

The objective of this paper is to unbundle the effects of QE by focusing on a specific channel: market segmentation in the form of *preferred habitat*, which posits that certain investors have preferences for specific maturities. To this end, we take the following approach. First, we identify shifts in private demand for Treasuries that mimic QE but are independent of all other plausible channels of QE. Next, we analyze the propagation of these demand shocks across financial markets. In particular, we assess the ability of preferred habitat theory to rationalize the observed responses, and we confirm key predictions regarding pass-through of demand shocks in and out of financial crises. Informed by our empirical analysis, we then develop a general equilibrium macroeconomic model designed to study QE policies. We find that the preferred habitat channel accounts for the bulk of the observed response to QE in financial markets. Our model suggests that the first rounds of QE had modest stimulative effects on output and inflation, but given the relative health of financial markets in the present period, QT alone is unlikely to be successful in reducing inflation to target. Finally,

as seminar participants at University of California Berkeley, Federal Reserve Bank of New York, Federal Reserve Bank of Cleveland, Federal Reserve Bank of San Francisco, Federal Reserve Bank of Chicago, Arizona State University, Dubrovnik Economic Conference, Claremont McKenna College, Bilkent University, George Washington University, Texas A&M University, and Northwestern University for comments on an early version of the paper. We are grateful to Maxime Sauzet for excellent research assistance. This paper was edited by Harald Uhlig.

[1] For surveys, see Campbell et al. (2012), Martin and Milas (2012), and Bhattarai and Neely (2020).

we explore alternative QE implementations that can help improve the design of future asset purchases.

Our analysis starts with the key insight that the mechanism through which market segmentation and preferred habitat forces operate is not the source of Treasury quantity shocks per se but rather how marginal investors absorb these shocks. We utilize the primary market for Treasuries to identify demand shifts that are independent of all QE propagation mechanisms besides preferred habitat. Although the primary market is the venue through which the Treasury issues debt (a supply-side action), the institutional structure of Treasury auctions has a number of desirable features for identifying demand-side shocks. Because all of the supply information is announced by the Treasury in advance of each auction, the release of the auction results reveal unexpected shifts in demand alone, allowing us to rule out a host of confounding factors. By utilizing intraday changes in Treasury yields around the close of Treasury auctions, we construct a novel measure of Treasury demand shocks.

We document that demand shocks are reasonably large and persistent, with effects on yields typically lasting for many weeks following the auction. Furthermore, the surprise movements in demand are driven by institutional investors, such as foreign monetary authorities, investment funds, insurance companies, and the like. We show that these shocks are not driven by market-wide changes in expectations about inflation, output, or other general macrofinancial conditions. Therefore, variation in Treasury yields around the release of Treasury auction results can help us to isolate the effect of idiosyncratic purchases in specific asset segments on the level and shape of the yield curve, which is difficult to achieve by examining only QE events. Because Treasury auctions are frequent and information spans many decades, we can study state dependence in the effect of targeted purchases of assets (e.g., crisis vs. noncrisis states), which is instrumental for understanding how QE programs can work in normal times. Importantly, because Treasury auctions for specific maturities are spread over time, we can identify changes in demand for government debt of specific maturities and trace the effect of these changes on other parts of the yield curve. In this sense, we have natural experiments that can mimic targeted purchases of the Fed during QE programs.

We use our auction demand shocks to empirically evaluate the localization hypothesis, a characteristic prediction of preferred habitat theory and a key input into our subsequent quantitative analyses. Using a simple regression specification informed by theory, we document strong evidence in favor of localized yield curve effects during financial crisis periods (i.e., the location of the demand shock in maturity space matters, and the effects on the yield curve are larger for bonds of similar maturities). These results imply potentially powerful effects of QE on targeted yields in crises. However, we cannot reject the null of no localization effects during normal times, which suggests a limited use of QE as a conventional policy tool.

Building on seminal work by Vayanos and Vila (2021), we develop a general equilibrium model with risk-free and risky debt to rationalize the empirical responses of the yield curve to shocks in demand for Treasury securities and to better understand the effects of QE on financial markets and the broader macroeconomy. We calibrate the model to match a variety of moments for yields and macroeconomic variables as well as the responses of yields to surprise movements in private demand during Treasury auctions in crisis and noncrisis times. When fed a shock mimicking QE1 in size and duration, our model generates movements in the yield curve remarkably close to those observed in the data. This is consistent with the view that QE works mainly via market segmentation and preferred habitat and that the net effect of other channels is small. Given the disruption in financial markets, our calibration implies that QE1 stimulated output and inflation by as much as a 50–75 basis point rate cut.

Policy experiments in our model indicate that these relatively modest macroeconomic effects are sensitive to implementation details. In particular, holding securities on the balance sheet longer (and making this clear to markets on announcement) boosts the stimulative power of QE significantly. The expansionary effects of QE fall precipitously when undertaken during periods when bond markets are relatively healthy. We also show that QE may have unintended consequences: uncertainty surrounding the Fed's asset purchases may lead to excess macroeconomic volatility, thus calling for policy guidance. Finally, QE programs tilted toward risky assets (e.g., mortgage-backed securities [MBS]) are more effective in stimulating the economy, particularly when these assets are more volatile. Our results indicate that QE should remain in policy makers' toolkit but must be utilized with caution and realistic expectations. For example, while our model predicts that the Fed's ongoing QT program is disinflationary, the effects are similar to a 25–50 basis point rate hike and thus dwarfed by the current rise in inflation.

Our paper makes three primary contributions. Theoretically, we embed a financial model of the entire term structure of interest rates for risky and safe assets within a dynamic model of the macroeconomy and thus can provide an integrated analysis of QE.[2] This is important because previous studies of QE largely focus separately on either financial variables or macroeconomic variables. For example, Krishnamurthy and Vissing-Jorgensen (2012) and Chodorow-Reich (2014) study the effects of

[2] Ray (2019) develops additional analytical and normative results in a version of this model.

QE announcements on financial markets but do not quantify the macroeconomic effects of QE. Greenwood and Vayanos (2014), Greenwood, Hanson, and Vayanos (2016), King (2019a), Vayanos and Vila (2021), and related work explore the financial market implications of preferred habitat theory but are silent on any potential effects on output or inflation. On the other hand, recent developments in macroeconomic theory (e.g., Cúrdia and Woodford 2011; Gertler and Karadi 2011, 2013; Chen, Cúrdia, and Ferrero 2012; Carlstrom, Fuerst, and Paustian 2017; Ippolito, Ozdagli, and Perez-Orive 2018; Karadi and Nakov 2020; Sims and Wu 2020) concentrate on aggregate variables but are unable to capture the rich dynamics in bond markets that we document. Moreover, many of these theories rely on reserve requirements or moral hazard/enforcement constraints on banks as the key channel through which QE works. In contrast, we focus on the interaction of preferred habitat with limited risk-bearing capacity in bond markets, which we argue is a core mechanism behind QE effects. Our quantitative model merges these literatures: we utilize financial data and high-frequency identification to discipline our model, which we then use to quantify the macroeconomic effects of QE.

Empirically, we develop a novel, high-frequency measure of demand shocks for Treasuries by exploiting the institutional structure of Treasury auctions. A well-identified shock series is crucial for testing various macrofinance theories, and we hope our measure will be analyzed and extended in future work in a fashion similar to the high-frequency monetary shocks pioneered by Kuttner (2001), Bernanke and Kuttner (2005), Gürkaynak, Sack, and Wright (2007), and others. Our approach is a natural complement to existing empirical approaches estimating affine term structure models (Hamilton and Wu 2012; Kaminska and Zinna 2020) or demand systems (Koijen et al. 2021) using lower-frequency data.

Combining the insights of our model with our high-frequency demand shocks (and building on extensive research studying how QE purchases impacted the yield curve; e.g., Cahill et al. 2013; D'Amico and King 2013; Li and Wei 2013; King 2019b), we provide strong empirical evidence for state-dependent (i.e., crisis vs. noncrisis) localization effects in the spillovers across maturities and asset classes of surprise movements in Treasury demand. Put differently, we use QE-like events (rather than QE directly, in the spirit of Fieldhouse, Mertens, and Ravn [2018] and Di Maggio, Kermani, and Palmer [2020]) in order to exploit rich cross-sectional and time series variation to obtain sharp identification and precise estimates of demand shock spillovers. This new finding confirms one of the characteristic predictions of our model (and preferred habitat theory more generally), and thus our paper shows that these mechanisms are crucial in understanding the Treasury market and channels through which QE affects yields.

Our paper is related to a broad literature studying Treasury auctions, investigating how yields move around Treasury auctions (e.g., Lou, Yan, and Zhang 2013; Fleming and Liu 2016) and respond to variation in demand (e.g., Cammack 1991; Beetsma et al. 2016, 2018; Forest 2018). However, our focus is not Treasury auctions in and of themselves; rather, we exploit the institutional structure of the primary market for Treasuries to better understand the mechanisms behind QE. Relative to this early important work, we structurally link demand shocks identified from Treasury auctions to a general equilibrium preferred habitat model.

II. Data and Institutional Details

In this section, we describe the primary sources of our data and present basic statistics. First, we describe the US Treasury auctions for US government notes and bonds (coupon-bearing nominal securities). Second, we describe the details of the data regarding intraday secondary market Treasury prices.

A. *Primary Market for Treasury Securities*

The Treasury sells newly issued securities to the public on a regular basis through auctions. In recent years, 2-, 3-, 5- and 7-year notes are auctioned monthly. Ten-year notes and 30-year bonds are auctioned in February, May, August, and November, with reopenings in the other 8 months. The frequency of auctions has changed over time. For example, 30-year bonds were not issued between 1999 and 2006 and were issued only twice a year between 1993 and 1999, and 20-year bonds were auctioned in May 2020, the first time since 1986.[3]

There are two types of bids: noncompetitive and competitive. Noncompetitive bidders agree to accept the terms settled at the auction and are typically limited to $5 million per bidder. Competitive bidders submit the amount they would like to purchase, not exceeding 35% of the amount offered at auction, and the price (the interest rate) at which they would like to make the purchase.

Auction participants include primary dealers, other nonprimary brokers and dealers, investment funds (e.g., pension, hedge, mutual), insurance companies, depository institutions, foreign and international entities (governmental and private), the Federal Reserve System Open Market Account (SOMA), and individuals. These participants are classified into

[3] In our analysis, we exclude inflation-protected securities and floating rate notes because these securities have different structural arrangements than simple coupon-bearing nominal securities. We also exclude Treasury bills (zero coupon securities with maturity 1 year or less) because the QE programs mainly bought long-maturity nominal US government debt.

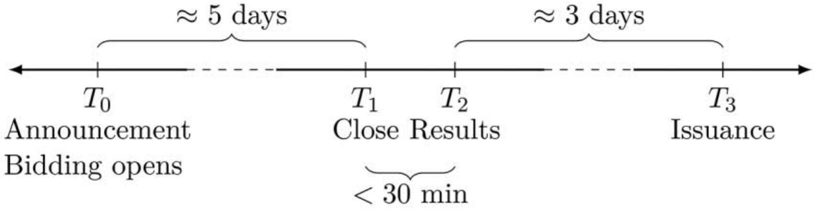

FIG. 1.—Auction timing.

three groups: primary dealers, direct bidders, and indirect bidders. Primary dealers (brokers and banks) trade on their account with the Federal Reserve Bank of New York; they are required to participate in every Treasury auction and typically buy the largest share of auctioned debt. Direct bidders are nonprimary dealers who submit bids for their own proprietary accounts. Indirect bidders submit competitive bids via a direct submitter, including foreign and international monetary authorities placing bids through the Federal Reserve Bank of New York. Additionally, the Treasury divides investors into the following classes: investment funds, pension and retirement funds and insurance companies, depository institutions, individuals, dealers and brokers, foreign and international, Federal Reserve System, and other.[4]

Figure 1 depicts the stages of a Treasury auction.[5] First, the Treasury releases all pertinent information regarding an upcoming auction a few days prior to the auction date. An announcement includes the amount offered, additional security information (e.g., maturity, CUSIP [Committee on Uniform Securities Identification Procedures], coupon schedule), and other information describing the rules of the auction. After the announcement, investors may submit bids up until the auction closing time. For typical note and bond auctions, noncompetitive bids may be submitted by 12:00 p.m., while the deadline for competitive bids is 1:00 p.m.

After the auction closes, competitive bids are accepted in ascending order (in terms of yields) until the quantity meets the amount offered minus the amount of noncompetitive bids. Winning bidders receive the same

[4] Data for announcement and results of each auction since late 1979 are available from https://treasurydirect.gov. Data regarding amounts accepted and tendered by bidder type (primary dealer, direct, and indirect) are available starting in 2003. The Treasury provides information regarding allotment by investor class starting in 2000 (for a breakdown by types and class of bidders in greater detail, see Fleming 2007). Appendix F (apps. B–F are available online) provides detailed information on all data sources.

[5] Figure B2 (figs. B1–B31 are available online) presents a typical announcement of an auction and its results.

yield as the highest accepted bid. Once the winning bids are determined, auction results are released immediately; beginning in the early 2000s, results are released within minutes of the close of the auction (see Garbade and Ingber 2005). Besides the winning yield, the Treasury announces various aggregate statistics regarding the bidding, such as the total demand (bids tendered) for the security and the composition of bids and winners by investor and bid type. One particularly salient piece of information is the bid-to-cover ratio, the ratio of all bids received to all bids accepted. A few days after the close of an auction, the Treasury delivers the securities and charges the winning bidders for payment of the security.

Panel A of table 1 presents summary statistics for note and bond auctions from 1995–2017, the period for which we have intraday Treasury yields (fig. B1 plots the number and size of note and bond auctions split by maturity). Since 1995, a typical offering of \$20 billion generates more than \$50 billion in demand. Primary dealers account for the largest source of demand (bid-to-cover ratio ≈2), but other types of bidders also account for a large fraction of auction offerings. Primary dealers purchase ≈60% of auctioned Treasuries, with the rest split between investment funds and foreign buyers. There is considerable variation in the offered amounts (standard deviation ≈\$9 billion) as well as the level and composition of demand (standard deviation of bid-to-cover ratio ≈0.5, and standard deviation of bid-to-cover ratio for primary dealers ≈0.35).

B. *Intraday Treasury Yields*

Once a Treasury security auction is complete, the security is issued to the winning bidders and the security is free to trade in the over-the-counter (OTC) secondary market. Following the announcement of an auction but before issuance, there is a forward market for newly auctioned Treasuries. The forward contracts mature on the same day as the securities are issued, and hence this market is referred to as the when-issued market. Our data on secondary market yields (including when-issued yields) comes from GovPX, which provides comprehensive intraday coverage of all outstanding US Treasuries for the period 1995–2017. We use changes in intraday Treasury yields in order to construct market-based measures of demand surprises occurring during Treasury auctions.[6]

[6] An early draft of this paper used Treasury futures to construct auction demand shocks. Treasury futures provide a natural market-based measure of such shifts, but in practice movements in the secondary market around auctions are not predictable. In addition, Treasury futures markets are less liquid and cannot be not tied to a specific CUSIP-level bond. Thus, we construct demand shocks using secondary market Treasury yields, but our results are robust to using futures (see Gorodnichenko and Ray 2017).

TABLE 1
AUCTION AND SHOCK SUMMARY STATISTICS

	Mean	Median	Standard Deviation	Minimum	Maximum	Observations
			A. Auction Summary Statistics			
Offering amount (US$ billions)	22.48	22.00	9.02	5.00	44.00	1,047
Total tendered (US$ billions)	61.98	57.91	29.91	11.35	160.96	1,047
Term (years)	8.23	5.00	8.68	2.00	30.25	1,047
High yield	3.01	2.68	1.83	.22	7.79	1,047
Bid-to-cover ratio	2.60	2.58	.45	1.22	4.07	1,047
Bid-to-cover ratio by type:[a]						
Direct bidders	.22	.21	.16	.00	.80	827
Indirect bidders	.54	.53	.17	.03	1.05	827
Primary dealers	1.92	1.87	.34	.98	3.12	827
Fraction accepted by type:[b]						
Depository institutions	.59	.15	2.09	.00	31.79	896
Individuals	.98	.23	2.13	.00	18.93	896
Dealers	53.44	53.07	16.22	12.08	97.73	896
Pensions	.11	.01	.74	.00	20.91	896
Investment funds	25.81	23.05	16.95	.23	73.91	896
Foreign	19.22	18.10	8.74	.34	61.01	896
Other	.25	.03	.91	.00	13.46	896
			B. Shock Summary Statistics			
D_t	.03	.00	2.12	−8.30	17.80	1,047
$D_t^{(2Y)}$	−.06	.00	1.39	−5.20	4.40	256
$D_t^{(3Y)}$	.15	.10	1.32	−5.30	5.00	134
$D_t^{(5Y)}$	.20	.30	1.94	−5.55	8.50	233
$D_t^{(7Y)}$	−.28	.05	2.10	−8.30	5.50	106
$D_t^{(10Y)}$	−.04	−.05	2.55	−6.80	12.20	186
$D_t^{(30Y)}$	.09	−.15	3.28	−7.40	17.80	132
			C. Shocks across Regimes			
Nonzero lower bound	.08	.10	2.02	−7.40	12.20	563
Zero lower bound	−.03	.00	2.23	−8.30	17.80	484
Expansion	.03	.00	2.03	−8.30	17.80	963
Recession	−.07	−.05	2.95	−7.10	9.50	84
			D. Nonauction Shocks			
$\tilde{D}_t^{(2Y)}$	−.07	.00	1.32	−19.20	11.00	1,225
$\tilde{D}_t^{(5Y)}$	−.01	.00	1.13	−17.40	12.90	1,924
$\tilde{D}_t^{(10Y)}$	−.01	.00	1.03	−13.30	8.90	2,151
$\tilde{D}_t^{(30Y)}$	.01	.00	.92	−8.50	5.60	1,914

NOTE.—Panel A presents summary statistics for Treasury note and bond auctions from 1995–2017 for which we have intraday data. Panel B summarizes demand shocks $D_t = y_{t,\text{post}} - y_{t,\text{pre}}$, the intraday change in Treasury yields before and after the close of an auction (in basis points). For newly issued securities, the yields are from when-issued trades. For reopened securities, the yields are from secondary market trades. Panel B also reports statistics for shocks $D_t^{(m)} = y_{t,\text{post}}^{(m)} - y_{t,\text{pre}}^{(m)}$ separately for major maturities $m = 2, 3, 5, 7, 10, 30$ years. Panel C reports statistics separately for binding/nonbinding zero lower bound periods and recessions/expansions. Panel D reports synthetic shocks on nonauction dates $\tilde{D}_t^{(m)} = y_{t,\text{post}}^{(m)} - y_{t,\text{pre}}^{(m)}$ using the same intraday windows around 1:00 p.m. on nonauction dates. The yields are from secondary market trades for on-the-run securities for maturities $m = 2, 5, 10, 30$ years.

[a] Indicates that the moments are computed from 2003 onward, the period for which these data are available.

[b] Indicates that the moments are computed for 2000 onward, the period for which these data are available.

III. Quantifying Demand Shocks

This section describes our procedure to measure surprise movements in Treasury yields around Treasury auctions and documents properties of these surprises. Our key assumption is that in a small window around the release of Treasury auction results, shifts in Treasury yields reflect unexpected changes in market beliefs about the demand for Treasuries with a specific maturity. Indeed, the Treasury announces an offered amount well before an auction happens, thus fixing supply in advance of investor bidding. Hence, between the announcement and close of the auction, Treasury yields should move only in response to unexpected changes in demand conditions. Our high-frequency approach isolates variation only due to unexpected shifts in demand arising from a specific auction.

