The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXV, NO. 5 • OCTOBER 2020

The Banking View of Bond Risk Premia

VALENTIN HADDAD and DAVID SRAER*

ABSTRACT

Banks' balance sheet exposure to fluctuations in interest rates strongly forecasts excess Treasury bond returns. This result is consistent with optimal risk management, a banking counterpart to the household Euler equation. In equilibrium, the bond risk premium compensates banks for bearing fluctuations in interest rates. When banks' exposure to interest rate risk increases, the price of this risk simultaneously rises. We present a collection of empirical observations that support this view, but also discuss several challenges to this interpretation.

BANKS ARE LARGE SOPHISTICATED INTERMEDIARIES in the market for interest rate risk, but are absent from standard studies of the yield curve.[1] In this paper, we show that banks' balance sheet exposure to fluctuations in interest rates strongly forecasts excess Treasury bond returns. We interpret this result through the lens of banks' risk management decisions, which tightly connect their exposure to interest rate risk with the price of this risk. This connection represents a banking counterpart to the classic household Euler equation. In equilibrium, an increase in future bond returns compensates any increase in banks' exposure to interest rate risk.[2] This paper establishes this relationship

*Valentin Haddad is with UCLA and NBER. David Sraer is with UC Berkeley, NBER, and CEPR.We gratefully acknowledge the useful comments and suggestions of Stefan Nagel; an Associate Editor; two anonymous referees; Tobias Adrian; Mikhail Chernov; Anna Cieslak; John Cochrane; Arvind Krishnamurthy; Augustin Landier; Giorgia Piacentino; Monika Piazzesi; David Thesmar; Dimitri Vayanos; as well as seminar participants at Kellogg, Princeton, the University of Michigan, Stanford, the Federal Reserve Bank of New York, the UNC Junior Faculty Roundtable, the NBER Summer Institute, CITE, the MFS meeting, the Adam Smith Conference, and the SED Annual Meeting. We have read *The Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

Correspondence: Valentin Haddad, Anderson School of Management, UCLA, 110 Westwood Plaza, Los Angeles, CA 90095; e-mail: valentin.haddad@anderson.ucla.edu.

[1] In 2014, private depository institutions (U.S.-chartered depository institutions, foreign banking offices, banks in U.S.-affiliated areas, and credit unions) held 3.2% of all outstanding Treasuries, 25% of agency and government-sponsored enterprise-backed securities, 12.3% of municipal securities, 33.6% of mortgages, and 49.5% of all consumer credit.

[2] Importantly, this statement describes an equilibrium relation rather than a causal relationship. The price and quantity of interest rate risk are jointly determined in equilibrium. However, we sometimes follow the tradition of the literature on the household Euler equation, which tends to describe equilibrium relations using a more causal language.

DOI: 10.1111/jofi.12949

empirically, presents a collection of facts that further support this view, and highlights challenges to this interpretation.

We start by constructing a measure of the average bank exposure to interest rate risk. At the bank level, we follow Gomez et al. (Forthcoming) and use the *income gap* as our measure of interest rate risk exposure. The income gap of a financial institution corresponds to the difference between the book value of all assets that either reprice or mature within one year and the book value of all liabilities that mature or reprice within a year, normalized by total assets. This measure, commonly used by both banks and bank regulators, is readily available at the quarterly frequency for the 1986 to 2014 period through FR Y-9C filings of Bank Holding Corporations (BHC) to the Federal Reserve. The income gap provides a relevant quantification of the *net* exposure of banks' income to interest rate risk. Gomez et al. (Forthcoming) show that the sensitivity of banks' profits to interest rates increases significantly with their income gap.[3] We use the average income gap across banks with more than \$1bn of total assets as our measure of financial intermediaries' interest rate risk exposure.