A. Shock Construction

Let $y^{(m)}_{t,\text{pre}}$, $y^{(m)}_{t,\text{post}}$ be the Treasury yields before and after the close of the auction on date t with maturity m. We measure the surprise movements in Treasury yields as

$$D^{(m)}_t = y^{(m)}_{t,\text{post}} - y^{(m)}_{t,\text{pre}}. \tag{1}$$

For all auctions, $y^{(m)}_{t,\text{pre}}$ is the last yield observed 10 minutes before the close of the auction, while $y^{(m)}_{t,\text{post}}$ is the first yield observed 10 minutes following the release of the auction results. If the date t auction is a reopening of a previously issued security, we use secondary market yields to construct our shocks. If the date t auction is instead a newly issued security, we use yields from the when-issued market. In our sample, auctions typically close at 1:00 p.m. or less frequently at 11:30 a.m. However, the time between the close of the auction and the release of the results is a function of how long it takes the Treasury to compile the results. The Treasury began releasing results within minutes in the early 2000s but in the 1990s frequently took longer. We collect wire reports from Bloomberg, which gives a tight upper bound on the release time.[7]

Summary statistics of our constructed shock measures $D^{(m)}_t$ are presented in panel B of table 1 (time series are plotted in fig. B3). The shock means are close to zero (and statistical tests do not reject the null of zero means). Moreover, there is essentially no serial correlation in $D^{(m)}_t$

[7] Lou, Yan, and Zhang (2013) and Fleming and Liu (2016) show that there are predictable price movements in the days and hours before and after the auction, but price movements are unpredictable very near the close of the auction. Hence, the use of small intraday windows is key to identifying unanticipated demand shocks.

($\rho = -0.03$). Hence, these summary statistics indicate that surprises are not systematic and do not contain predictable movements. In our sample period, the standard deviation of $D_t^{(m)}$ increases in maturity m and ranges from 1.4 basis points for 2-year maturity to 3.3 basis points for 30-year maturity. For comparison, Chodorow-Reich (2014) reports that the largest intraday movements in yields following a QE announcement occurred on March 18, 2009, when Treasury (5-year) yields fell by 23 basis points. Panel C presents statistics separately for periods of a binding/nonbinding zero lower bound and for expansion/recessions. Regardless of regime, on average our demand shocks are close to zero.

To verify that these shocks are not spurious, we also report movements in Treasury yields on nonauction days (panel D of table 1; for days without auctions, the same pre and post windows are used as auctions in the same period). In all cases, the variance of the shocks on auction dates is larger than on nonauction dates. This pattern further suggests that surprise auction results influence secondary market Treasury yields.

B. Narrative Evidence

To provide a better understanding of what forces are behind these surprise movements, figure 2 plots 30-year Treasury yields during two 30-year Treasury bond auctions. The dashed lines denote the close of the auction and the release of results, respectively. The first auction occurred on December 9, 2010. This auction was a reopening of previously issued 30-year bonds from the month prior. Yields are relatively stable in the lead up to the close of the auction but drop sharply and immediately following the release of the results. The Financial Times wrote:

> Large domestic financial institutions and foreign central banks were big buyers at an auction of 30-year US Treasury bonds on Thursday. “Investors weren’t messing around . . . You don’t get the opportunity to buy large amounts of paper outside the auctions and ‘real money’ were aggressive buyers.” (Mackenzie 2010)

The second is from an auction of newly issued bonds on August 11, 2011. When-issued yields were relatively stable prior to auction close, but after the close and release of the auction results, yields immediately rose. The Financial Times wrote:

> An auction of 30-year US Treasury bonds saw weak demand . . . bidders such as pension funds, insurers and foreign governments shied away. “There’s not too many ways you can slice this one, it was a very poorly bid auction.” (Demos 2011)

FIG. 2.—Intraday treasury yield movements. The figure shows intraday movements in 30-year Treasury yields around the announcement of auction results for two selected auction dates. Dashed lines denote the close of an auction and release of results. In both auctions considered, auctions closed at 1:00 p.m., and results were released a few minutes thereafter.

We interpret these two examples as follows. Before the auctions closed, the market information set consists of all the supply information for both outstanding securities and the amount on offer for the current 30-year auction. The 30-year Treasury yields reflect beliefs about the expected path of short-term interest rates, inflation expectations, and demand for long-maturity Treasury securities. After the auctions closed and the results were released, the only update to the information set is the news regarding the bidding that took place in the auction, which solely reflects demand for Treasury debt. The change in the 30-year yields is a reaction to the unexpected shift in Treasury demand revealed at the auction (of course, this shift reflects many factors, including individual investors' beliefs about interest rates and economic fundamentals).

These articles also highlight why auctions can have important elements of price discovery: when investors wish to purchase large amounts of Treasuries to meet their needs, they may prefer to use auctions rather than make large transactions on the secondary market. As a result, auctions reveal new information about demand for government debt that is not already reflected in OTC secondary market trades.

C. Demand Determinants

The institutional structure of the primary market and narrative evidence from the financial press suggest that $D_t^{(m)}$ captures unexpected shifts in the demand for Treasuries. But because $D_t^{(m)}$ is an equilibrium response, it is important to establish that these shifts are related to observable measures of demand. Table 2 presents formal evidence by regressing our shocks on measures of demand reported at the auction:

$$D_t^{(m)} = \alpha^{(m)} + \beta^{(m)} X_t^{(m)} + \varepsilon_t^{(m)}, \tag{2}$$

where $X_t^{(m)}$ are various estimates of the shift in demand at a given auction. Columns 1–6 present estimates separately for auctions of different maturities, while column 7 pools across maturities. Panel A uses changes in the bid-to-cover ratio (the change is taken relative to the most recent note or bond auction), which is a natural measure of demand shifts, since an increase in the bid-to-cover ratio indicates higher demand relative to the amount of Treasuries offered. Consistent with our interpretation of $D_t^{(m)}$ as a reaction to unexpected demand, panel A shows that an increase in the bid-to-cover ratio predicts a larger intraday fall in Treasury yields following the close of the auction.[8]

Our results show that the effect of typical surprise increases in demand is economically large. For example, a 1 standard deviation (0.45) increase in the bid-to-cover ratio in a Treasury auction for 10-year notes leads to a $2.66 \times 0.45 \approx 1.2$ basis point decline in 10-year Treasury yields. We can back out a simple estimate for the sensitivity of yields as a function of the change in quantity demanded (in terms of dollars). A typical offering amount in a 10-year Treasury note auction is between \$20 and \$30 billion. Hence, our estimates imply that an increase in demand for 10-year Treasuries by \$10 billion decreases 10-year yields by $2.66 \times (10/30, 10/20) \approx (0.89, 1.33)$ basis points.

In order to assess sensitivity of our demand shocks to changes in demand by bidder type, panel B reports estimates of equation (2) using the change in the bid-to-cover ratio of indirect bidders, direct bidders, and primary dealers. The sensitivity of surprises $D_t^{(m)}$ to unexpected demand of indirect bidders increases with maturity. For example, a unit increase in the bid-to-cover ratio for indirect bidders decreases the yields of 2-year Treasuries by 2.24 basis points and the yields of 30-year Treasuries by 11.3 basis points. Direct bidders exhibit the same pattern, although

[8] Figure B4 reports binscatter plots of these results. Table 2 uses changes, as there is predictable low-frequency movement in the bid-to-cover ratio (figs. B7, B8). However, our results are not sensitive to this choice. Figures B5 and B6 repeat our analysis using the bid-to-cover ratio in levels or a residualized measure of the bid-to-cover ratio from a univariate AR(4) model.

TABLE 2
Demand Shocks and Measures of Demand

	$D_t^{(2Y)}$ (1)	$D_t^{(3Y)}$ (2)	$D_t^{(5Y)}$ (3)	$D_t^{(7Y)}$ (4)	$D_t^{(10Y)}$ (5)	$D_t^{(30Y)}$ (6)	Pool D_t (7)
	A. Total Bid-to-Cover Ratio						
Bid-to-cover ratio	−.74***	−1.51***	−1.25***	−5.67***	−2.66***	−4.22***	−1.64***
	(.18)	(.27)	(.23)	(1.22)	(.45)	(1.23)	(.17)
Observations	255	134	233	106	186	132	1,046
R^2	.08	.23	.10	.28	.17	.21	.12
	B. Bid-to-Cover Ratio by Bidder Type						
Indirect	−2.24***	−3.83***	−4.22***	−8.34***	−5.59***	−11.30***	−5.10***
	(.42)	(.92)	(.75)	(1.89)	(.97)	(2.30)	(.57)
Direct	−.41	−.08	−1.14	−7.40***	−1.03	−3.29	−1.63**
	(.74)	(1.00)	(1.19)	(2.06)	(1.03)	(3.24)	(.73)
Primary	−.56**	−.89**	−.76**	−1.28	−1.74***	−1.32	−.86***
	(.23)	(.41)	(.35)	(1.82)	(.54)	(1.38)	(.24)
Observations	163	120	172	106	150	115	826
R^2	.20	.25	.17	.34	.28	.37	.23
	C. Fraction Accepted by Bidder Type						
Investment fund	−3.03***	−3.34***	−3.93***	−11.76***	−3.67**	−12.96***	−5.37***
	(.90)	(1.13)	(1.20)	(2.32)	(1.53)	(3.09)	(.73)
Foreign	−2.13**	−5.80***	−4.21***	−15.23***	−5.35***	−16.07***	−5.80***
	(.98)	(2.02)	(1.43)	(4.37)	(1.42)	(3.92)	(.79)
Miscellaneous	−1.81	−3.62	−3.14	−13.91***	−4.37	−4.97	−3.96**
	(2.29)	(3.74)	(2.20)	(3.50)	(3.72)	(10.13)	(1.61)
Observations	200	121	185	106	164	119	895
R^2	.08	.14	.07	.30	.11	.28	.13

Note.—The table shows regressions of demand shocks $D_t^{(m)}$ on the change in bid-to-cover ratio and fractions accepted, total and broken up by bidder type. Columns 1–6 restrict the sample to include only auctions of maturities $m = 2, 3, 4, 5, 7, 19, 30$ years; col. 7 pools across all auctions. Changes in the bid-to-cover ratio and fraction accepted are computed relative to the most recent note or bond auction. Newey-West (nine lags) standard errors are in parentheses.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

the coefficients are smaller. The sensitivity to changes in the bid-to-cover ratio coming from primary dealers is also smaller. When we pool across maturities, demand of direct and especially indirect bidders generates ceteris paribus more variation in Treasury yields than demand of primary dealers, although for all bidder types an increase in bidder demand implies a decline in intraday yields $D_t^{(m)}$.

Panel C uses additional investor allotment data from the Treasury to break down the amount accepted by types of bidders: investment funds, foreign, dealers, and remaining smaller investors classes (aggregated into a miscellaneous category). Since the fractions by group add up to 1, we set dealers as the leave-out category. Our estimates suggest that as the fraction accepted for investment funds and foreign buyers increases, $D_t^{(m)}$ declines. Estimates for the miscellaneous category are generally smaller and less robust.

Our results indicate that movements in demand, proxied by the bid-to-cover ratio, are a key determinant of $D_t^{(m)}$. Furthermore, we observe that the demand from institutional investors is important in accounting for variation in $D_t^{(m)}$. We stress that our identifying assumption relies on observing only the changes in yields immediately following the close and release of auction results. Our empirical approach does not require a measure of the unanticipated movement in quantity demanded. Hence, our empirical analysis in the remainder of the paper relies on only our constructed measure $D_t^{(m)}$.

D. *Comovement across Markets*

We now turn to analyzing how our demand shocks for Treasuries propagate across other financial markets. We measure the impact of demand shocks on other assets by estimating univariate regressions of the form

$$y_t = \gamma + \phi D_t + u_t, \tag{3}$$

where y_t is the change in the price or yield of some asset on auction date t and D_t is our auction demand shock. We pool across all auction maturities to simplify presentation, but our results are robust to estimating equation (3) separately by maturity groups.

Where available, we use intraday changes measured in the same window as our shocks. We also examine daily changes, both because of data limitations for some series and because daily changes may pick up responses that do not occur immediately. A strong relationship between D_t and y_t indicates either that D_t and y_t have a common determinant (e.g., changes in inflation expectations alter the behavior of bids in

Treasury auctions and change prices of inflation swaps) or that there is a propagation channel from D_t to y_t (e.g., higher revealed demand for Treasuries results in repricing of other debt securities).

Panel A of table 3 reports results for debt markets. The dependent variable in the first row is the intraday change in the price (comoves negatively with the yield) of the Exchange Traded Fund (ETF) LQD, the iShares iBoxx $ Investment Grade Corporate Bond tracking investment grade corporate bonds. The estimated coefficient $\hat{\phi}$ is interpreted as the impact in log points of a 1 basis point increase in D_t. We observe a strong reaction to the Treasury demand shock, accounting for more than 50% of variation observed in corporate bond ETF prices during the short windows around the close and release of the Treasury auction results. The second row reports the results using HYG, the iShares iBoxx $ High Yield Corporate Bond ETF tracking high-yield corporate bonds. Although the sign of the estimated coefficient is as expected, the pass-through from our demand shocks to high-yield corporate bonds is weaker than that of investment grade bonds. The estimated coefficient is not statistically significant, and the magnitude is smaller than for investment grade debt by a factor of 10. Moreover, the R^2 is much lower, suggesting that our demand shocks account for only 1% of the observed variation in high-yield corporate bond ETF prices during these intraday periods. The next two rows examine pass-through to mortgage rates, as measured by the ETFs MBB and VMBS (from iShares and Vanguard, respectively; both ETFs track investment-grade mortgage-backed pass-through securities guaranteed by US government agencies). Similar to our findings for investment-grade corporate debt, we find economically large and statistically significant pass-through of demand shocks to mortgage borrowing rates. We also find relatively high R^2s of 34% and 13%, respectively.

The next rows repeat the above analysis using daily measures of corporate bond yields, as measured by Moody's Aaa, Moody's Baa, and Bank of America's C corporate yield indexes. Consistent with the intraday results, our demand shocks have a strong effect on safe (Aaa) corporate bonds. Moreover, the pass-through of our demand shocks to corporate bond yields is nearly one-to-one. However, using daily rather than intraday changes as the dependent variable leads to a decline in R^2, which underscores the benefits of using intraday data. While our demand shocks still have large effects on moderately safe debt (Baa), there appears to be much smaller transmission to highly risky corporate debt (C), consistent with our intraday findings.

Secondary market yields react not only strongly but also persistently. Figure 3 plots the contemporaneous reaction of 10-year Treasury spot rates (*A*) and the Aaa corporate bond yields (*B*) to our shocks D_t as well as the reactions up to 60 days in the future. The reaction remains strongly statistically significant over 3 weeks later, while the point estimate is quite

TABLE 3
Asset Price Reactions to Demand Shocks

Asset Type	Estimate (1)	Standard Error (2)	Observations (3)	R^2 (4)	Sample (5)
	A. Corporate and Private Debt				
LQD	−3.93***	(.15)	830	.59	2002–17
HYG	−.34	(.30)	678	.01	2007–17
MBB	−1.42***	(.18)	662	.34	2007–17
VMBS	−1.46***	(.33)	371	.13	2009–17
Corporate Aaa[a]	.94***	(.10)	1,040	.14	1995–2017
Corporate Baa[a]	.96***	(.10)	1,040	.15	1995–2017
Corporate C[a]	.23	(.39)	973	.00	1997–2017
	B. Equities				
SPY	−.23	(.52)	1,033	.00	1995–2017
IWM	.18	(.58)	876	.00	2000–2017
S&P 500[a]	3.61	(2.74)	974	.00	1995–2016
Russell 2000[a]	6.26*	(3.25)	974	.01	1995–2016
	C. Swaps, Commodities, and Spreads				
GLD	−1.16***	(.36)	775	.02	2004–17
GSCI[a]	−1.20	(2.74)	974	.00	1995–2016
10-year inflation swap[a]	.02	(.08)	724	.00	2004–16
2-year inflation swap[a]	.04	(.17)	724	.00	2004–16
LIBOR-OIS[a]	.03	(.04)	737	.00	2003–16
Auto CDS[a]	−.60	(2.65)	733	.00	2004–16
Bank CDS[a]	−.23	(.16)	733	.01	2004–16
VIX[a]	−3.82	(3.37)	1,040	.00	1995–2017
	D. Federal Funds Futures				
1-month ahead[a]	.03	(.02)	1,040	.00	1995–2017
2-month/1-month slope[a]	.00	(.01)	1,040	.00	1995–2017

Note.—The table shows regressions of asset price changes on demand shocks D_t (pooled across maturities). Intraday changes are from ETFs that track securities or indexes: LQD (investment grade corporate debt); HYG (high-yield corporate debt); MBB and VMBS (mortgage indexes); SPY (S&P 500); IWM (Russell 2000); GLD (gold bullion). Daily series: Aaa, Baa, C (Moody's and Bank of America corporate debt yield indexes); S&P 500, Russell 2000 (equity indexes); GSCI (S&P Total Commodity Index); 10- and 2-year inflation swaps; auto and bank credit default swap indexes; LIBOR-OIS (3-month USD LIBOR–overnight index swap spread); VIX (implied volatility index); Federal funds futures (h-month ahead continuous contracts). Newey-West (nine lags) standard errors are in parentheses. All regressions are estimated using ordinary least squares. We report very similar instrumental variable estimates (using changes in the bid-to-cover ratio as instruments) in table B1, available online.

[a] Denotes daily series; otherwise, intraday changes are measured over the same window as auction demand shocks.

* Statistically significant at the 10% level.

*** Statistically significant at the 1% level.

stable even 2 months later. To provide a perspective on the magnitude of this persistence, in figure 3 we also plot the change in yields following the QE1 announcement on March 18, 2009 (normalized such that the change on impact is equal to 1). Consistent with Wright (2012) and

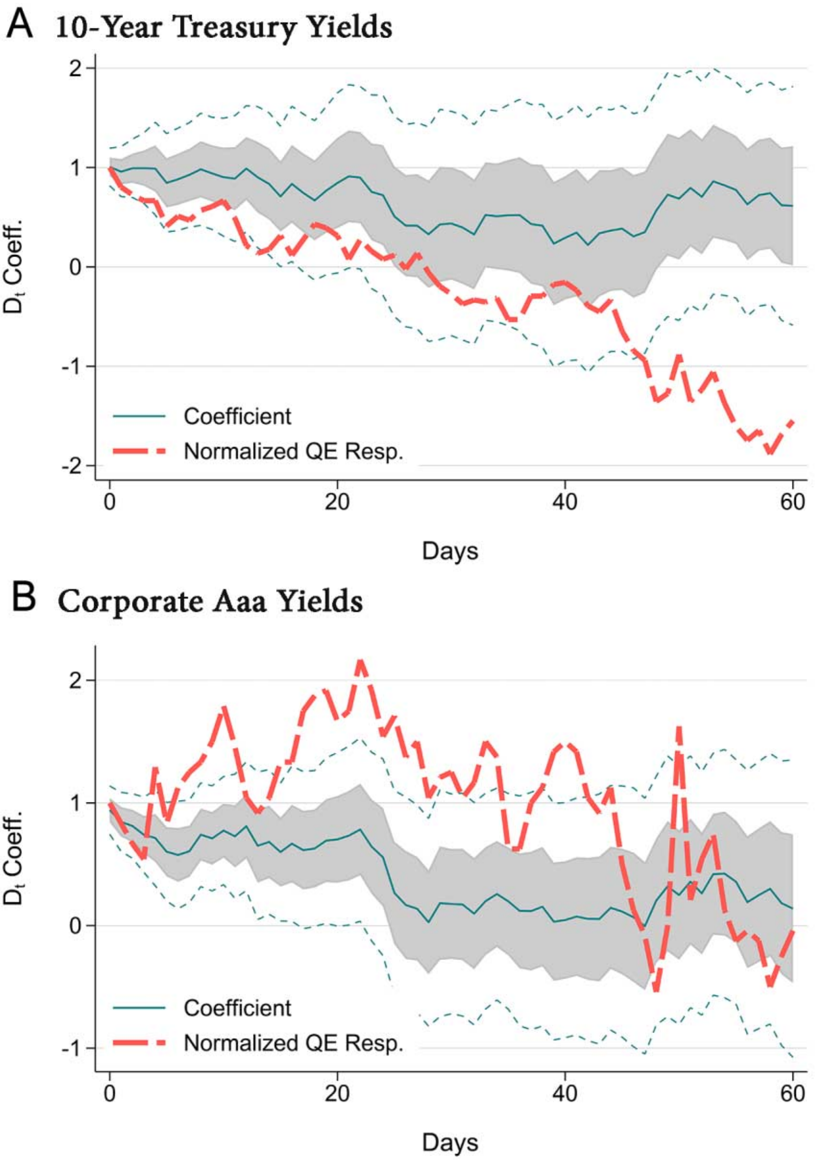

FIG. 3.—Long-difference response to demand shocks. The figure shows responses of 10-year Treasury spot rates (*A*) and Moody's Aaa yields (*B*) to demand shocks D_t (pooled across maturities). We compute long-difference regressions: on an auction date t, the dependent variable is $y_{t+h} - y_{t-1}$, the change h days after the auction relative to the day before the auction. The solid line plots the coefficients from regressions for $h = 0, \ldots, 60$; the shaded region and thin dashed lines correspond to 1 and 2 standard error (Newey-West, nine lags) confidence bands. The thick dashed line compares the long-horizon effects of the March 2009 QE1 FOMC announcement (normalized so that the impact on announcement is 1).

Greenlaw et al. (2018), yields had returned to roughly their starting point within a few months after QE1.

One may be concerned that these reactions in the bond market are driven by some omitted factor rather than idiosyncratic changes in institutional investors' demand for Treasuries. Our high-frequency identification goes a long way toward assuaging this concern, but it is still possible that our demand shocks reflect systematic changes in bidders' expectations of macroeconomic or financial fundamentals. This is an issue for our interpretation if two conditions hold: (1) this information was not already reflected in market prices and instead is revealed only at the auction; (2) the market as a whole updates their beliefs about fundamentals by observing the results of the auction. If this information factor drives the bond market reactions we observe, we should find strong comovement between D_t and indicators capturing beliefs about current or future states of the economy and financial markets. The information factor can take any number of forms, so it is difficult to conclusively rule out this channel; we focus on a battery of key indicators to assess the quantitative significance of this alternative explanation.

Panel B of table 3 reports estimates for the response of equities to our demand shocks. Rows 1 and 2 report the results for the intraday change in ETFs tracking the Standard and Poor's (S&P) 500 and the Russell 2000 indexes. Rows 3 and 4 are for the daily changes in these indexes. Although the estimated slope is generally positive, the estimate is typically insignificant. Moreover, the quantitative importance is small, as the auction demand shocks account for a trivial share of variation in equities. Thus, there is little evidence on average for a common factor driving Treasury and equity prices during small windows around the release of auction results or for strong propagation of our demand shocks to equity prices.

Panel C of table 3 presents results for a wider array of assets. The dependent variable in row 1 is the intraday change in the ETF GLD, which tracks the price of gold bullion. Row 2 reports results for the daily change in the S&P Total Commodity Index. For the Commodity Index, we do not find a significant correlation with D_t. For gold, while the relationship is statistically significant, the R^2 is very low. Rows 3 and 4 use the daily change in inflation expectations implied by inflation swaps at the 10- and 2-year horizons. Our demand shocks are not associated with a significant comovement in inflation expectations. Row 5 uses the daily change in the 3-month USD London Interbank Offered Rank (LIBOR)–overnight index swap spread (a common measure of credit risk in the banking sector). Rows 6 and 7 use daily changes in two credit default swap (CDS) indexes from Credit Market Analysis that track the automotive industry (a highly cyclical industry) and banks (a proxy for the financial sector). These three measures proxy for expectations about future output and market conditions. Our demand shocks have no tangible relationship with these measures, consistent with our interpretation that the shocks do not capture superior information of Treasury auction bidders about future recessions and the like. Finally, row 8 documents that D_t shocks are not associated with the Volatility Index (VIX; a measure of market perceptions of future volatility). In short, these null results suggest that our demand shocks are not driven by changes in expectations regarding inflation, output, liquidity, default risk, or volatility.

As a final test, panel D of table 3 reports results with expected federal funds rates, which are derived from federal funds futures contracts. We view this as a catchall test of the information factor story: if our demand shocks reflect changes in market expectations of fundamentals, then the market should also expect a Fed response (to the extent that these fundamentals matter for the macroeconomy). We find that our demand shocks are not associated with changes in expected federal funds rates: the estimated coefficients on our demand shocks are very close to zero in these specifications, and demand shocks explain virtually none of the variation in federal funds futures in our sample.[9]

E. Localization Hypothesis

The results of tables 2 and 3 allow for some broad observations. First, given our high-frequency approach and the institutional structure of Treasury auctions, our constructed shocks are likely driven by only new information regarding the demand side of the market. Second, these shifts are largely driven by shifts in the demand arising from institutional investors. Third, these demand shocks from the primary market for Treasuries propagate to private borrowing rates. We do not observe the underlying sources of demand shocks, and indeed it is likely driven by a host of factors, including shifts in a given investor's idiosyncratic beliefs about the macroeconomy. However, these demand shifts are not driven by systematic or market-wide shifts in macroeconomic expectations (e.g., flight to quality or inflation expectations) that may move demand for Treasuries at all maturities.

We now turn to testing the key predictions of preferred habitat theory, first discussed in Modigliani and Sutch (1966). Preferred habitat theory posits that certain investor clienteles specialize in bonds of specific maturities; idiosyncratic positive (negative) shocks to their demand leads to decreases (increases) in bond yields. This prediction contrasts with the neutrality result of standard theories and the expectations hypothesis: holding fixed the expected path of short rates, idiosyncratic demand

[9] We plot rolling regressions for select asset prices in fig. B9. Figure B10 plots rolling regressions for h-month ahead federal funds futures h up to 12 months.

shocks for specific bonds should have no effect on yields. Our results in table 2 are consistent with the basic preferred habitat theory (and are inconsistent with the neutrality result of the expectations hypothesis): increases in idiosyncratic demand (proxied by the bid-to-cover ratio at a given auction) imply decreases in yields following the close of a given auction.

However, a naive preferred habitat view where yields for a given bond are determined solely by idiosyncratic demand shocks is unrealistic, as this would imply large arbitrage profits to be made by term structure arbitrageurs. Thus, modern preferred habitat theory (Vayanos and Vila 2021) formalizes the interaction between such arbitrageurs and preferred habitat investors. One of the characteristic predictions of modern preferred habitat theory is the localization hypothesis. When arbitrageur risk-bearing capacity is high, habitat demand shocks have global effects on the yield curve. That is, the relative response of the interest rates across the yield curve does not depend on where in maturity space the demand shock occurs. However, as arbitrageur risk-bearing capacity declines, the spillovers of demand shocks become more localized. That is, the relative response of interest rates becomes more concentrated on parts of the yield curve that are closer in maturity space to where the demand shock occurs.

Intuitively, as risk-bearing capacity falls, term structure arbitrageurs find it more and more costly to integrate the markets for bonds of different maturities. Thus, demand shocks in one maturity segment of the yield curve will have relatively larger effects for bonds with similar maturities to those directly affected by the shift in demand. Yields for maturities that are far removed from those of the bonds directly affected by the demand shock will instead exhibit relatively small movements.

Thus, testing the localization hypothesis requires empirical measures of the conditional response of yields to demand shifts for short- and long-maturity Treasuries, holding all other risk factors constant. Our high-frequency identification strategy is designed to isolate the reactions of the yield curve to Treasury demand shocks alone, effectively ruling out any other shocks during these small windows around auctions. This insight allows us to arrive at a simple regression specification that can test the localization hypothesis using only high-frequency changes in yields around auctions $D_t^{(\tau)} = y_t^{(\tau),\text{post}} - y_t^{(\tau),\text{pre}}$:

$$D_t^{(\tau)} = \alpha^{(\tau)} + \gamma_s^{(\tau)} \mathbf{I}(m_t = short) D_t^{(\tau^*)} + \gamma_\ell^{(\tau)} \mathbf{I}(m_t = long) D_t^{(\tau^*)} + \varepsilon_t^{(\tau)}, \quad (4)$$

where $\mathbf{I}(m_t = short)$ and $\mathbf{I}(m_t = long)$ are indicator variables equal to 1 if there is an auction of short-maturity and long-maturity bonds on date t, respectively. We estimate equation (4) for all maturities τ, holding some baseline maturity τ^* fixed. Coefficients $\hat{\gamma}_s^{(\tau)}$ and $\hat{\gamma}_\ell^{(\tau)}$ measure the relative effect of demand shocks for short-maturity and long-maturity auctions respectively. Now we can restate the localization hypothesis formally as a function of the coefficients $\gamma_s^{(\tau)}$, $\gamma_\ell^{(\tau)}$ as follows: (1) when risk-bearing capacity is high, $|\gamma_s^{(\tau)} - \gamma_\ell^{(\tau)}| \to 0$ for all maturities τ; (2) when risk-bearing capacity is low, then $\gamma_s^{(\tau)} > \gamma_\ell^{(\tau)}$ if $\tau < \tau^*$ and $\gamma_s^{(\tau)} < \gamma_\ell^{(\tau)}$ if $\tau > \tau^*$.

F. *Empirical Localization Results*

To test the state-dependent localization hypothesis, we estimate specification (4) separately for two subsamples: a noncrisis period (when risk-bearing capacity is high) and a crisis period (when risk-bearing capacity is low). Testing for the equality of coefficients $\hat{\gamma}_s^{(\tau)}$ and $\hat{\gamma}_\ell^{(\tau)}$ in each subsample shows whether demand shocks have more localized effects when risk-bearing capacity is low. In our baseline estimates, we take the crisis period to be 2008–12. We divide the auctions into short maturity (maturity of 5 years or less) and long maturity (maturity of 7 years or greater). The results are robust to alternative choices.

Since we construct $D_t^{(\tau)}$ from secondary market yields, we do not have measures for all maturities τ at all times. Our approach is to run rolling regressions (across maturities), including yields for maturities within ±2 years for each maturity $\tau \leq 20$ years and within ±4 years for each maturity $\tau > 20$ year. Given that the choice of benchmark maturity τ^* is arbitrary, we set $\tau^* = 3$, with an eye toward applying our results to QE in order to focus on the differential effects of intermediate and long-maturity yields.

Figure 4*A* plots the estimates of $\hat{\gamma}_s^{(\tau)}$ and $\hat{\gamma}_\ell^{(\tau)}$ for the noncrisis sample. The estimated responses follow a similar hump-shaped pattern, peaking at intermediate maturities (around roughly 5–7 years) and then declining and stabilizing for longer-term maturities (with a possible slight uptick for very long [20+ years] maturities). In general, we cannot reject equality of $\hat{\gamma}_s^{(\tau)}$ and $\hat{\gamma}_\ell^{(\tau)}$.

When we estimate specification (4) on the crisis sample, we observe a strikingly different pattern (fig. 4*B*). In response to shocks in demand for short- and long-maturity Treasuries, both $\hat{\gamma}_s^{(\tau)}$ and $\hat{\gamma}_\ell^{(\tau)}$ increase rapidly for $\tau < \tau^*$, with the $\hat{\gamma}_s^{(\tau)}$ estimates increasing somewhat more quickly than $\hat{\gamma}_\ell^{(\tau)}$. Once we move to maturities greater than τ^*, the estimates quickly diverge. Specifically, in response to shocks in demand for short-maturity Treasuries, the yields for maturities greater than τ^* fall relative to the yield response for τ^* (as shown by the estimates of $\hat{\gamma}_s^{(\tau)}$). On the other hand, in response shocks in demand for long-maturity Treasuries, yields for maturities $\tau > \tau^*$ continue to increase relative to the benchmark maturity τ^* (as shown by the estimates of $\hat{\gamma}_\ell^{(\tau)}$). For maturities of 20 to 30 years, the relative response of yields to long-maturity demand shocks ($\hat{\gamma}_\ell^{(\tau)}$) is over

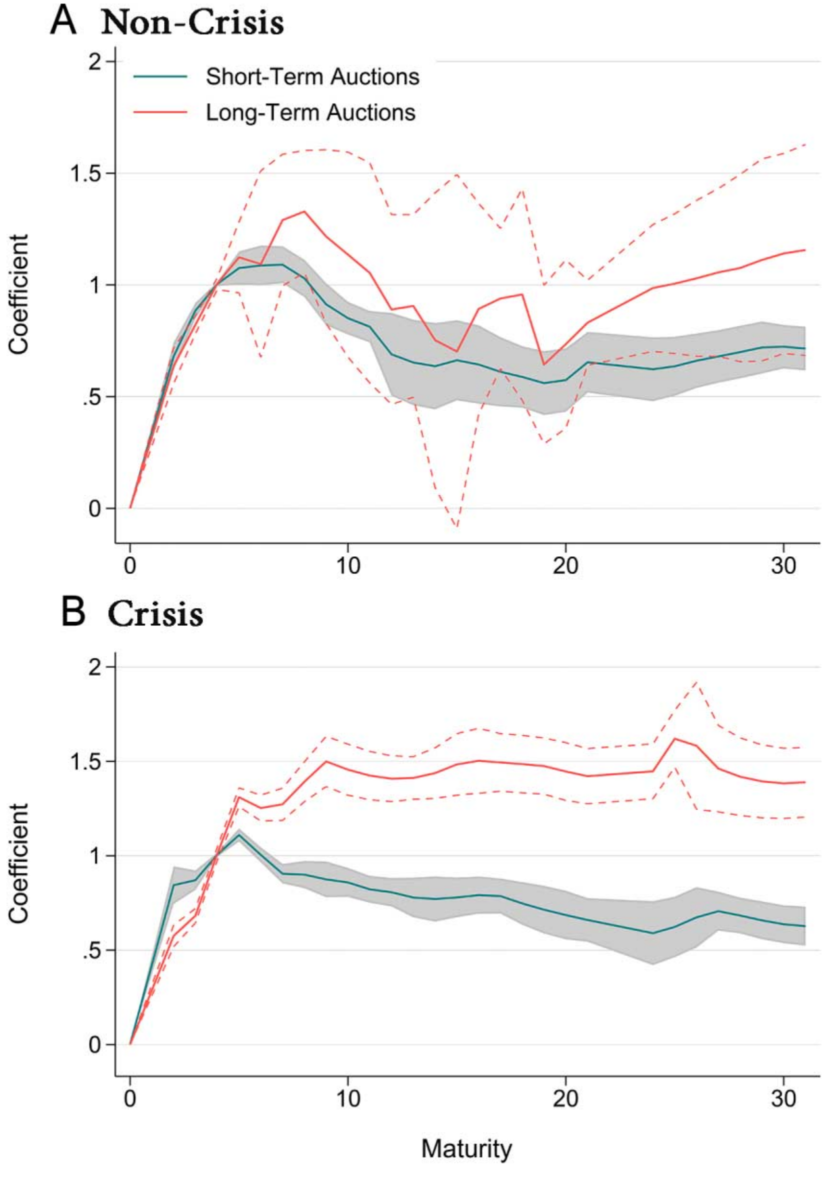

FIG. 4.—Localization regression results. The figure shows plots of the regression coefficients from regression equation (4). *A* compares estimates from short- and long-maturity auctions during noncrisis periods, and *B* compares these estimates during crisis periods, with 2 standard error (Newey-West, nine lags) confidence intervals included. Figure B11 reports *p*-values for the equality of estimates.

twice as large as the response of yields to short-maturity demand shocks ($\hat{\gamma}_s^{(\tau)}$). We can strongly (at the 0.1% level) reject the null of $\hat{\gamma}_s^{(\tau)}$ and $\hat{\gamma}_\ell^{(\tau)}$ being equal.

These results are consistent with the key predictions of preferred habitat theory: during normal periods when financial risk-bearing capacity is high, demand shocks for short- and long-maturity securities have relatively similar impacts on the yield curve. During periods of financial distress when risk-bearing capacity is low, the impacts are more localized: the impact of short-maturity demand shocks are largest for short maturities, while the impact of long-maturity demand shocks peaks at the long end of the term structure. These results provide support for the view that during financial crises, arbitrageurs are less willing or able to integrate bond markets.

The localization results are robust to a variety of alternative assumptions. In the interest of space, we provide summaries for a few key robustness specifications in this section. First, we explore the sensitivity of our results to using alternative measures of financial distress. Specifically, we consider two alternatives: (1) an aggregate intermediary capital ratio, a market-based measure of financial distress (low intermediary capital ratios are associated with lower risk-bearing capacity) described in He, Kelly, and Manela (2016); and (2) a narrative-based measure of financial crisis (higher values of the crisis indicator are associated with lower risk-bearing capacity) from Romer and Romer (2017). Both of these alternatives generate results similar to our baseline findings (see figs. B12–B16). Second, our results are robust to using different cutoffs for separating auctions into short and long maturities. For example, when we use 10 years (rather than 7 years in the baseline) as the cutoff for long-maturity auctions, the results (fig. B17) still strongly support the localization hypothesis and, if anything, the point estimates for the noncrisis period are even more similar than in our baseline specification. Third, as we discussed above, the choice of the benchmark maturity τ^* is arbitrary in our framework. To verify that this choice is indeed immaterial for our results, we experiment with values other than $\tau^* = 3$ and find similar results (e.g., fig. B18 reports estimates for $\tau^* = 6$). Fourth, our results are nearly identical when dropping auctions that occurred during the weeks of QE announcements (fig. B19).

Note that specification (4) does not require us to observe the underlying shock to demand quantity but instead relies on only the change in yields conditional on such (unobserved) changes in demand. However, the change in the bid-to-cover ratio is a proxy for demand shifts. Although we do not have a high-frequency market-based measure of the expected bid-to-cover ratio before and after an auction, we nevertheless considered an alternative regression specification to test the localization hypothesis:

$$D_t^{(\tau)} = \alpha^{(\tau)} + \nu_s^{(\tau)}\mathbf{I}(m_t = short)\tilde{\beta}_t + \nu_\ell^{(\tau)}\mathbf{I}(m_t = long)\tilde{\beta}_t + \varepsilon_t^{(\tau)}, \quad (5)$$

where $\tilde{\beta}_t$ is the change in the bid-to-cover ratio following an auction at date t (as in table 2). Figure 5 reports the results (after flipping the sign of the coefficients for comparability with our baseline specification). Unlike specification (4), the estimates $\hat{\nu}_s^{(\tau)}, \hat{\nu}_\ell^{(\tau)}$ measure the absolute response of the yield curve to shifts in demand (proxies) as opposed to the relative response. Hence, the formal test of the localization hypothesis is not as straightforward. With this caveat, we see clear patterns of localization in the crisis estimates but little evidence of localization in the noncrisis estimates. In the noncrisis subsample estimates (fig. 5*A*), the response to both short- and long-maturity demand follows a similar hump-shaped pattern: the response is increasing from short to intermediate maturities, peaking around $\tau = 5$ to $\tau = 10$ before declining (though there is some evidence that long demand shocks have larger effects on very long-maturity yields). On the other hand, the estimates from the crisis subsample (fig. 5*B*) exhibit significant differences across short and long estimates: the response to short demand shocks remains hump shaped but peaks for short maturities before quickly declining. Conversely, the response to long demand shocks increases almost without fail as the maturity increases, peaking at very long maturities.[10]

We additionally explore state dependence in the pass-through of Treasury demand shocks to private borrowing rates. Although preferred habitat models typically feature only risk-free bonds, the logic of the localization hypothesis also suggests that the pass-through to risky assets of Treasury demand shocks should lessen when risk-bearing capacity is low. Panel A of table 4 estimates a version of equation (4) where we regress the intraday change of various ETF prices on $D_t^{(\tau^*)}$, interacted with whether the auction took place during the crisis period. For very safe or very risky corporate borrowing (LQD and HYG in cols. 1 and 2, respectively), we find little evidence of localization. Although the sign of the crisis interaction coefficient in both cases is positive, it is not statistically significant. On the other hand, mortgage borrowing rates (MBB and VMBS in cols. 3 and 4, respectively) seem to exhibit substantial localization. In both cases, the interaction coefficient is positive, economically large compared with the noncrisis coefficient, and highly statistically significant. Next, we explore the pass-through to equity indexes. For the S&P 500 (SPY; col. 5), we find little evidence for localization; while for the Russell 2000 (IWM; col. 6), there is some support for localization: the coefficient is positive in noncrisis

[10] Figures B20 and B21 repeat the estimation of specification (5) but use the level of the bid-to-cover ratio or the residuals of a univariate AR(4) model. Figures B22–B24 conduct additional robustness exercises using measures of the bid-to-cover ratio only from indirect bidders. In all cases, the patterns observed are very similar to our results in fig. 5.