We run regressions of one-year excess returns on Treasuries—borrow at the short rate, buy a long-term bond—on the average income gap available at the beginning of the period. The estimated coefficient is significant for all bond maturities. With this single predictor, we find R^2 values of 20% on average across maturities. A battery of robustness checks shows that this result does not spuriously derive from the persistence of our forecasting variable in a small sample. Additionally, the forecasting power of the average income gap for Treasuries' excess returns is not affected by the inclusion of macroeconomic factors known to predict bond returns (Ludvigson and Ng (2009)). The robust correlation between bonds' excess returns and the average income gap, depicted in Figure 1, is the main contribution of the paper. This finding offers *prima facie* evidence of the role of financial intermediaries in asset pricing (e.g., He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014)).

We interpret this finding through the lens of a simple equilibrium restriction on the yield curve following Greenwood and Vayanos (2014). This equilibrium restriction must hold in a large family of economies. In the model, banks trade assets of different maturities to maximize their expected profits while managing their risk. When banks hold more long-term assets, they must absorb additional interest rate risk. They will do so only if the market compensation for this risk increases. Such compensation can be observed in, for instance, Treasury bond returns.[4] In equilibrium, banks' income gap, that is, the sensitivity of banks' profits to variation in the short rate, is negatively correlated

[3] Purnanandam (2007), Begenau, Piazzesi, and Schneider (2015), and English, den Heuvel, and Zakrajsek (2012) also document that financial intermediaries do not fully hedge their exposure to interest rate risk. Di Tella and Kurlat (2017) build a model to explain why banks optimally expose their balance sheets to movements in interest rates.

[4] While a large share of the exposure of banks to interest rate risk comes from non-Treasury assets, Treasuries constitute a simple and stable way to measure this price of risk. Hanson (2014) and Malkhozov et al. (2016) follow a similar measurement approach in the context of mortgage-backed security supply.

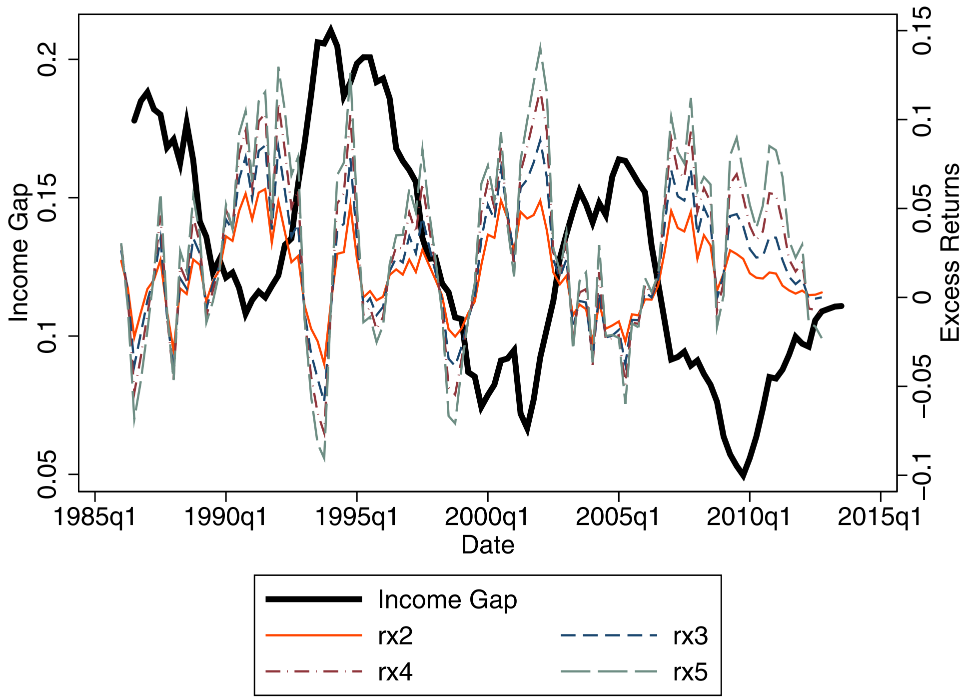

Figure 1. Average income gap and future bond excess returns. This figure plots the time series of banks' income gap and future bond excess returns. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year and the dollar amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. "Income Gap" is the average income gap, computed across all U.S. bank holding companies with total consolidated assets of $1 bn or more. rx n is the excess one-year return of zero-coupon bonds of maturity n, using data from Gurkaynak, Sack, and Wright (2007) (Color figure can be viewed at wileyonlinelibrary.com)