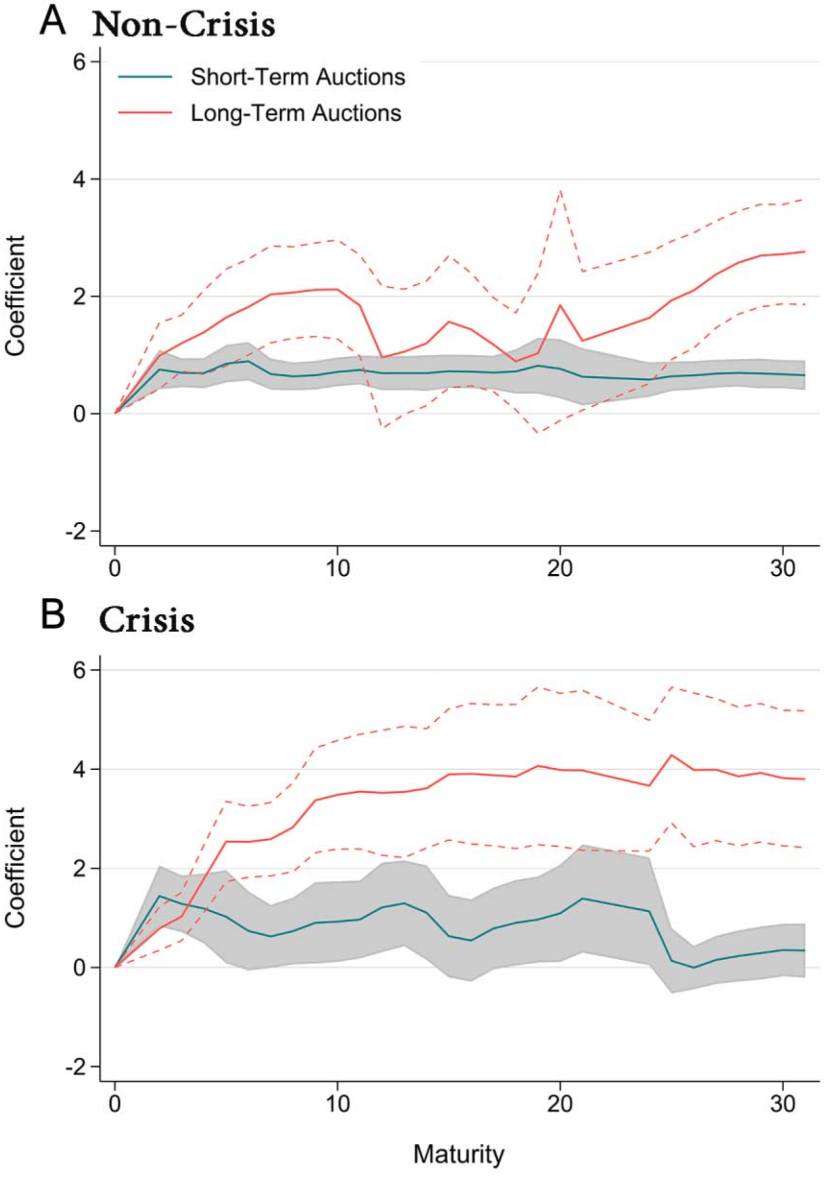

FIG. 5.—Alternative localization: bid-to-cover ratio. The figure shows estimates of the alternative localization regression (5), using the change in the bid-to-cover ratio as a proxy of structural demand shocks. We further flip the sign of the coefficients in order to make the estimates more comparable to our baseline specification. *A* compares estimates from short- and long-maturity auctions during noncrisis periods, and *B* compares these estimates during crisis periods, with 2 standard error (Newey-West, nine lags) confidence intervals included.

TABLE 4
LOCALIZATION REGRESSION RESULTS, PRIVATE BORROWING

	LQD (1)	HYG (2)	MBB (3)	VMBS (4)	SPY (5)	IWM (6)
	A. Crisis Localization					
$D_t^{(\tau^*)}$	−4.98***	−.18	−3.41***	−2.83***	−.03	2.85***
	(.58)	(.56)	(.41)	(.42)	(1.22)	(1.02)
Crisis × $D_t^{(\tau^*)}$	.30	.19	1.80***	4.51***	.97	−2.64*
	(.72)	(.69)	(.43)	(1.58)	(1.56)	(1.42)
Observations	794	661	647	376	941	811
R^2	.40	.00	.35	.15	.01	.01
Sample	2002–17	2007–17	2007–17	2009–17	1995–2017	2001–17
	B. Crisis and Maturity Localization					
$D_t^{(\tau^*)}$	−1.54***	−5.20***	−8.74***	−.84***	−2.98***	−4.32***
	(.18)	(.37)	(.93)	(.17)	(.25)	(1.24)
Long × $D_t^{(\tau^*)}$	−.22	−.88	−4.45*	.23	−.58	−4.37
	(.38)	(1.46)	(2.31)	(.28)	(.62)	(3.37)
Crisis × $D_t^{(\tau^*)}$	.19	.95	1.72	−.15	.79*	−2.23
	(.25)	(.58)	(1.36)	(.30)	(.46)	(2.85)
Long × crisis × $D_t^{(\tau^*)}$	.04	−2.72*	−5.73*	−.11	−.54	−21.22**
	(.51)	(1.61)	(3.15)	(.52)	(.81)	(10.82)
Observations	664	653	569	647	642	401
R^2	.36	.62	.38	.14	.28	.19
Sample	2007–17	2007–17	2007–17	2007–17	2007–17	2010–17

NOTE.—The table shows regressions of intraday changes in ETFs (proxying for various private borrowing rates) on the intraday change in intermediate yields $D_t^{(\tau^*)}$ (as in fig. 4). Panel A includes interactions with whether the auction took place in a noncrisis or crisis period, while panel B additionally includes interactions with the maturity of the auction (using the same definitions as in fig. 4). The ETFs used in panel A are the same as in table 3. In panel B, the Vanguard ETFs BSV, BIV, and BLV track investment-grade bonds with short, intermediate, and long maturities, respectively. The BlackRock iShares ETFs CSJ, CIU, and CLY track investment-grade corporate bonds with short, intermediate, and long maturities, respectively. Newey-West (nine lags) standard errors are in parentheses.
* Statistically significant at the 10% level.
** Statistically significant at the 5% level.
*** Statistically significant at the 1% level.

times, but the interaction coefficient suggests that during crisis times the pass-through drops to zero. However, in both cases, the R^2 is low.

Finally, we explore state- and maturity-dependent localization for safe corporate borrowing rates. We estimate a version of equation (4) but use intraday changes in various corporate bond ETFs as the dependent variable interacted with both the regime (crisis/noncrisis) and maturity (short/long) of the auction. We utilize two sets of ETFs: BSV, BIV, and BLV from Vanguard and CSJ, CIU, and CLY from iShares. These ETFs track investment-grade bonds with short, intermediate, and long maturities, respectively.[11] The results for this exercise (panel B of table 4) are

[11] The Vanguard ETFs were introduced in 2007 but include Treasuries in addition to corporate bonds. The iShares ETFs include only corporate bonds but have shorter time series.

not directly comparable to our baseline Treasury results from figure 4 because the ETFs are measured as (log) price changes and because we do not observe the entire term structure for corporate debt. Nevertheless, the results are consistent with our findings in figure 4. The first row shows that demand shocks in Treasury auctions also have large effects on corporate bond prices across the term structure. The second row shows that this pass-through is similar for demand shocks originating in long-term auctions in noncrisis times (with the one exception of the coefficient for BLV). The third and fourth rows report the interaction coefficients for the crisis periods. The final row shows that the pass-through of demand shocks for long-maturity Treasuries occurring during crisis periods has the same kind of localization pattern as figure 4. The coefficient becomes monotonically more negative as we move from short-maturity to long-maturity bond ETFs. In particular, the coefficient for the long-maturity ETFs is negative and statistically significant.

IV. A New Keynesian Preferred Habitat Model

The results of section III show that our shock series offer useful variation to study empirically the financial effects of Treasury demand shocks. However, our end goal is to make sense of the transmission of QE both to asset prices and to the broader macroeconomy. Thus, in order to make more progress understanding shock transmission along these dimensions, we need to develop a theoretical framework with meaningful interactions between financial markets and the macroeconomy. This will allow us to be precise about the transmission mechanisms of demand shocks as well as help us extrapolate from private shocks observed at auctions to large-scale asset purchases (or sales). Informed by our empirical findings in section III, this section develops such a model.

A. *Overview: Agents and Financial Assets*

The financial block of our model builds on Vayanos and Vila (2021), extended to allow for both riskless and risky assets. Formally, we assume two sets of assets: riskless and risky zero coupon bonds with maturity $\tau \in (0, T)$. A τ-maturity riskless bond has price $P_t^{(\tau)}$ and pays \$1 at maturity with certainty in period $t + \tau$. A risky τ-maturity bond has price $\tilde{P}_t^{(\tau)}$ and instead pays $D_{t+\tau} \equiv e^{d_{t+\tau}}$ at maturity in period $t + \tau$, where D_t is a stochastic process. We will assume that both sets of assets are in zero net supply, but for completeness we denote the supply of assets by $S_t^{(\tau)}$ and $\tilde{S}_t^{(\tau)}$, respectively.

As in Vayanos and Vila (2021), financial assets are traded between two sets of investors: arbitrageurs and preferred habitat funds. Crucially, there exist clientele investors who have idiosyncratic demand (preferred habitat) for specific assets and maturities. For example, pension funds can have a preference for long-maturity assets to better match the maturity structure of pension liabilities. The other side of the market are risk-averse arbitrageurs, such as hedge funds and dealers, who smooth out these demand shocks.

The macro side of our model is largely built on a standard New Keynesian framework. A representative household consumes and supplies labor to firms producing differentiated goods while facing pricing frictions. Equilibrium aggregate dynamics will depend on two familiar equations: an investment-saving (IS) curve that determines the dynamics of the output gap x_t and a Phillips curve that determines the dynamics of inflation π_t. The links between the financial and the macroeconomic blocks of our model are twofold. First, the central bank sets the policy rate i_t in reaction to aggregate fluctuations. In equilibrium, i_t will be the return on a riskless bond as maturity $\tau \to 0$; therefore, aggregate fluctuations will result in changes in asset prices. Second, we depart from the usual New Keynesian setup and assume that households cannot access financial markets directly but instead must borrow through the preferred habitat funds. Through the household Euler equation, output in equilibrium will depend on a weighted average of the returns on both short- and long-term assets rather than only on the policy rate. Therefore, asset price movements will lead to changes in macroeconomic conditions. Our model formalizes this intuition and characterizes the general equilibrium solution. We stress that our aim is to develop a tractable model and discuss the simplifying assumptions necessary for our model.

Households.—The representative household chooses a consumption bundle C_t and labor supply N_t given nominal wealth A_t in order to maximize lifetime discounted utility. The household faces a nominal wage $\mathcal{W}_t$ and aggregate price index P_t. Our point of departure from a textbook New Keynesian model is that households can invest with only the preferred habitat funds and allocate a fixed fraction of their wealth across these funds. The household problem is given by

$$\max E_0 \int_0^{\infty} e^{-\rho t} \Psi_t \left(\frac{C_t^{1-\varsigma}}{1-\varsigma} - \frac{N_t^{1+\phi}}{1+\phi} \right) \mathrm{d}t, \text{ subject to} \tag{6}$$

$$\mathrm{d}A_t = [\mathcal{W}_t N_t - P_t C_t]\mathrm{d}t + A_t \int_0^T \left(\eta(\tau) \frac{\mathrm{d}P_t^{(\tau)}}{P_t^{(\tau)}} + \tilde{\eta}(\tau) \frac{\mathrm{d}\tilde{P}_t^{(\tau)}}{\tilde{P}_t^{(\tau)}} \right) \mathrm{d}\tau + \mathrm{d}F_t, \tag{7}$$

and A_0 given. The parameter ϕ is the Frisch elasticity of labor supply, and the parameter ρ is the household discount factor. We introduce aggregate demand shocks through discount rate shocks Ψ_t. The term $\mathrm{d}F_t$

captures all (flow) nominal transfers to the household. Households therefore borrow or lend their wealth A_t at an effective borrowing rate (rather than the short-term policy rate), given by the integral term in the household budget constraint. The weight functions satisfy $\int_0^T(\eta(\tau) + \tilde{\eta}(\tau))\mathrm{d}\tau = 1$. Thus, household borrowing depends on a weighted average of the returns across the term structure and asset types. Denote the drift of the effective borrowing rate by

$$\hat{r}_t \equiv E_t \int_0^T \left(\eta(\tau) \frac{\mathrm{d}P_t^{(\tau)}}{P_t^{(\tau)}} + \tilde{\eta}(\tau) \frac{\mathrm{d}\tilde{P}_t^{(\tau)}}{\tilde{P}_t^{(\tau)}} \right) \mathrm{d}\tau, \tag{8}$$

which is endogenous but taken as given by households.

We think of the weighting functions $\eta(\tau)$, $\tilde{\eta}(\tau)$ as capturing a type of preferred habitat for households (e.g., due to demographics, such as age) without explicitly introducing heterogeneous households. This also captures imperfect access to financial markets, as households typically do not directly buy and sell financial assets, such as corporate bonds; instead, investments are done through pension funds and mutual funds, which tend to allocate their portfolio using fixed weights and rebalance sluggishly (e.g., Koijen and Yogo 2019, 2022; Bretscher et al. 2022; Peng and Wang 2022).

To be clear, these weighting functions abstract from many important features of household finance, such as borrowing constraints and heterogeneity of households as well as financial intermediation that offers a variety of products ranging from mortgages to credit card debt (for recent surveys, see, e.g., Zinman 2015; Gomes, Haliassos, and Ramadorai 2021). This creates some distance between the model and the data. For example, QE can change access to credit at different maturities by affecting banks' balance sheets, inducing endogenous changes in the weighting functions $\eta(\tau)$ and $\tilde{\eta}(\tau)$ (e.g., Rodnyansky and Darmouni 2017; Di Maggio, Kermani, and Palmer 2020). In a similar spirit, borrowing constraints and household heterogeneity call for more cross-sectional and time series variation in $\eta(\tau)$ and $\tilde{\eta}(\tau)$ (e.g., Cui and Sterk 2021).

Our key assumption that we can take the weighting functions as model primitives is strong. However, this assumption permits a tractable solution and keeps the macroeconomic dynamics as close as possible to a standard New Keynesian model while still allowing fluctuations in long-term borrowing rates to affect household decisions. Enriching the model with elements necessary to capture additional financial complexities is a fruitful avenue for future research.

Firms.—The production block of the model is standard. Intermediate goods are produced by monopolistically competitive firms indexed by

$j \in [0, 1]$. Each firm hires workers and produces using linear technology $Y_{j,t} = N_{j,t}$. Firms face Rotemberg pricing frictions when choosing prices and must pay a cost $\Theta(\pi_{j,t}) \equiv (\theta/2)\pi_{j,t}^2 P_t Y_t$, where $\mathrm{d}P_{j,t}/P_{j,t} = \pi_{j,t}\,\mathrm{d}t$ is the inflation rate of firm j. Firms are owned by households and hence maximize expected real profits discounted using the household real stochastic discount factor $e^{-\rho t}Q_t$. The firm problem is therefore given by

$$\max E_0 \int_0^\infty e^{-\rho t} Q_t \frac{\Pi_{j,t}}{P_t}\,\mathrm{d}t, \text{ where } \Pi_{j,t} = P_{j,t}Y_{j,t} - \mathcal{W}_t N_{j,t} - \Theta(\pi_{j,t}). \quad (9)$$

Intermediate goods are sold to competitive final goods producers with time-varying demand elasticity ε_t. This implies the usual demand and price index for differentiated goods (see, e.g., Gali 2015)): $Y_{j,t} = (P_{j,t}/P_t)^{-\varepsilon_t} Y_t$, where the price index $P_t = \left(\int_0^1 P_{j,t}^{1-\varepsilon_t}\right)^{1/(1-\varepsilon_t)}$.

Arbitrageurs.—Arbitrageurs have mean-variance preferences and allocate their wealth W_t to holdings $X_t^{(\tau)}$, $\tilde{X}_t^{(\tau)}$ across τ riskless and risky bonds, respectively. They solve

$$\max E_t\,\mathrm{d}W_t - \frac{a}{2}\mathrm{Var}_t\,\mathrm{d}W_t, \text{ subject to} \quad (10)$$

$$\mathrm{d}W_t = W_t i_t\,\mathrm{d}t + \int_0^T X_t^{(\tau)}\left(\frac{\mathrm{d}P_t^{(\tau)}}{P_t^{(\tau)}} - i_t\,\mathrm{d}t\right) + \tilde{X}_t^{(\tau)}\left(\frac{\mathrm{d}\tilde{P}_t^{(\tau)}}{\tilde{P}_t^{(\tau)}} - i_t\,\mathrm{d}t\right)\mathrm{d}\tau, \quad (11)$$

where the parameter a governs the risk-return trade-off that arbitrageurs face. This parameter can be taken literally as a risk aversion parameter or, more generally, can be thought of as a proxy for factors that lead to the imperfect risk-bearing capacity of arbitrageurs. Arbitrageurs transfer all profits or losses to the household each period. Taking a as time invariant is important for tractability; however, we might expect a to increase in periods of financial distress (e.g., Kyle and Xiong 2001; He and Krishnamurthy 2013). We analyze how the predictions of the model depend on the risk-bearing capacity of arbitrageurs.

Preferred habitat funds.—On the other side of the market is a continuum of habitat investors, who specialize in (riskless and risky) bonds of specific maturities. We follow Vayanos and Vila (2021) and assume that these investors respond to prices through the following demand curves:

$$Z_t^{(\tau)} = -\alpha(\tau)\log P_t^{(\tau)} - \beta_t(\tau),\ \tilde{Z}_t^{(\tau)} = -\tilde{\alpha}(\tau)\log\tilde{P}_t^{(\tau)} - \tilde{\beta}_t(\tau). \quad (12)$$

The functions $\alpha(\tau)$, $\tilde{\alpha}(\tau)$ are the semielasticities of a τ-habitat investor's demand for riskless and risky assets, respectively. The time-varying demand intercept is given by $\beta_t(\tau)$, $\tilde{\beta}_t(\tau)$, where we assume a factor structure such that

$$\beta_t(\tau) = \sum_{k=1}^{K}\theta^k(\tau)\beta_t^k - \zeta(\tau), \tilde{\beta}_t(\tau) = \sum_{k=1}^{\tilde{K}}\tilde{\theta}^k(\tau)\tilde{\beta}_t^k - \tilde{\zeta}(\tau). \quad (13)$$

The functions $\theta^k(\tau)$, $\tilde{\theta}^k(\tau)$ govern how demand factors β_t^k, $\tilde{\beta}_t^k$ lead to changes in demand from τ-habitat investors. In general, there are no restrictions on the dynamics of the demand factors (up to linearizing). For our quantitative analysis, we assume that the habitat demand factors are independent from one another but may respond to the short rate (as in King 2019a). Hence, for each demand factor, we have

$$\mathrm{d}\beta_t^k = -(\kappa_\beta^k\beta_t^k + \phi_{i,\beta}^k i_t)\,\mathrm{d}t + \sigma_\beta^k\,\mathrm{d}B_{\beta^k,t}, \quad (14)$$

where the parameters κ_β^k, $\phi_{i,\beta}^k$, and σ_β^k govern the mean reversion, response to short rates, and volatility of each demand factor, respectively (and with analogous parameterization and dynamics for demand factors $\tilde{\beta}_t^k$).

The τ-maturity riskless and risky funds have wealth $W_t^{(\tau)}$ and $\tilde{W}_t^{(\tau)}$. In order to fund their positions in their respective asset classes, the fund receives investment from the household sector on the basis of the weights $\eta(\tau)$, $\tilde{\eta}(\tau)$. Any remaining wealth is invested at the short rate. The budget constraint of a τ-maturity riskless fund is therefore

$$\mathrm{d}W_t^{(\tau)} = Z_t^{(\tau)}\frac{\mathrm{d}P_t^{(\tau)}}{P_t^{(\tau)}} + \left[W_t^{(\tau)} - Z_t^{(\tau)} + \eta(\tau)A_t\right]i_t\,\mathrm{d}t - \eta(\tau)A_t\frac{\mathrm{d}P_t^{(\tau)}}{P^{(\tau)}},$$

and the risky fund faces an analogous budget constraint. Each preferred habitat fund also transfers profits or losses to the household each period.

Interest rate and payoff processes.—To close the model, we need to specify how the instantaneous interest rate i_t and the payoff process d_t are determined. We assume that the central bank sets the nominal interest rate i_t according to the following policy rule:

$$\mathrm{d}i_t = -\kappa_i(i_t - \phi_\pi\pi_t - \phi_x x_t - i^*)\mathrm{d}t + \sigma_i\,\mathrm{d}B_{i,t}, \quad (15)$$

where the response of the policy rate to inflation and the output gap are governed by ϕ_π, ϕ_x, κ_i is a mean reversion parameter, and i^* is the central bank's target policy rate. Note that if $\kappa_i \to \infty$, we recover a standard Taylor rule.

Finally, because risky assets are in zero net supply ($\tilde{S}_t^{(\tau)} = 0$), they can be interpreted as synthetic securities that are created by the financial sector (since arbitrageurs and preferred habitat funds take nonzero positions in risky assets across all maturities). We assume that the payoff processes for these synthetic securities evolves according to

$$\mathrm{d}d_t = -\kappa_d(d_t - \psi_x x_t - \psi_\pi\pi_t - d^*)\,\mathrm{d}t + \sigma_d\,\mathrm{d}B_{d,t}. \quad (16)$$

This reduced-form approach to modeling payoffs of risky assets allows us to capture salient features of how private borrowing rates fluctuate over the business cycle. The process is mean reverting (with inertia κ_d) and may depend on movements in output and inflation through parameters ψ_x, ψ_π (which are primitives of the model). Outside of the model, assets besides financial institution securities influence the borrowing decisions of the household sector. In reality, this also includes risky private bonds, mortgage debt, equities, and a host of other assets. Rather than complicate the model by specifying a corporate structure for firms, we collapse all of these sources of payoff risk into a single process $D_t \equiv e^{d_t}$, which we allow to respond in a flexible manner to aggregate dynamics.[12]

B. *Equilibrium*

The household and firm problems are highly similar to a textbook New Keynesian model; detailed derivations are in appendix C. The optimizing behavior of these agents and market clearing give rise to a Phillips curve and IS curve for inflation and the output gap π_t, x_t, where the output gap is the deviation of output from the natural level that would prevail in the absence of nominal rigidities and fluctuations in desired markups ($\theta = 0$, $\varepsilon_t = \bar{\varepsilon}$). The only modification to the standard (linearized) New Keynesian dynamics is that the effective borrowing rate $\hat{r}_t$ replaces the policy rate i_t. We have

$$E_t \mathrm{d}\pi_t = (\rho\pi_t - \delta x_t - z_{\pi,t})\mathrm{d}t, \tag{17}$$

$$E_t \mathrm{d}x_t = \varsigma^{-1}(\hat{r}_t - \pi_t - \bar{r} - z_{x,t})\mathrm{d}t, \tag{18}$$

where the model is linearized around a steady state with $\bar{\pi} = \bar{x} = 0$. The parameter δ measures the aggregate degree of price rigidity, and $\bar{r}$ is the natural borrowing rate. The aggregate supply and demand shocks are assumed to evolve according to

$$\mathrm{d}z_{\pi,t} = -\kappa_{z_\pi} z_{\pi,t}\,\mathrm{d}t + \sigma_{z_\pi}\mathrm{d}B_{z_\pi,t}, \tag{19}$$

$$\mathrm{d}z_{x,t} = -\kappa_{z_x} z_{x,t}\,\mathrm{d}t + \sigma_{z_x}\,\mathrm{d}B_{z_x,t}. \tag{20}$$

The parameters κ_{z_π}, κ_{z_x} govern the persistence in these processes, and the terms $B_{z_\pi,t}$, $B_{z_x,t}$ are standard independent Brownian motions, with respective volatility terms σ_{z_π}, σ_{z_x}.

We collect the state variables, jump variables, and Brownian terms into vectors $y_t = [\, i_t \;\; d_t \;\; \cdots \;\; \beta_t^k \;\; \cdots \,]^\top$, $\mathbf{x}_t = [\, \pi_t \;\; x_t \,]^\top$, and $\mathbf{B}_t = [\, B_{i,t} \;\; B_{d,t} \;\; \cdots \;\; B_{\beta^k,t} \;\; \cdots \,]^\top$,

[12] Our results are not sensitive to this precise microfoundation of the payoff process. We could have chosen a formulation whereby the payoff process D_t captures claims on a portion of firm profits not paid directly and immediately to households. Under the zero net supply assumption and with the payoff process given by eq. (16), the macroeconomic dynamics of the model are identical. Although this formulation brings the interpretation of D_t closer to corporate bonds, it does not explicitly model the financing frictions leading to this distribution of profits (and our model does not feature firm investment).

respectively. The following Lemma describes the aggregate dynamics of the model, taking as given the effective borrowing rate. All proofs are in appendix A.

LEMMA 1 (Aggregate dynamics). Suppose that the effective borrowing rate drift is given by

$$\hat{r}_t = \hat{\mathbf{A}}^\top \mathbf{y}_t + \hat{C}. \tag{21}$$

Then the rational expectations equilibrium is given by

$$\mathrm{d}\mathbf{y}_t = -\boldsymbol{\Gamma}(\mathbf{y}_t - \bar{\mathbf{y}})\,\mathrm{d}t + \boldsymbol{\sigma}\,\mathrm{d}\mathbf{B}_t, \tag{22}$$

$$\mathbf{x}_t - \bar{\mathbf{x}} = \boldsymbol{\Omega}(\mathbf{y}_t - \bar{\mathbf{y}}), \tag{23}$$

where the matrices $\boldsymbol{\Gamma}$, $\boldsymbol{\Omega}$ are a function of the eigendecomposition of the linearized dynamics of the model (and therefore functions of $\hat{\mathbf{A}}$).

If $\hat{\mathbf{A}} = \mathbf{e}_i$, the vector that "selects" the policy rate $\mathbf{e}_i^\top \mathbf{y}_t = i_t$, then the effective borrowing rate responds one-for-one with the policy rate i_t and the aggregate dynamics of the model reduce to a standard New Keynesian model.[13]

Next, we turn to characterizing the behavior of asset prices. We focus on a solution to the model in which (log) asset prices are affine functions of the state variables, given by (endogenous) coefficient functions:

$$-\log P_t^{(\tau)} = \mathbf{A}(\tau)^\top \mathbf{y}_t + C(\tau), -\log \tilde{P}_t^{(\tau)} = \tilde{\mathbf{A}}(\tau)^\top \mathbf{y}_t + \tilde{C}(\tau). \tag{24}$$

LEMMA 2 (Asset prices). Suppose that the state variables $\mathbf{y}_t$ evolve according to equation (22). Then the affine coefficients in equation (24) are given by

$$\mathbf{A}(\tau) = [\mathbf{I} - e^{-\mathbf{M}\tau}]\mathbf{M}^{-1}\mathbf{e}_i, \tilde{\mathbf{A}}(\tau) = [\mathbf{I} - e^{-\mathbf{M}\tau}]\mathbf{M}^{-1}\mathbf{e}_i - e^{-\mathbf{M}\tau}\mathbf{e}_d, \tag{25}$$

where $\mathbf{e}_i$, $\mathbf{e}_d$ are vectors such that $\mathbf{e}_i^\top \mathbf{y}_t = i_t$ and $\mathbf{e}_d^\top \mathbf{y}_t = d_t$. The matrix $\mathbf{M}$ solves the fixed point problem:

$$\begin{aligned}\mathbf{M} = \boldsymbol{\Gamma}^\top - a\int_0^T &[-\alpha(\tau)\mathbf{A}(\tau) + \boldsymbol{\Theta}(\tau)]\mathbf{A}(\tau)^\top \\ &+ \left[-\tilde{\alpha}(\tau)\tilde{\mathbf{A}}(\tau) + \tilde{\boldsymbol{\Theta}}(\tau)\right]\tilde{\mathbf{A}}(\tau)^\top \mathrm{d}\tau\boldsymbol{\sigma}\boldsymbol{\sigma}^\top,\end{aligned} \tag{26}$$

where $\boldsymbol{\Theta}(\tau)$, $\tilde{\boldsymbol{\Theta}}(\tau)$ stack the habitat demand functions $\theta^k(\tau)$, $\tilde{\theta}^k(\tau)$ into vectors.

The matrix $\mathbf{M}$ can be thought of as the risk-adjusted dynamics of the state. Note that when arbitrageurs are risk neutral ($a = 0$), we have

[13] In this case, the central bank can ensure determinacy when the response to inflation $\phi_\pi > 1$. However, in general $\hat{\mathbf{A}} \neq \mathbf{e}_i$ and the determinacy conditions are more complicated. In our quantitative exercises, we verify that these determinacy conditions are satisfied numerically.

$\mathbf{M} = \boldsymbol{\Gamma}^{\top}$. However, when $a \neq 0$, $\mathbf{M}$ appears on both sides of equation (26) (through the affine coefficients $\mathbf{A}(\tau)$, $\tilde{\mathbf{A}}(\tau)$).

With the results in lemmas 1 and 2, we can characterize the equilibrium of the model.

Proposition 1 (Existence and uniqueness). An affine equilibrium is one in which the state and jump variables evolve according to equations (22) and (23) and asset prices are determined by the solution to the expressions (25) and (26). In this case, the effective borrowing rate is given by equation (21), where $\hat{\mathbf{A}}$ solves the fixed point problem

$$\hat{\mathbf{A}} = \mathbf{e}_i + (\boldsymbol{\Gamma}^{\top} - \mathbf{M})\int_0^T \left[\eta(\tau)\mathbf{A}(\tau) + \tilde{\eta}(\tau)\tilde{\mathbf{A}}(\tau)\right]\mathrm{d}\tau. \tag{27}$$

In a neighborhood of risk neutrality ($a \approx 0$), the equilibrium exists and is (locally) unique.

Note that the dynamics matrix of the state $\boldsymbol{\Gamma}$ depends on the effective borrowing rate coefficients $\hat{\mathbf{A}}$, which itself is a function of the risk-adjusted dynamics matrix $\mathbf{M}$. Thus, equilibrium is determined as a fixed point that produces asset price dynamics consistent with equilibrium dynamics of the macroeconomy and vice versa. In general, an affine equilibrium of this type may not exist, or there may be multiple solutions to this fixed point problem. However, when $a = 0$, the model reduces to a standard New Keynesian model. The result in proposition 1 shows that this equilibrium persists and is locally unique as we depart from risk neutrality. We use this insight in our numerical continuation solution algorithm, described in appendix D.

C. *Localization and the Channels of Quantitative Easing*

Our model makes precise the channels through which QE may have financial and macroeconomic effects. The immediate and direct effect of a QE purchase shock is a reduction in the amount of assets held by arbitrageurs through market clearing (a decline in $X_t^{(\tau)}$ or $\tilde{X}_t^{(\tau)}$ in eq. [11], depending on the type of assets purchased during QE). If arbitrageurs are risk neutral ($a = 0$), we recover a typical neutrality result: QE changes the arbitrageurs' portfolio but not asset prices (and thus has no affect on real activity).

On the other hand, when arbitrageurs are risk averse ($a > 0$), the change in arbitrageurs' portfolio implies changes in risk exposure, which is priced. We can broadly distinguish between three main sources of risk: duration or short-rate risk (from fluctuations in i_t), payoff risk (from fluctuations in d_t), and habitat demand risk (from fluctuations in demand factors β_t^k). By reducing arbitrageur exposure to these sources of risk, in equilibrium QE reduces the risk compensation arbitrageurs require, which manifests as a decline in expected returns. As a result, the household effective borrowing rate $\hat{r}_t$ falls, which through the intertemporal decisions of the household leads to lower savings and a jump in consumption and, from the pricing decisions of firms, leads to an increase in inflation.

However, the precise mapping from a QE shock (or private demand shocks) to the repricing of risk depends qualitatively on the risk-bearing capacity of arbitrageurs. Our model helps clarify the intuition behind the empirical localization results found in section III. Intuitively, when arbitrageurs are nearly risk neutral ($a \approx 0$), macroeconomic fundamentals affecting the path of the short rate (duration risk) are the dominant factor in determining the term structure of interest rates. Hence, when arbitrageurs hold bonds, short-rate fluctuations are their main source of risk. Demand shocks that reallocate bonds away from arbitrageurs reduce their exposure to short-rate risk and hence decrease the compensation arbitrageurs require to hold bonds. Since all bonds are sensitive to short-rate risk, any such demand shock will push down yields of all bonds. Importantly, this mechanism is independent of the location (in maturity space) of the demand shock.

As arbitrageur risk aversion increases ($a > 0$), habitat demand shocks become more prominent as additional sources of risk. Arbitrageurs try to limit their exposure to these sources of risk, leading to less propagation from the location of the demand shock to other parts of the term structure. Arbitrageurs become less willing to integrate bond markets across maturities, and hence the response of the yield curve becomes more localized around the location (in maturity space) of a given demand shock.

We can now use the model to formalize the logic of the localization effects documented in section III. Take two demand factors: a short maturity factor β_t^s and a long maturity factor β_t^ℓ with the same overall magnitude across maturities ($\int_0^T \theta^s(\tau)\mathrm{d}\tau = \int_0^T \theta^\ell(\tau)\mathrm{d}\tau$), but the short factor is more concentrated in bonds of short maturities relative to the long factor ($\exists\ \tau' : \theta^s(\tau) < \theta^\ell(\tau) \Longleftrightarrow \tau > \tau'$). In this context, the localization hypothesis involves the differential responses of the entire yield curve to movements in the short and long demand factors. We formally state a version of the localization hypothesis:

$$\text{if } a \approx 0 : \frac{\partial y_t^{(\tau)} / \partial \beta_t^s}{\partial y_t^{(\tau^*)} / \partial \beta_t^s} \approx \frac{\partial y_t^{(\tau)} / \partial \beta_t^\ell}{\partial y_t^{(\tau^*)} / \partial \beta_t^\ell}, \tag{28}$$

$$\text{if } a \gg 0 : \tau > \tau^* \Longleftrightarrow \frac{\partial y_t^{(\tau)} / \partial \beta_t^s}{\partial y_t^{(\tau^*)} / \partial \beta_t^s} < \frac{\partial y_t^{(\tau)} / \partial \beta_t^\ell}{\partial y_t^{(\tau^*)} / \partial \beta_t^\ell}, \tag{29}$$

where τ^* is an arbitrary baseline maturity (and unrelated to τ) and $\partial y_t^{(\tau)} / \partial \beta_t^k$ is the response of τ-maturity yields to the demand shock $k \in \{s, \ell\}$. When arbitrageur risk-bearing capacity is high ($a \approx 0$), the relative response of the yield curve is the same for both factors. On the other hand, when risk-bearing capacity is low ($a \gg 0$), then long demand shocks have relatively larger effects on long-maturity yields and vice versa for short demand shocks. These expression make statements about the movements of the yield curve to short and long demand shocks relative to movements in some fixed maturity τ^*.[14] Thus, the model maps to our specification (4), and the qualitative predictions of the model line up with our results. In section IV.D, we move from the qualitative to quantitative features of the model.

D. Model Calibration

To study the quantitative properties of the model, we need to pick functional forms and assign values to parameters. First, we must choose functional forms regarding the effective borrowing weights and the habitat elasticity and demand functions. We assume similar exponential functions as in Vayanos and Vila (2021):

$$\text{habitat elasticity function: } \alpha(\tau) = \alpha_0 \exp(-\alpha_1 \tau), \tag{30}$$

$$\text{habitat demand functions: } \theta^k(\tau) = \theta_0 \tau (\theta_1^k)^2 \exp(-\theta_1^k \tau), \tag{31}$$

$$\text{effective borrowing weights: } \eta(\tau) = \eta_0 \tau \eta_1^2 \exp(-\eta_1 \tau), \tag{32}$$

with analogous functional forms for $\tilde{\alpha}(\tau)$, $\tilde{\theta}^k(\tau)$, $\tilde{\eta}(\tau)$ (and $\tilde{\eta}_0 = 1 - \eta_0$ so that the weights sum to 1). Equation (30) implies that the habitat investor elasticity with respect to (log) price is declining with maturity and is single peaked with respect to yields. Similarly, equations (31) and (32) imply that the demand factor and borrowing weights are single-peaked functions. In addition to improving numerical properties of the model, these exponential functional forms allow some flexibility in capturing key modeling features, such as demand shocks targeted in specific areas of maturity space, without significantly increasing the dimensionality of the problem. The parameters α_0, θ_0 govern the overall size of the habitat elasticity and demand functions, while α_1, θ_1^k, η_1 govern the shape as a function of maturity: a lower value of these parameters implies that more of the weight of these functions lies at longer maturities.

Next, to map the model to our empirical estimates in section III, we assume that there are three demand factors, corresponding to short- and long-maturity Treasury auctions, and a risky asset. We denote these demand factors by β_t^s, β_t^ℓ, and $\tilde{\beta}_t$, respectively. In order to impose more discipline on

[14] Appendix E discusses sufficient conditions and derives further localization predictions of the model.

the model, we assume that the habitat demand shocks are identical, except for the shape of the demand factor functions $\theta^s(\tau)$, $\theta^\ell(\tau)$, and $\tilde{\theta}(\tau)$. Specifically, we set $\theta_1^s = 0.5$ and $\theta_1^\ell = 0.2$ to match our regression analysis in section III. These choices imply that the short factor is concentrated in short maturities less than 5 years (and the mean of $\theta^s(\tau)$ is 4 years), while the long factor has more weight in intermediate and long maturities above 7 years (and the mean of $\theta^\ell(\tau)$ is 10 years). We set the parameter $\alpha_1 = 0.1$ such that the habitat elasticity with respect to yields ($\equiv \tau \cdot \alpha(\tau)$) is maximized for the 10-year maturity, since this is often taken as a key benchmark yield by investors. We further assume that the habitat demand and elasticity functions are identical for Treasuries and risky assets: $\alpha(\tau) = \tilde{\alpha}(\tau)$ and $\theta^s(\tau) = \tilde{\theta}(\tau)$. Finally, we assume that the dynamics parameters κ_β, $\phi_{i,\beta}$, σ_β for each demand factor in equation (14) are identical.

An important input into our model is the effective borrowing rate weights $\eta(\tau)$, $\tilde{\eta}(\tau)$. If household borrowing responds only to movements in the instantaneous short rate, then the macroeconomic dynamics of the model are identical to a textbook New Keynesian model. Thus, the effective borrowing weights can be thought of broadly as capturing the sensitivity of the macroeconomy to borrowing rates besides the federal funds rate. For our baseline calibration, we assume that the effective borrowing rate depends on only risky rate ($\eta_0 = 0$ and hence $\eta(\tau) = 0$ and only $\tilde{\eta}(\tau)$ is nonzero). Further, we set $\eta_1 = 2$, such that $\tilde{\eta}(\tau)$ peaks at 0.5 years and has an average value of 1 year. Our baseline choice puts a heavy weight on short and intermediate maturities because trading volume in debt markets is concentrated at shorter-term maturities. For instance, estimates from the Securities Industry and Financial Markets Association (SIFMA) show that over the past two decades, trading volume for Treasury securities with maturity of 3 years or less made up about 45% of total trading volume on average (and this ratio was relatively stable, with a minimum and maximum annual value of 41% and 52%, respectively). Even within coupon-bearing securities (excluding T-bills), trading volume for Treasury notes with maturity of 3 years or less made up 33% of the trading volume in coupon-bearing Treasury notes and bonds. In comparison, trading volume for Treasury bonds with a maturity of greater than 11 years made up only 8% of the trading volume in coupon-bearing Treasury notes and bonds. Trading volume for corporate debt markets across maturities is less readily available. However, comparing trading volume in corporate bond markets and commercial paper markets (from SIFMA and the Federal Reserve, respectively) further suggests an outsized role of short-maturity rates. For instance, for 2020–2021 the trading volume in commercial paper markets was on average 4 times that of trading volume in corporate debt markets (for securities with maturity of at least one year). In short, our calibration implies that household borrowing is a function of risky borrowing rates (rather than the policy rate), largely concentrated at short and intermediate maturities. We explore the sensitivity of our results to alternative calibrations below.

Finally, in order to reduce the number of parameters needed to be calibrated, we simplify and assume that the central bank responds to only inflation ($\phi_x = 0$ in eq. [15]), and the payoff process responds to only the output gap ($\psi_\pi = 0$ in eq. [16]). We explore sensitivity to these assumptions below.