with bond risk premia. Since long-term Treasuries are more sensitive to the interest rate than short-term Treasuries, this correlation between banks' income gap and risk premia is larger, in absolute value, for bonds of longer maturities. These qualitative predictions echo our main findings. We confirm that they hold quantitatively as well. Fitting the model to the data also allows us to estimate banks' willingness to take risk, a key input for our theory and more generally for macroeconomic models with financial intermediation.

Our analysis departs from the classic, frictionless view of the market for interest rate risk. This view has received mitigated empirical success so far.[5] In contrast, several recent papers provide convincing evidence that not all

[5] See Duffee (2018), Gürkaynak and Wright (2012), Beeler and Campbell (2012), and Schneider (2017) for discussions of these issues.

investors are marginal in Treasury markets.[6] In such a setting, understanding the investment decisions of marginal investors is key to the determination of asset prices. Banks are natural candidates for this role. They hold a sizable share of assets exposed to interest rates. Their modest holdings of Treasuries understate their prominence in the broader fixed income markets (mortgages, consumer credit, and agency-backed securities). Banks are also likely sophisticated in managing their interest rate risk exposure (e.g., Drechsler, Savov, and Schnabl (2018)). The tight empirical relationship between banks' balance sheet exposure and bond excess returns supports the view that banks are marginal investors in Treasury markets. The remainder of the paper tests this hypothesis further.

We present a collection of evidence consistent with this banking view of bond risk premia. First, we show that, by itself, the average exposure of banks' *assets* to interest rate risk does not forecast bond risk premia in a significant way. The same holds for the average exposure of banks' *liabilities* to interest rate risk. Only financial institutions' overall holding of interest rate risk, that is, the average income gap, significantly predicts future bond excess returns. This finding is consistent with the interpretation that bond risk premia only appear in banks' *overall* portfolio holdings. Second, we show that, over our sample period, standard measures of liquidity risk do not forecast bond risk premia, in contrast to our measure of interest rate risk exposure. Third, we show that in the time series, the average income gap responds to several measured changes in the supply and demand for interest rate risk in the economy, such as the total amount of fixed-rate mortgages net of adjustable-rate mortgages (ARMs), the total supply of Treasuries, or the amount of noninterest-bearing deposits. However, these shocks to the supply and demand for interest rate risk have no forecasting power for bond risk premia above and beyond the income gap. This result is again consistent with our interpretation since bond risk premia should be captured entirely by banks' net position, measured in our analysis by the average income gap, and not by any particular components of their net position. Finally, we exploit our bank-level data to provide evidence consistent with interest rate risk-sharing among heterogeneous banks. We split our sample of banks into 10 size-sorted groups and compute the time series of the average income gap for these 10 groups. Despite their heterogeneity, we show that these 10 groups share a similar evolution of their average income gap over time. We find similar evidence of risk-sharing among banks with different leverage or among banks located in different geographic areas. All of these results are consistent with our simple theory.