We jointly estimate the remaining parameters through a moment-matching exercise.[15] We target volatility and cross correlations of Treasury and corporate yields $y_t^{(\tau)}$, $\tilde{y}_t^{(\tau)}$ as well as inflation and the output gap π_t, x_t.[16] Specifically, we target the following set of moments: (1) the volatility of 1-year Treasury yields, 1-year corporate-Treasury spreads, the output gap, and inflation (in levels, 1-year changes, and 1-month changes); (2) the correlation of inflation and 1-year Treasury yields, the correlation of output gap and 1-year corporate Treasury spreads, and the correlation of inflation and the output gap (in levels); (3) the volatility of the entire term structure of Treasury yields (1-year changes and 1-month changes); (4) the correlation of the entire term structure of Treasury yields and 1-year Treasury yields (1-year changes and 1-month changes); and (5) the localization regression coefficients from figure 4. Besides the localization regression coefficients, the data are monthly. Inflation is defined as the year-over-year log change in the Personal Consumption Expenditures Price Index. The output gap is defined as the cyclical component of industrial production using a Hodrick-Prescott filter. Zero coupon Treasury yields are from Gürkaynak, Sack, and Wright (2007). Zero coupon corporate bond yields are from the Treasury's High Quality Market (HQM) Corporate Bond Yield Curve data. We utilize the HQM data because it is a long-running measure of the term structure of zero coupon private borrowing rates that are important for macroeconomic activity; however, these rates reflect relatively safe corporate rates (rated A or above). Besides the localization coefficients, moments are computed starting in

[15] Collecting unknown parameters into a vector $\boldsymbol{\rho}$, we estimate the model by choosing $\hat{\boldsymbol{\rho}}$ to minimize the weighted sum of squares: $(L(\boldsymbol{\rho}) = \Sigma_{n=1}^{N} w_n(\hat{m}_n - m_n(\boldsymbol{\rho}))^2)$, where $\{\hat{m}_n\}_{n=1}^{N}$ are empirical moments and $\{m_n(\boldsymbol{\rho})\}_{n=1}^{N}$ are model-implied counterparts as a function of $\hat{\boldsymbol{\rho}}$. The terms $\{w_n\}_{n=1}^{N}$ are weights placed on each target moment. These weights are set to $1/\mathcal{N}_T$ for moments that are a function of maturity, where $\mathcal{N}_T$ is the number of maturities, and 1 otherwise, so that our estimates are not overly influenced by maturity-based moments. Appendix D derives the model-implied moments.

[16] The habitat elasticity functions and demand functions enter the solution multiplicatively with risk aversion and demand shock volatility, so they are identified only up to a scaling factor. If shocks to quantity demand for Treasuries and risky assets can be measured directly across all periods (including outside auction times), one can separately identify these parameters. Because such measures are not available, the literature (e.g., Vayanos and Vila 2021) typically chooses some normalization, such as $\sigma_\beta = 1$. In our auction data, the bid-to-cover ratio captures this information partially; this information is not necessary for our calibration strategy, but we utilize this information in our QE exercises.

1986 (when 30-year yields are available in Gürkaynak, Sack, and Wright 2007).

By targeting the localization regression coefficients from figure 4, our model is informed by our empirical localization results, which are well suited to identify the preferred habitat parameters of our model. We allow only the habitat risk-adjusted parameters ($a \cdot \alpha_0$, $a \cdot \sigma_\beta \cdot \theta_0$, and $a \cdot \phi_{i,\beta}$) to be different in crisis, which we target to match the crisis localization regression coefficients from figure 4. All other parameters are assumed to be the same across crisis and noncrisis periods. Given that the number of observations in the noncrisis period is vastly larger than the number of observations in the crisis period, we first estimate all of the parameters from the noncrisis sample.

The fit of the model is summarized in figure 6 (moments that depend on maturity) and panel A of table 5 (additional moments). Panel B of table 5 reports our baseline parameter values. Our parameterized model not only picks up the qualitative features of the data but also is successful at matching the data quantitatively. Besides matching a wide variety of volatility and correlation moments across borrowing rates and aggregate variables, the model localization coefficients are close to the data localization coefficients in crisis and noncrisis periods (fig. 6*A*–6*D*).[17] Nonetheless, given the challenges associated with estimating dynamic stochastic general equilibrium models, one should treat this parameterization with caution, and we run extensive sensitivity analysis in section IV.H.

E. *Response of the Yield Curve to QE1*

To assess the contribution of preferred habitat theory to the observed reaction of yields to QE, we feed a QE1 shock into our model, compute predicted responses of yields at different maturities, and compare predictions to actual changes in yields. We focus on QE1 because it was arguably the "cleanest" QE shock: there is a clear set of policy announcement events, and the observed response to these events is unlikely to be plagued by anticipation issues relative to later rounds of QE. But in order to accurately capture the relevant features of QE1, we do not treat QE as precisely equivalent to a habitat demand shock in our model. Besides

[17] We also examine localization across asset classes. Figure B26 plots various measures of risky borrowing rate localization in the model. We find evidence in the model of increased maturity localization in crisis periods: the relative effects of our short and long riskless demand shocks $\beta_t^{(s)}$, $\beta_t^{(\ell)}$ have larger differential effects on risky rates across maturities in the model calibrated to the crisis period compared with the noncrisis period (consistent with panel B of table 4). Additionally, we find that compared with the response of the riskless yield curve, the pass-through of riskless demand shocks to risky borrowing rates decreases in the crisis calibration (consistent with panel A of table 4). However, unlike our baseline Treasury localization results, the mapping from our empirical results to the model is not as precise, and so these results should be taken as qualitative.

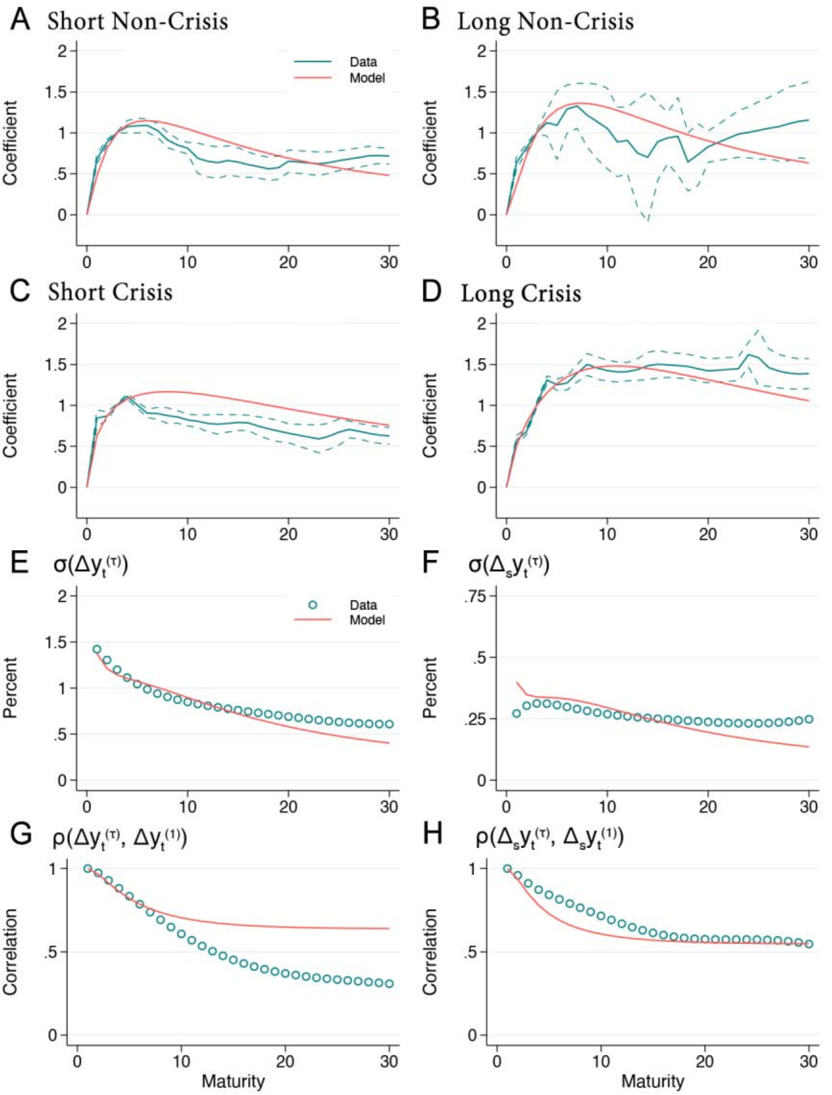

Fig. 6.—Model fit. *A–D* plot the model-implied and empirical coefficients from regression equation (4). *A* and *B* plot the coefficients from the model calibrated to noncrisis periods, while *C* and *D* plot the coefficients from the model calibrated to crisis periods. *E–H* compare additional targeted volatility and correlation moments in the model (solid lines) with the data (scatter points) as a function of maturity.

TABLE 5
Model Calibration and Targeted Moments

Target Moment/Parameter	(1)	(2)
	A. Matched Moments	
	Data	Model
$\sigma(y_t^{(1)})$	1.993	1.978
$\sigma(\tilde{y}_t^{(1)} - y_t^{(1)})$	.542	.549
$\sigma(\pi_t)$	.973	.882
$\sigma(x_t)$	1.992	1.996
$\sigma(\Delta y_t^{(1)})$	1.423	1.376
$\sigma(\Delta(\tilde{y}_t^{(1)} - y_t^{(1)}))$	.615	.617
$\sigma(\Delta\pi_t)$	.900	.935
$\sigma(\Delta x_t)$	2.190	2.137
$\sigma(\Delta_s y_t^{(1)})$	.272	.399
$\sigma(\Delta_s(\tilde{y}_t^{(1)} - y_t^{(1)}))$	.258	.218
$\sigma(\Delta_s\pi_t)$	.232	.362
$\sigma(\Delta_s x_t)$	.499	.745
$\rho(\tilde{y}_t^{(1)} - y_t^{(1)}, x_t)$	−.171	−.158
$\rho(y_t^{(1)}, \pi_t)$	.549	.615
$\rho(\pi_t, x_t)$	.268	.271
	B. Calibrated Parameters	
	Value (Crisis)	Description
σ_i	2.567	Monetary policy volatility
κ_i	1.082	Monetary policy inertia
σ_d	1.116	Risky payoff volatility
κ_d	.879	Risky payoff inertia
$\sigma_{z,\pi}$	2.039	Cost-push shock volatility
$\kappa_{z,\pi}$	.801	Cost-push shock inertia
$\sigma_{z,x}$	1.749	Aggregate demand shock volatility
$\kappa_{z,x}$	.253	Aggregate demand shock inertia
ϕ_π	3.096	Inflation Taylor coefficient
ψ_x	.393	Risky payoff output coefficient
δ	.705	Nominal rigidity
ρ	.04	Discount factor
ς^{-1}	1.00	Intertemporal elasticity
κ_β	1.367	Habitat demand inertia
$a \cdot \alpha_0$	.008 (.018)	Habitat elasticity size
$a \cdot \sigma_\beta \cdot \theta_0$	2.509 (5.123)	Habitat demand size
$a \cdot \phi_{i,\beta}$	.491 (4.620)	Habitat demand short rate response
θ_1^s	.50	Short treasury factor maturity weight
θ_1^ℓ	.20	Long treasury factor maturity weight
$\tilde{\theta}_1$	.50	Risky factor maturity weight
α_1	.10	Habitat elasticity maturity weight
	C. QE/QT Parameters	
	Value	Description
θ_1^{QE}	.35	QE1 maturity weight
κ_{QE}	.20	QE1 inertia
θ_1^{QT}	.50	QT maturity weight
$\kappa_{QT,A}$	.20	QT inertia, active component

TABLE 5 (*Continued*)

Target Moment/Parameter	(1)	(2)
$\kappa_{\mathrm{QT,P}}$	2.25	QT inertia, passive component
$\gamma_{\mathrm{QT,A,P}}$	1.75	QT passive component response

NOTE.—Panel A presents the targeted moments that are not a function of maturity that we use to calibrate the model. The targeted moments are volatility and correlations across short-term rates and macroeconomic variables. Δ denotes a 1-year change, while Δ_s denote a 1-month change. $\sigma(\cdot)$ denotes standard deviation (reported in percentage points), while $\rho(\cdot,\cdot)$ denotes correlation. Panel B presents calibrated parameters used in our quantitative exercises. Panel C reports the parameterization of our QE and QT shocks.

the obvious fact that QE1 was significantly larger than typical private demand shocks, QE1 also occurred unexpectedly, and markets had different expectations about the dynamic properties of QE1 compared with typical private demand shocks. Moreover, QE1 simultaneously purchased Treasuries and MBS, and the profile of maturities was different than in auctions. Thus, in our model QE1 is not a simple rescaling of the habitat demand shocks during auctions. We discuss in detail how we mimic the actual QE1 shock in the model.

Persistence.—We assume that this shock was completely unexpected (MIT shock) and that afterward markets expected purchases to be unwound slowly and deterministically:

$$d\beta_t^{\mathrm{QE}} = -\kappa_{\mathrm{QE}}\beta_t^{\mathrm{QE}}\, dt. \tag{33}$$

We set the inertia parameter $\kappa_{\mathrm{QE}} = 0.2$. This magnitude roughly implies that markets expected the Fed to unwind its purchases somewhat faster than holding to maturity (more precisely, the half-life of the purchases is roughly 3.5 years). Ex post, the MBS holdings roughly followed this process until the reintroduction of MBS purchases during QE3. On the other hand, the Treasuries purchased as part of the Fed's QE programs were held on the balance sheet for a very long time, and holdings remained elevated well beyond that implied by our parameterization. However, this does not imply that markets expected this ex ante.[18] We explore the sensitivity of this assumption below.

[18] The actual purchases of MBS and Treasuries during QE1 were planned to take place over the 6 months following the March 2009 announcement. The Federal Open Market Committee (FOMC) statements during this time do not make any reference to selling these securities off but state that the FOMC will "carefully monitor the size and composition of the Federal Reserve's balance sheet in light of evolving financial and economic developments." Eventually, the FOMC made clear their policy of Treasury reinvestment. For instance, in Chairman Bernanke's July 2010 report to Congress, he stated that "the proceeds from maturing Treasury securities are being reinvested in new issues of Treasury securities with similar maturities."

Composition.—QE1 involved purchasing a total of roughly $1.25 trillion MBS and $300 billion Treasuries, concentrated on intermediate and long-term maturity purchases (over 5 years).[19] Given this focus, we assume the same functional form as in equation (31) for the QE demand function $\theta^{\mathrm{QE}}(\tau)$ and set $\theta_1^{\mathrm{QE}} = 0.35$, such that the model-implied QE1 purchases are of relatively long-term maturity but in between the maturities of the preferred habitat short and long factor described above (hence the mean of $\theta^{\mathrm{QE}}(\tau)$ is roughly 6 years). These purchases are split up such that roughly 80% of QE1 purchases are risky assets in the model, while the remaining 20% are of safe bonds, in order to match the fraction of actual MBS and Treasury purchases during QE1.

Size.—Matching the overall magnitude of QE1 requires a few additional steps. As discussed above, our calibration strategy does not separately identify the size of preferred habitat demand shifts and arbitrageur risk aversion. This allows us to compute the response of yields to a unit demand shock $\partial y_t^{(\tau)}/\partial \beta_t^k$ for short- or long-term auctions. However, our moment-matching exercise thus far provides no information about the dollar magnitude of a unit demand shock in the model, since we can only estimate the product $a \cdot \sigma_\beta \cdot \theta_0$. Our model is meant to capture all the volatility of demand shocks, not just those that occur during auctions. But by utilizing additional information from Treasury auctions, we can pin down the dollar magnitude of demand shocks.

The results of our alternative localization specification in figure 5 measures the response of the yield curve to a unit change in the bid-to-cover ratio across short-term and long-term auctions and in normal versus crisis times. During the QE1 period, a unit change in the bid-to-cover ratio corresponded to about $30 billion for long-term auctions, while short-term auctions were 50% larger. Thus, we estimate this quantity of long-term auctions (which we denote by $\Delta\beta_t^{(\mathrm{auc})}$) by targeting the results of our alternative localization regression; results are in figure B25. Finally, since QE1 was roughly $1.5 trillion, we have that $\Delta\beta_t^{(\mathrm{QE})} \approx 50 \times \Delta\beta_t^{(\mathrm{auc})}$. To be clear, this does not imply that the yield curve response to QE1 is hardwired to be the same as a (scaled-up) auction demand shock. The maturity composition, asset composition, and the stochastic properties of the QE1 shock in the model differ from the estimated demand factors. Thus, the yield response to the QE shock versus the demand factors will not be proportional: $\partial y_t^{(\tau)}/\partial \beta_t^{(\mathrm{QE})} \not\propto \partial y_t^{(\tau)}/\partial \beta_t^k$.

Figure 7*A* plots the cumulative change observed in the data and predicted in the model. We use the cumulative change in the yield curve reported in Krishnamurthy and Vissing-Jorgensen (2011) following the five

[19] In our analysis, we exclude purchases of agency debt (which accounts for a small fraction of the total purchases during QE1), given the complexities associated with the health of Freddie Mac and Fannie Mae.

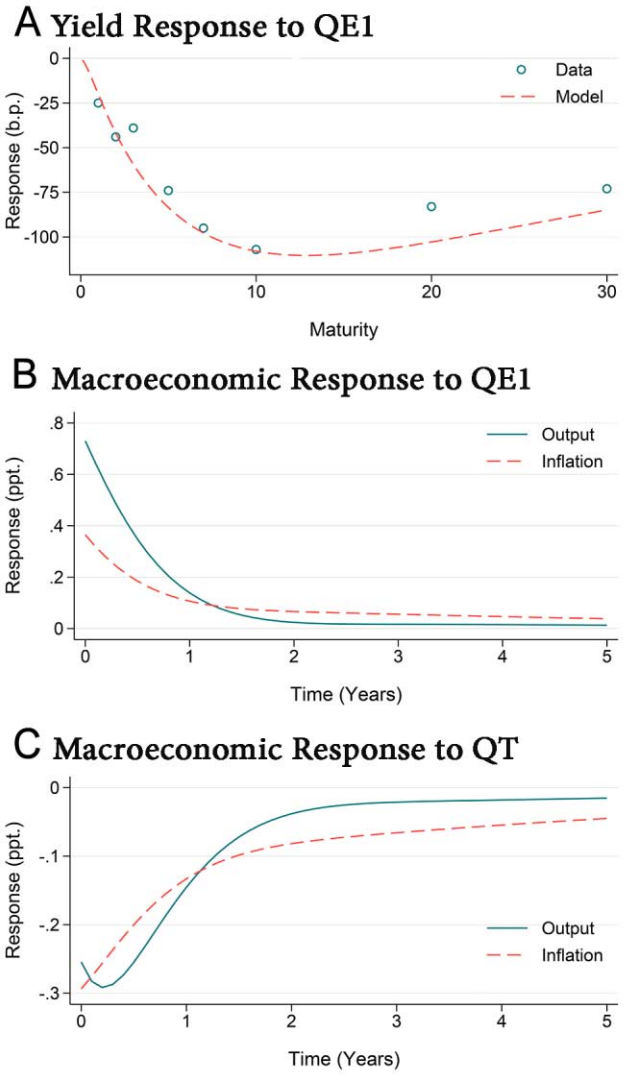

FIG. 7.—Model-implied response to QE1 and QT. *A* compares the observed (cumulative) yield reactions and the model-implied change in yields following a QE1 shock (in basis points). *B* and *C* plot the model-implied impulse response functions of output and inflation following QE1 and QT, respectively (in percentage points).

major QE1 event dates as the empirical response of the yield curve to QE1. The remarkable consistency between the responses suggests that the actual market reaction to QE1 announcements aligns with the predictions of our model in response to observed shifts in private demand for Treasuries. This finding implies that the net effect of other channels of

QE (e.g., inflation expectations, forward guidance, signaling) could be smaller than thought before.

F. *Macroeconomic Effects of Quantitative Easing*

While there is extensive research documenting the responses of financial markets to rounds of QE, little is known about how QE affected the broader economy because QE events are so infrequent. We cannot shed more light on this using regression analysis or similar tools, but we can use our calibrated model to quantify the macroeconomic effects of QE.

Figure 7*B* shows that our baseline calibration implies that in terms of macroeconomic effects, on impact QE1 increased output by roughly 0.73 percentage points and inflation by 0.36 percentage points. These effects then monotonically fall toward zero over the next few years. The cumulative effects are $\int_0^\infty E_0[x_t]\mathrm{d}t \approx 0.56$ percentage points and $\int_0^\infty E_0[\pi_t]\mathrm{d}t \approx 0.64$ percentage points. For comparison, these stimulative effects of QE1 are comparable to a rate cut of roughly 50–75 basis points in the model, undertaken when financial markets are not distressed (and note that conventional monetary policy rate cuts have a faster degree of mean reversion in our model).[20]

G. *Comparing Quantitative Easing and Tightening*

To rein in inflation, central banks are beginning the implementation of QT (shrinking their balance sheets through asset sales or a reduction in reinvestment of maturing assets). In our linearized model, QE and QT are mirror images of one another. Differences in the macroeconomic impacts of QE and QT can arise only from differences in the financial environment or in the design of the program. These differences are salient in the context of the Fed's QT policy, which began in June 2022. First, the Fed is conducting QT in a relatively passive, back-loaded fashion: QT began with a monthly reduction of up to $30 billion in Treasuries and $17.5 billion in MBS, with both caps doubling after one quarter. The increasing pace of QT (which was announced in advance) differs from how the Fed conducted QE1. Second, the Fed's QT has focused more on a

[20] In our model, there is little difference between (1) the Fed buying assets (and thus reducing assets available to the private market) and (2) the Treasury issuing less debt (and thus reducing assets available to the private market). Thus, while comparing our estimated effect of QE1 and actual data, one should keep in mind that the Treasury increased issuance of debt by approximately $2.7 trillion between December 2008 and October 2010 (the duration of QE1) relative to the precrisis trend. The average maturity of debt issued over this period was 6.4 years. If we assume that this shock to the supply of government debt is as persistent as QE1 (i.e., debt is held to maturity), then QE1 was roughly offsetting the increased supply of Treasuries and the combined macroeconomic effect of changes in government debt due to fiscal and monetary policies was a wash.

reduction of Treasury holdings than MBS, again in contrast with QE1. Third, the maturity composition of QT is shorter than QE1. Fourth, although the Fed has not specified the size of QT, this program is likely to be larger than the purchases undertaken during QE1. Finally, risk-bearing capacity of financial markets is higher now than during QE1.

While there is substantial uncertainty associated with how the Fed will conduct QT going forward, we attempt to parsimoniously capture these salient differences within the model. We assume that the QT sales shock β_t^{QT} is composed of 65% Treasuries and 35% MBS, mimicking the composition of the Fed's QT balance sheet reduction caps. In line with the maturity composition of the Fed's balance sheet, we also assume that $\theta_1^{QT} = 0.5$ so that the maturity of the assets sold during QT is shorter than the maturity of assets purchased during QE1 (and the mean of $\theta^{QT}(\tau)$ is 4 years; fig. B27 compares the change in the Fed's Treasury holdings by maturity following QE1 and QT). Capturing the passive, back-loaded approach of QT while maintaining the mean-reverting properties necessary in the model requires some additional changes. We posit that β_t^{QT} is a mixture of an active component ($\beta_t^{QT,A}$) and a passive component ($\beta_t^{QT,P}$) such that $\beta_t^{QT} = \lambda_{QT}\beta_t^{QT,A} + (1 - \lambda_{QT})\beta_t^{QT,P}$, where

$$d\beta_t^{QT,A} = -\kappa_{QT,A}\beta_t^{QT,A}dt, \tag{34}$$

$$d\beta_t^{QT,P} = -\left(\gamma_{QT,A,P}\beta_t^{QT,A} + \kappa_{QT,P}\beta_t^{QT,P}\right)dt. \tag{35}$$

Thus, QT sales initially are driven by the active component $\beta_t^{QT,A}$, which reverts back to steady state just as our QE1 shock (and we choose the same inertia $\kappa_{QT,A} = \kappa_{QE}$). However, the passive factor $\beta_t^{QT,P}$ then kicks in, capturing the slow increase in the magnitude of QT sales following the announcement of QT (taken into account by market participants). We choose $\gamma_{QT,A,P} = 1.75$, $\kappa_{QT,P} = 2.25$, and $\lambda_{QT} = 0.25$ so that the QT shock doubles in magnitude after one quarter, peaks after 1 year, and then reverts to half of the peak after roughly 4–5 years. The overall magnitude of the initial QT shock is chosen to be double that of QE1: $\Delta\beta_t^{QT,A} = 2 \times \Delta\beta_t^{QE}$. Thus, while QT sales are smaller than QE1 purchases at the beginning of the program, QT sales eventually surpass QE1 in magnitude, remaining elevated so that the cumulative size of QT is larger than QE1.[21]

[21] Figure B28 compares the dynamics of QE1 and QT shocks in the model. The Fed has been very clear about the initial timing of the Treasury and MBS caps (e.g., see the March 2022 FOMC statement; Federal Reserve Board 2022). Going forward, the FOMC has stated that they intend to "maintain securities holdings in amounts needed to implement monetary policy efficiently and effectively in its ample reserves regime," although the FOMC has been reticent regarding what this means quantitatively. But if we take the size of the balance sheet pre-COVID as this target level, then the overall size of QT relative to gross domestic product would be roughly twice the size of QE1 (see also Ennis and Kirk

Finally, we choose the risk-adjusted parameters to be equal to 75% of our baseline QE1 calibration. On the one hand, despite increased volatility during COVID-19, many broad measures of financial distress (e.g., the Chicago Fed National Financial Conditions Index) remain significantly below the peaks observed during the Great Recession. On the other hand, some measures of intermediary health, like He, Kelly, and Manela (2016), that are more closely related to the risk-bearing capacity of arbitrageurs show signs of deterioration in the last year (see fig. B12). Thus, we split the difference for our baseline QT experiment. However, we view this as a lower bound on risk-bearing capacity (and hence the macroeconomic reactions to QT should be viewed as upper bounds).

Figure 7*C* plots the impulse response of inflation and the output gap to our QT shock. Despite the fact that QT is larger in magnitude than QE1, the macroeconomic effects on impact are somewhat smaller. For instance, the initial response of output and inflation is a decline of between 0.25 and 0.3 percentage points. However, the persistent passive component also generates a longer-lasting effect of QT compared with QE1 and implies that the response of the output gap to QT is hump shaped rather than monotonic.

The design of QT is in part responsible for the relatively muted response of output and inflation. However, the key difference between QE1 and QT is the risk-bearing capacity of financial markets. As discussed above, when risk-bearing capacity increases ($a \to 0$), our model collapses to a textbook New Keynesian model and recovers the standard neutrality results regarding asset purchases or sales.

In summary, our model suggests that QT will not induce substantial downward pressure on inflation. The differences in the model-implied effects of QE and QT motivate our subsequent sensitivity analysis for the design of these programs.

H. Alternative Designs of Quantitative Easing

We now explore how variations in structural parameters and in the implementation of QE influence the ability to move output and inflation. Because in our model output and inflation comove strongly in response to QE shocks, we focus on the sensitivity analysis for output responses (fig. 8) and relegate results for inflation to the online appendix (fig. B29). For all sensitivity analyses, we study how the macroeconomic effects vary in our baseline calibration (solid line) and also compare with counterfactual risky-only and Treasury-only QE policies (dashed and dotted lines,

2022). Hence, these calibration choices represent our best translation into a mean-reverting shock in the model. With the linearity of the model, responses scale one-for-one with the size of shocks. All other sensitivity analyses are conducted in sec. IV.H.

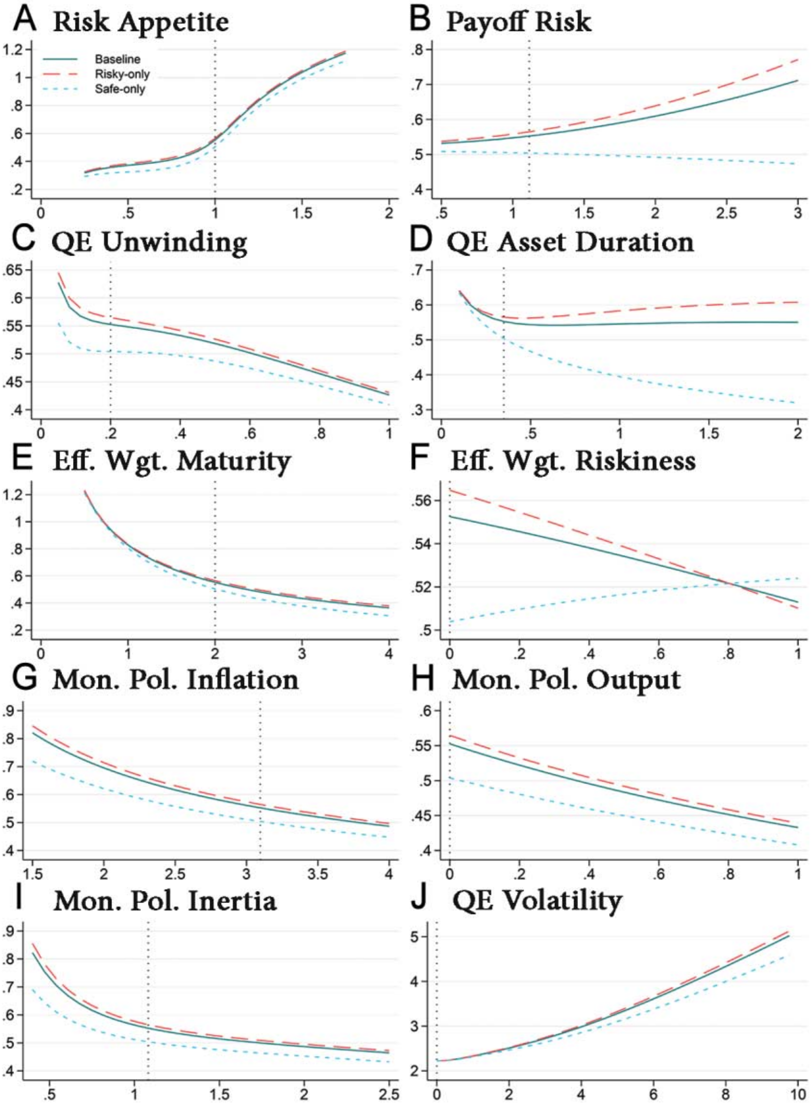

FIG. 8.—QE sensitivity analysis: output. The figure shows the response of output to a QE shock as a function of various model parameters. *A*, Arbitrageur risk aversion (a). *B*, Risky payoff uncertainty (σ_d). *C*, QE mean reversion (κ_{QE}). *D*, QE maturity duration (θ_1^{QE}). *E*, Effective borrowing maturity weights (η_1). *F*, Effective borrowing risky/safe weights (η_0). *G*, Inflation Taylor coefficient (ϕ_π). *H*, Output gap Taylor coefficient (ϕ_x). *I*, Inertia Taylor coefficient (κ_i). *J*, QE volatility (σ_{QE}). In each panel, the dotted line corresponds to the baseline QE calibration. The orange dashed line shows results for a hypothetical QE policy that purchases risky assets only; similarly, the blue dashed line is for Treasury-only purchases. *J* reports the long-run volatility $\text{Var}[x_t]$, while all other panels report the cumulative effect $\int_0^\infty E_0[x_t]\mathrm{d}t$.

respectively), which are identical to our QE1 calibration but purchase only risky or safe assets.

Risk appetite.—As we discussed above, the power of QE to affect yields at target maturities is high in crisis times and weak in noncrisis times. Consistent with this insight, figure 8*A* shows that the transmission is highly sensitive to the risk aversion of arbitrageurs. Given that our baseline calibration for crisis corresponds to the Great Recession, one has to have a financial crisis of truly unprecedented proportions to materially increase the power of QE-based tools. As a result, the balance of risks appears to be somewhat one-sided: it is difficult to raise the power of QE beyond what was achieved during the Great Recession, but it is relatively easy to reduce the power as soon as financial panics calm down. This logic suggests that QE is less effective as a conventional tool to the extent that “conventional” entails well-functioning financial markets (and was reflected in our analysis of QT). This finding is true regardless of the design of QE purchases: the counterfactual risky-only and Treasury-only QE policies have the same pattern as a function of risk aversion.

Risky asset uncertainty.—We next discuss how the transmission of QE depends on the riskiness of risky assets. Figure 8*B* plots the long-run effect of QE on output as we increase σ_d from the baseline calibration. We find that QE policy has larger macroeconomic effects as risky assets become riskier. The intuition for this can be seen by examining the counterfactual risky-only and Treasury-only QE policies. Unlike the previous sensitivity exercise, these policies act differently as a function of risky asset uncertainty. As uncertainty increases, the effect of risky asset purchases on output increases. However, the impact of Treasury purchases decreases. Intuitively, when risky asset payoff uncertainty is very low, Treasuries and risky assets are good substitutes (in the limit of no payoff risk, these two assets become perfect substitutes). In this case, risky-only and Treasury-only QE policies have similar effects on household borrowing rates, and therefore macroeconomic transmission is similar because the payoff risk extraction channel discussed in section IV.C disappears. However, as risky asset uncertainty increases, Treasuries and risky assets become less substitutable. Now, risky-only QE policies are highly effective at moving household borrowing rates (and therefore have larger macroeconomic effects), while Treasury-only QE policies have smaller effects.

Unwinding.—One practical question for policy makers is how long central banks should hold assets accumulated during QE rounds (e.g., Karadi and Nakov 2020; Sims and Wu 2020). To this end, we vary κ_{QE} in equation (33), which governs the unwinding process. Figure 8*C* shows that the effects on output nearly double when comparing very fast and very slow unwinding (the *x*-axis reports results from $\kappa_{QE} \in (0.05, 1.0)$, implying a half-life of about 14 years to three quarters, respectively). Intuitively, if the Fed makes large purchases of assets but then quickly sells those securities, we should not observe large and long-term macroeconomic effects. Rather, it is the cumulative size of QE over time that matters. Hence, policy makers should be clear about how long the central bank expects to hold the securities on its balance sheet. Providing a type of forward guidance regarding the expected path of purchases can potentially increase the immediate effectiveness of these policies.

Asset maturity.—The maturity of the assets purchased during QE is one of the key decisions policy makers must make. The *x*-axis of figure 8*D* is the maturity of the assets purchased (where a higher value implies shorter maturity). The baseline QE policy has larger macroeconomic effects as the maturity of the purchased assets increases. Long-maturity assets are riskier in terms of duration risk, and hence purchasing these assets amplifies the duration risk extraction channel. However, we again see a dichotomy between the counterfactual Treasury-only and risky-only QE policies. This is because the payoff risk extraction channel is still active when the duration of the risky asset purchases is short. Hence, while Treasury purchases become ineffective as the maturity becomes very short term, risky asset purchases still have sizable macroeconomic effects.

Effective borrowing weights.—Figures 8*E* and 8*F* explore the robustness of our assumptions regarding the effective borrowing weight function $\eta(\tau)$, which as discussed above is an important input into our model. Figure 8*E* plots the long-run macroeconomic effects of QE as the effective borrowing rate weights become more concentrated at short-term rates (larger values of the *x*-axis) or at long-term rates (smaller values of the *x*-axis). As expected, if the macroeconomy is more sensitive to longer-term borrowing rates, QE has a larger impact on output. Because the short end of the term structure is pinned down by the Fed's policy rate, QE will have larger effects on intermediate Treasury yields. Moreover, given the localization results discussed above, longer maturity purchases will have larger effects on long-term rates when the risk-bearing capacity of financial markets is relatively low.

Figure 8*F* varies the weights placed on safe and risky borrowing rates. The baseline effective borrowing rate depends entirely on risky rates (corresponding to the *x*-axis value of 0). We compare with alternative weights placed on safe borrowing rates, up until the effective rate is a function of only safe rates (*x*-axis value of 1). As expected, when household borrowing is more sensitive to safe rates, the counterfactual Treasury-only QE program becomes more powerful. However, the magnitude of the differences are small.

Interaction with conventional policy.—Next we turn to how the macroeconomic effects of QE interact with conventional policy via the standard Taylor rule. Figures 8*G*, 8*H*, and 8*F* vary the Taylor coefficients on inflation ϕ_π, the output gap ϕ_x, and inertia κ_i, respectively. In general, if markets expect the Fed to react to the macroeconomy more quickly and aggressively with the policy rate (corresponding to larger values on the x-axis), then the macroeconomic effects of QE are attenuated. However, for a wide range of coefficient values, the macroeconomic effects of QE are quantitatively similar.

QE uncertainty.—Finally, we assumed that QE1 was completely unanticipated by markets. This may be an accurate representation of the first round of QE, but over a decade later it is likely that markets now consider the possibility of future QE shocks. To model the recurrent nature of QE, we modify equation (33) and instead assume that QE shocks are described by

$$\mathrm{d}\beta_t^{\mathrm{QE}} = -\kappa_{\mathrm{QE}}\beta_t^{\mathrm{QE}}\mathrm{d}t + \sigma_{\mathrm{QE}}\mathrm{d}B_{\mathrm{QE},t}. \tag{36}$$

Hence, whenever $\sigma_{\mathrm{QE}} > 0$, QE shocks themselves may lead to additional macroeconomic volatility. Further, risk-averse arbitrageurs must hedge against QE risk in addition to fundamental sources of risk in the economy.

Figure 8*J* explores the change in long-run macroeconomic volatility as a function of the volatility of QE shocks. $\sigma_{\mathrm{QE}} = 0$ is our baseline estimate (marked by the vertical dotted line); the x-axis is the volatility of QE shocks relative to the volatility of habitat demand shocks $\sigma_{\mathrm{QE}}/\sigma_\beta$; the y-axis is the long-run (unconditional) variance of output $\mathrm{Var}[x_t]$. The results provide an important note of caution for central bankers: increased uncertainty regarding QE leads to increased macroeconomic volatility. In other words, although policy makers may desire the added flexibility of discretionary QE tools, the downside is that this increases the uncertainty surrounding these policy tools. By communicating clearly the expected path of QE purchases, policy makers should be able to reduce market uncertainty and prevent volatility spillovers from QE into the real economy. This provides additional support for the use of QE rules or forward guidance regarding asset purchases in the spirit of Greenwood, Hanson, and Vayanos (2016).[22]

V. Concluding Remarks

QE programs during the Great Recession were a massive policy experiment. The redeployment of QE in response to the COVID-19 crisis and the current reversal through QT is a testament to policy makers' belief in their effectiveness. But the precise channels through which QE policies work are still not clear, and so this paper seeks to unbundle QE by isolating one specific channel. Our focus is preferred habitat, which hypothesizes that there are investor clienteles with specific preferences for bonds within a given maturity segment. Utilizing Treasury auctions, we identify shifts in private demand for Treasuries that mimic QE but are independent of all other plausible channels of QE. We develop and confirm a key localization test of preferred habitat theory: when risk-bearing

[22] Additional robustness exercises are reported in figs. B30 and B31.

capacity is low, purchase or sale shocks of specific bonds in a given maturity segment have larger effects on bonds within this range. We then develop a general equilibrium preferred habitat model, using our demand shock results to discipline the model calibration. Our quantitative results are consistent with the view that QE programs worked largely through preferred habitat forces and that QE is a useful (but modest) policy tool for macroeconomic stabilization during crises. However, the effects of QE weaken if financial markets are well functioning, the holding period of assets purchases is short, or the program focuses on safer, shorter-duration assets. Further, uncertainty about future QE/QT rounds can lead to excess macroeconomic volatility. Given the design of QT as well as the prevailing financial environment, our model predicts rather small effects on output and inflation.

Our results suggest several promising directions for future work. First, the demand shocks we construct are a potentially important input to evaluate a host of quantity-based asset pricing theories. Second, while we have focused on preferred habitat in the market for Treasuries and other debt securities, this channel may be relevant for other arenas too. For example, we document suggestive evidence supporting state dependence of demand shock spillovers to MBS and possibly to equities. In a similar spirit, Greenwood et al. (2020) and Gourinchas, Ray, and Vayanos (2022) explore preferred habitat effects in currency markets and international bond markets. Finally, preferred habitat is one of many financial frictions relevant for economic agents. It will be interesting to explore how preferred habitat interacts with debt versus equity structure of firms, liquidity constraints, and banking frictions. Apart from clarifying the distributional aspects of QE-like policies, this extension could also shed more light on how preferred habitat can affect investment.

Data Availability

Code and data for replication can be found in Ray, Droste, and Gorodnichenko (2024) in the Harvard Dataverse, https://doi.org/10.7910/DVN/KHZI1L. This also includes information about the proprietary data used in this article.

Appendix A

Proofs

A1. Proof of Lemma 1

Collect all state and jump variables in a vector $\mathbf{Y}_t = [\mathbf{y}_t^\top \ \mathbf{x}_t^\top]^\top$. The interest rate process (15), risky asset payoff process (16), and habitat demand factor processes (14) are all affine functions of $\mathbf{Y}_t$. Moreover, the (linearized) Phillips curve (17) and IS equation (18) are also affine functions of $\mathbf{Y}_t$, since from equation (21) $\hat{r}_t = \hat{\mathbf{A}}^\top \mathbf{y}_t + \hat{C}$ is affine in the state variables. Aggregate dynamics can thus be written

$$\mathrm{d}\mathbf{Y}_t = -\Upsilon(\mathbf{Y}_t - \bar{\mathbf{Y}})\mathrm{d}t + \mathbf{S}\,\mathrm{d}\mathbf{B}_t. \tag{A1}$$

Note that Υ depends on $\hat{\mathbf{A}}$ but which we currently take as given. Then the rational expectations equilibrium is found immediately from Buiter (1984). Partition the eigenvalues and eigenvectors as follows:

$$\Upsilon = \mathbf{Q}\boldsymbol{\Lambda}\mathbf{Q}^{-1}, \boldsymbol{\Lambda} = \begin{bmatrix} \boldsymbol{\Lambda}_1 & 0 \\ 0 & \boldsymbol{\Lambda}_2 \end{bmatrix}, \mathbf{Q} = \begin{bmatrix} \mathbf{Q}_{11} & \mathbf{Q}_{12} \\ \mathbf{Q}_{21} & \mathbf{Q}_{22} \end{bmatrix},$$

where the partitions correspond to the state $\mathbf{y}_t$ and jump $\mathbf{x}_t$ variables. If the number of "stable" eigenvalues (nonnegative real parts) equals the number of state variables, then the rational expectations equilibrium dynamics are given by (22) and (23), where

$$\boldsymbol{\Gamma} = \mathbf{Q}_{11}\boldsymbol{\Lambda}_1\mathbf{Q}_{11}^{-1}, \boldsymbol{\Omega} = \mathbf{Q}_{21}\mathbf{Q}_{11}^{-1}. \tag{A2}$$

QED

A2. Proof of Lemma 2

Since asset prices are given by (24) and the state evolves according to (22), Ito's lemma implies that $\mathrm{d}P_t^{(\tau)}/P_t^{(\tau)} = \mu_t^{(\tau)}\mathrm{d}t + \boldsymbol{\sigma}^{(\tau)}\mathrm{d}\mathbf{B}_t$, with $\boldsymbol{\sigma}^{(\tau)} = -\mathbf{A}(\tau)^\top\boldsymbol{\sigma}$ and

$$\mu_t^{(\tau)} = \mathbf{A}'(\tau)^\top\mathbf{y}_t + C'(\tau) + \mathbf{A}(\tau)^\top\boldsymbol{\Gamma}(\mathbf{y}_t - \bar{\mathbf{y}}) + \frac{1}{2}\mathbf{A}(\tau)^\top\boldsymbol{\Sigma}\mathbf{A}(\tau), \tag{A3}$$

where $\boldsymbol{\Sigma} \equiv \boldsymbol{\sigma}\boldsymbol{\sigma}^\top$ and with analogous expressions for $\tilde{P}_t^{(\tau)}$. Differentiating the arbitrageur budget constraint (10) with respect to holdings $X_t^{(\tau)}$ gives the optimality conditions

$$\mu_t^{(\tau)} - i_t = a\left[\int_0^T X_t^{(\tau)}\mathbf{A}(\tau) + \tilde{X}_t^{(\tau)}\tilde{\mathbf{A}}(\tau)\mathrm{d}\tau\right]^\top \boldsymbol{\Sigma}\mathbf{A}(\tau),$$

again with analogous conditions with respect to $\tilde{X}_t^{(\tau)}$.

Substituting equation (24) into the habitat demand equations (13), we can write $Z_t^{(\tau)} = [\alpha(\tau)\mathbf{A}(\tau) - \boldsymbol{\Theta}(\tau)]^\top\mathbf{y}_t$ and $\tilde{Z}_t^{(\tau)} = [\tilde{\alpha}(\tau)\tilde{\mathbf{A}}(\tau) - \tilde{\boldsymbol{\Theta}}(\tau)]^\top\mathbf{y}_t$, where

$$\boldsymbol{\Theta}(\tau) = \begin{bmatrix} \cdots & \theta^k(\tau) & \cdots \end{bmatrix}^\top, \tilde{\boldsymbol{\Theta}}(\tau) = \begin{bmatrix} \cdots & \tilde{\theta}^k(\tau) & \cdots \end{bmatrix}^\top. \tag{A4}$$

Then substitute market clearing conditions $X_t^{(\tau)} = -Z_t^{(\tau)}$, $\tilde{X}_t^{(\tau)} = -\tilde{Z}_t^{(\tau)}$ into the optimality conditions and collect terms that are linear in the state $\mathbf{y}_t$ to get

$$\mathbf{A}'(\tau) + \mathbf{M}\mathbf{A}(\tau) - \mathbf{e}_i = 0, \tilde{\mathbf{A}}'(\tau) + \mathbf{M}\tilde{\mathbf{A}}(\tau) - \mathbf{e}_i = 0, \tag{A5}$$

where $\mathbf{M}$ is given by equation (26). If we take $\mathbf{M}$ as given, this is a linear system of differential equations. To derive initial conditions, note that at maturity, the riskless bonds pay \$1, while the risky asset pays D_t, so the $\tau = 0$ prices are given by $P_t^{(0)} = 1$, $\tilde{P}_t^{(0)} = D_t$. Hence, we have $(\mathbf{A}(0) = \mathbf{0}, \tilde{\mathbf{A}}(0) = -\mathbf{e}_d)$. Then if we assume that $\mathbf{M}$ is diagonalizable and invertible, the solution is given by equation (25). QED

A3. Proof of Proposition 1

In an affine equilibrium where asset prices are given by equation (24), we have that $\hat{r}_t = \int_0^T (\eta(\tau)\mu_t^{(\tau)} + \tilde{\eta}(\tau)\tilde{\mu}_t^{(\tau)})\mathrm{d}\tau$. Substituting equations (A3) and (A5) into this expression and collecting terms that are linear in the state $\mathbf{y}_t$ gives equation (27). Equilibrium is the solution of the fixed point problem implicitly defined by equations (26) and (27). Rewrite these conditions in the following function:

$$f(\hat{\mathbf{A}};\mathbf{M};a) = \begin{bmatrix} \mathbf{e}_i + \left(\boldsymbol{\Gamma}(\hat{\mathbf{A}})^\top - \mathbf{M}\right)\boldsymbol{\nu}(\mathbf{M}) - \hat{\mathbf{A}} \\ \mathrm{vec}\left\{\boldsymbol{\Gamma}(\hat{\mathbf{A}})^\top - a\cdot\boldsymbol{\Lambda}(\mathbf{M}) - \mathbf{M}\right\} \end{bmatrix}, \tag{A6}$$

where $\boldsymbol{\Lambda}(\mathbf{M})$ and $\boldsymbol{\nu}(\mathbf{M})$ are the integral terms from equations (26) and (27). In both cases, dependence on $\mathbf{M}$ comes through the affine coefficients $\mathbf{A}(\tau), \tilde{\mathbf{A}}(\tau)$. We have also made explicit the dependence of $\boldsymbol{\Gamma}$ on $\hat{\mathbf{A}}$, which can be seen in the proof of lemma 1. If $J \equiv \dim \mathbf{y}_t$, then $\dim \mathbf{M} = J \times J$ and $\dim \hat{\mathbf{A}} = J$ and the function $f:\mathbb{R}^{J(J+1)+1} \to \mathbb{R}^{J(J+1)}$. For any value of a, equilibrium is defined by $f(\hat{\mathbf{A}};\mathbf{M};a) = 0$.

We now analyze the solution in a neighborhood around $a = 0$. For $a = 0$, clearly $\hat{\mathbf{A}} = \mathbf{e}_i$ and $\mathbf{M} = \boldsymbol{\Gamma}(\mathbf{e}_i)^\top$. The partial derivatives evaluated at this point are given by

$$\frac{\partial f}{\partial \hat{\mathbf{A}}_j} = \begin{bmatrix} \left[\partial\boldsymbol{\Gamma}^\top/\partial\hat{\mathbf{A}}_j\right]\boldsymbol{\nu}(\mathbf{M}) - \mathbf{e}_j \\ \mathrm{vec}\left[\partial\boldsymbol{\Gamma}^\top/\partial\hat{\mathbf{A}}_j\right] \end{bmatrix}, \frac{\partial f}{\partial \mathbf{M}_{kl}} = \begin{bmatrix} \mathbf{e}_k\mathbf{e}_l^\top\boldsymbol{\nu}(\mathbf{M}) \\ -\mathrm{vec}\ \mathbf{e}_k\mathbf{e}_l^\top \end{bmatrix},$$

where $\mathbf{e}_j$, $\mathbf{e}_k$, $\mathbf{e}_l$ are standard normal basis vectors. The matrix $\partial\boldsymbol{\Gamma}/\partial\hat{\mathbf{A}}_j$ is the derivative of the state dynamics matrix $\boldsymbol{\Gamma}$ with respect to the j-element of $\hat{\mathbf{A}}$; because this depends on derivatives of the eigendecomposion defined in the proof of lemma 1, even in the case of $a = 0$ this is a complicated expression. Nevertheless, from this we can show that the Jacobian of f with respect to $\hat{\mathbf{A}}$, $\mathbf{M}$ evaluated at the $a = 0$ solution has full rank. In fact, writing this Jacobian in block form, we have

$$Df \equiv \begin{bmatrix} \mathbf{D}_{11} & \mathbf{D}_{12} \\ \mathbf{D}_{21} & -\mathbf{I}_{J^2} \end{bmatrix}, \tag{A7}$$

and $\mathbf{D}_{12} = [\mathbf{I}_J\cdot\nu_1 \ldots \mathbf{I}_J\cdot\nu_J]$, where ν_j is the j-element of $\boldsymbol{\nu}(\mathbf{M})$. Because the elementary row operations that transform $\mathbf{D}_{12}$ into the zero matrix simultaneously transform $\mathbf{D}_{11}$ into $-\mathbf{I}_J$, $\det Df = 1$ and the result follows from the implicit function theorem. QED

References

Beetsma, R., M. Giuliodori, F. de Jong, and D. Widijanto. 2016. “Price Effects of Sovereign Debt Auctions in the Euro-Zone: The Role of the Crisis.” *J. Financial Intermediation* 25:30–53.

Beetsma, R., M. Giuliodori, J. Hanson, and F. de Jong. 2018. "Bid-to-Cover and Yield Changes around Public Debt Auctions in the Euro Area." *J. Banking and Finance* 87:118–34.

Bernanke, B. S., and K. N. Kuttner. 2005. "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *J. Finance* 60 (3): 1221–57.

Bhattarai, S., and C. J. Neely. 2020. "An Analysis of the Literature on International Unconventional Monetary Policy." Working Paper no. 2016-21, Fed. Reserve Bank St. Louis.

Bretscher, L., L. Schmid, I. Sen, and V. Sharma. 2022. "Institutional Corporate Bond Pricing." Working Paper no. 21-07, Swiss Finance Inst.

Buiter, W. H. 1984. "Saddlepoint Problems in Continuous Time Rational Expectations Models: A General Method and Some Macroeconomic Examples." *Econometrica* 52 (3): 665–80.

Cahill, M. E., S. D'Amico, C. Li, and J. S. Sears. 2013. "Duration Risk versus Local Supply Channel in Treasury Yields: Evidence from the Federal Reserve's Asset Purchase Announcements." Working Paper no. 2013-35, Finance and Econ. Discussion Series.

Cammack, E. B. 1991. "Evidence on Bidding Strategies and the Information in Treasury Bill Auctions." *J.P.E.* 99 (1): 100–130.

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano. 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers Econ. Activity* 43 (1): 1–80.

Carlstrom, C. T., T. S. Fuerst, and M. Paustian. 2017. "Targeting Long Rates in a Model with Segmented Markets." *American Econ. J. Macroeconomics* 9 (1): 205–42.

Chen, H., V. Cúrdia, and A. Ferrero. 2012. "The Macroeconomic Effects of Large-Scale Asset Purchase Programmes." *Econ. J.* 122 (564): 289–315.

Chodorow-Reich, G. 2014. "Effects of Unconventional Monetary Policy on Financial Institutions." *Brookings Papers Econ. Activity* 45 (1): 155–227.

Cui, W., and V. Sterk. 2021. "Quantitative Easing with Heterogeneous Agents." *J. Monetary Econ.* 123:68–90.

Cúrdia, V., and M. Woodford. 2011. "The Central-Bank Balance Sheet as an Instrument of Monetary Policy." *J. Monetary Econ.* 58 (1): 54–79.

D'Amico, S., and T. B. King. 2013. "Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply." *J. Financial Econ.* 108 (2): 425–48.

Demos, Telis. 2011. "Weak Demands for 30-Year US Treasuries." *Financial Times*, August 11.

Di Maggio, M., A. Kermani, and C. J. Palmer. 2020. "How Quantitative Easing Works: Evidence on the Refinancing Channel." *Rev. Econ. Studies* 87 (3): 1498–528.

Ennis, H., and K. Kirk. 2022. "Projecting the Evolution of the Fed's Balance Sheet." Working Paper no. 22-15, Fed. Reserve Bank Richmond.

Federal Reserve Board. 2022. "Plans for Reducing the Size of the Federal Reserve's Balance Sheet." Washington, DC: Fed. Reserve System.

Fieldhouse, A. J., K. Mertens, and M. O. Ravn. 2018. "The Macroeconomic Effects of Government Asset Purchases: Evidence from Postwar U.S. Housing Credit Policy." *Q.J.E.* 133 (3): 1503–60.

Fleming, M. J. 2007. "Who Buys Treasury Securities at Auction?" *Current Issues Econ. and Finance* 13 (1).

Fleming, M. J., and W. Liu. 2016. "Intraday Pricing and Liquidity Effects of US Treasury Auctions." Working paper.

Forest, J. J. 2018. "The Effect of Treasury Auction Results on Interest Rates: The 1990s Experience." Working paper, Univ. Massachusetts Amherst.

Gali, J. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton, NJ: Princeton Univ. Press.
Garbade, K. D., and J. F. Ingber. 2005. "The Treasury Auction Process: Objectives, Structure, and Recent Acquisitions." *Current Issues Econ. and Finance* 11 (2).
Gertler, M., and P. Karadi. 2011. "A Model of Unconventional Monetary Policy." *J. Monetary Econ.* 58 (1): 17–34.
———. 2013. "QE 1 vs. 2 vs. 3. . .: A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool." *Internat. J. Central Banking* 9 (1): 5–53.
Gomes, F., M. Haliassos, and T. Ramadorai. 2021. "Household Finance." *J. Econ. Literature* 59 (3): 919–1000.
Gorodnichenko, Y., and W. Ray. 2017. "The Effects of Quantitative Easing: Taking a Cue from Treasury Auctions." Working Paper no. 24122, NBER, Cambridge, MA.
Gourinchas, P.-O., W. Ray, and D. Vayanos. 2022. "A Preferred-Habitat Model of Term Premia, Exchange Rates, and Monetary Policy Spillovers." Working Paper no. 29875, NBER, Cambridge, MA.
Greenlaw, D., J. D. Hamilton, E. Harris, and K. D. West. 2018. "A Skeptical View of the Impact of the Fed's Balance Sheet." Working Paper no. 24687, NBER, Cambridge, MA.
Greenwood, R., S. Hanson, J. Stein, and A. Sunderam. 2020. "A Quantity Driven Theory of Risk Premiums and Exchange Rates." Working Paper no. 27615, NBER, Cambridge, MA.
Greenwood, R., S. G. Hanson, and D. Vayanos. 2016. "Forward Guidance in the Yield Curve: Short Rates versus Bond Supply." In *Monetary Policy through Asset Markets: Lessons from Unconventional Measures and Implications for an Integrated World*, edited by E. Albagli, D. Saravia, and M. Woodford, 11–62. Santiago: Central Bank Chile.
Greenwood, R., and D. Vayanos. 2014. "Bond Supply and Excess Bond Returns." *Rev. Financial Studies* 27 (3): 663–713.
Gürkaynak, R. S., B. Sack, and J. H. Wright. 2007. "The U.S. Treasury Yield Curve: 1961 to the Present." *J. Monetary Econ.* 54 (8): 2291–304.
Hamilton, J. D., and J. C. Wu. 2012. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." *J. Money, Credit, and Banking* 44:3–46.
He, Z., B. Kelly, and A. Manela. 2016. "Intermediary Asset Pricing: New Evidence from Many Asset Classes." Working Paper no. 21920, NBER, Cambridge, MA.
He, Z., and A. Krishnamurthy. 2013. "Intermediary Asset Pricing." *A.E.R.* 103 (2): 732–70.
Ippolito, F., A. K. Ozdagli, and A. Perez-Orive. 2018. "The Transmission of Monetary Policy through Bank Lending: The Floating Rate Channel." *J. Monetary Econ.* 95:49–71.
Kaminska, I., and G. Zinna. 2020. "Official Demand for U.S. Debt: Implications for U.S. Real Rates." *J. Money, Credit, and Banking* 52 (2–3): 323–64.
Karadi, P., and A. Nakov. 2020. "Effectiveness and Addictiveness of Quantitative Easing." *J. Monetary Econ.* 117:1096–117.
King, T. 2019a. "Duration Effects in Macro-Finance Models of the Term Structure." Working paper.
———. 2019b. "Expectation and Duration at the Effective Lower Bound." *J. Financial Econ.* 134 (3): 736–60.
Koijen, R. S., F. Koulischer, B. Nguyen, and M. Yogo. 2021. "Inspecting the Mechanism of Quantitative Easing in the Euro Area." *J. Financial Econ.* 140 (1): 1–20.
Koijen, R. S. J., and M. Yogo. 2019. "A Demand System Approach to Asset Pricing." *J.P.E.* 127 (4): 1475–515.

———. 2022. "Understanding the Ownership Structure of Corporate Bonds." Working Paper no. 29679, NBER, Cambridge, MA.

Krishnamurthy, A., and A. Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers Econ. Activity* 2011 (2): 215–65.

———. 2012. "The Aggregate Demand for Treasury Debt." *J.P.E.* 120 (2): 233–67.

Kuttner, K. N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *J. Monetary Econ.* 47 (3): 523–44.

Kyle, A. S., and W. Xiong. 2001. "Contagion as a Wealth Effect." *J. Finance* 56 (4): 1401–40.

Li, C., and M. Wei. 2013. "Term Structure Modeling with Supply Factors and the Federal Reserve's Large-Scale Asset Purchase Programs." *Internat. J. Central Banking* 9 (1): 3–39.

Lou, D., H. Yan, and J. Zhang. 2013. "Anticipated and Repeated Shocks in Liquid Markets." *Rev. Financial Studies* 26 (8): 1891–912.

Mackenzie, Michael. 2010. "Strong Treasuries Sale Calms Nerves." *Financial Times*, December 9.

Martin, C., and C. Milas. 2012. "Quantitative Easing: A Sceptical Survey." *Oxford Rev. Econ. Policy* 28 (4): 750–64.

Modigliani, F., and R. Sutch. 1966. "Innovations in Interest Rate Policy." *A.E.R.* 56 (1/2): 178–97.

Peng, C., and C. Wang. 2022. "Factor Demand and Factor Returns." Working paper.

Ray, W. 2019. "Monetary Policy and the Limits to Arbitrage: Insights from a New Keynesian Preferred Habitat Model." Working paper.

Ray, W., M. Droste, and Y. Gorodnichenko. 2024. "Replication package for: 'Unbundling Quantitative Easing: Taking a Cue from Treasury Auctions.'" Harvard Dataverse, https://doi.org/10.7910/DVN/KHZI1L.

Rodnyansky, A., and O. M. Darmouni. 2017. "The Effects of Quantitative Easing on Bank Lending Behavior." *Rev. Financial Studies* 30 (11): 3858–87.

Romer, C. D., and D. H. Romer. 2017. "New Evidence on the Aftermath of Financial Crises in Advanced Countries." *A.E.R.* 107 (10): 3072–118.

Sims, E., and J. C. Wu. 2020. "Evaluating Central Banks' Tool Kit: Past, Present, and Future." *J. Monetary Econ.* 118:135–60.

Vayanos, D., and J. Vila. 2021. "A Preferred-Habitat Model of the Term Structure of Interest Rates." *Econometrica* 89 (1): 77–112.

Wright, J. H. 2012. "What Does Monetary Policy Do to Long-Term Interest Rates at the Zero Lower Bound?" *Econ. J.* 122 (564): F447–F466.

Zinman, J. 2015. "Household Debt: Facts, Puzzles, Theories, and Policies." *Ann. Rev. Econ.* 7:251–76.