However, our preferred interpretation faces several challenges. In our theory, banks suffer when they hold significant balance sheet exposures and interest rates increase. This assumption underlies banks' risk management motive and

[6] For example, Greenwood and Vayanos (2014) provides such evidence at a low frequency, while Gagnon et al. (2011), Krishnamurthy and Vissing-Jorgensen (2011), Swanson (2011), Hamilton and Wu (2012), and D'Amico and King (2013) document such effects around quantitative easing interventions.

drives the relation between banks' average income gap and excess returns on Treasury bonds. Using banks' equity returns, we fail to find empirical support for this assumption. In the data, periods of low-income gaps are not positively related to the correlation between banks' equity returns and bond returns. Related, our mean-variance framework implies that bond risk premia should be proportional to the expected covariance of banks' equity returns with bond returns. Yet, the data show no significant relationship between bond excess returns and the predicted covariance between daily excess returns on long-term bonds and banks' stock returns. Finally, our model predicts that banks' balance sheet exposure should command a higher risk premium in periods of high interest rate risk. Using the realized variance of bond returns as a source of variation in interest rate risk beyond changes in balance sheet composition, we find no support for this prediction. All of these results challenge the interpretation that potential valuation losses drive the reluctance of financial institutions to bear risk, a standard feature of intermediary asset pricing models (e.g., Brunnermeier and Sannikov (2014), He and Krishnamurthy (2014)).

Related Literature. Our paper relates to the literature that seeks to understand the pricing of interest rate risk. One strand of the literature investigates how the price of interest rate risk relates to the information contained in the yield curve.[7] Another strand of this literature explores the role of macroeconomic variables in explaining excess returns on Treasuries.[8] Finally, a third strand of this literature emphasizes the role of segmentation in Treasury markets and shows that supply factors forecast bond risk premia.[9] Relative to this literature, our contribution is to shift the focus to financial institutions, which are major participants in the market for interest rate risk, and to use information on financial institutions' exposure to interest rate risk to forecast future bond returns.

In doing so, our paper also relates to the recent literature that emphasizes the crucial role of intermediaries for asset prices. Several theoretical contributions emphasize the role of intermediaries' balance sheets for equilibrium risk premia.[10] Empirically, the importance of financial intermediaries for the determination of asset prices has been investigated mostly in the context of equity markets (e.g., Adrian, Etula, and Muir (2014), Adrian, Moench, and Shin (2016) and He, Kelly, and Manela (2017)). Relative to this literature, our contribution shifts the focus away from equity markets to the market for Treasuries. Furthermore, our approach uses intermediaries' actual underlying risk exposure as a forecasting variable, instead of focusing on leverage as a proxy

[7] See, for example, Fama and Bliss (1987), Campbell and Shiller (1991), Cochrane and Piazzesi (2005), Duffee (2011), Adrian, Crump, and Moench (2013), and Cieslak and Povala (2015).

[8] See, for example, Piazzesi (2005), Ang and Piazzesi (2003), Ludvigson and Ng (2009), and Cooper and Priestley (2009).

[9] See, for example, Greenwood and Vayanos (2014), Gagnon et al. (2011), Krishnamurthy and Vissing-Jorgensen (2011), Swanson (2011), Hamilton and Wu (2012), D'Amico and King (2013), Baker, Greenwood, and Wurgler (2003), Hanson (2014), and Malkhozov et al. (2016).

[10] Prominent papers include, among others, Brunnermeier and Pedersen (2009), He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014).

for this exposure. Finally, while our paper highlights several empirical facts consistent with an intermediary asset pricing interpretation, we also present several challenges to this interpretation.

The rest of the paper is organized as follows. Section I describes the data that we use in our empirical analysis and discusses our main empirical results. Section II presents the model underlying our interpretation of this evidence. Section III offers a structural estimation of our model to quantify banks' risk management motive. Section IV provides further tests consistent with this banking view of bond risk premia. Section V concludes.

I. Banks' Income Gap and Bond Returns

A. Data

A.1. Income Gap

Income gap definition. The central object of our analysis is banks' net exposure to interest rate risk. Our main empirical counterpart to this quantity is the income gap, a standard measure of interest rate sensitivity used by banks and regulators. Our definition of the income gap follows the definition in Mishkin and Eakins (2009),

$$\textit{Income Gap} = (\textit{RSA} - \textit{RSL})/\textit{Total Assets}, \tag{1}$$

where *RSA* is a measure of the dollar amount of assets that either reprice or mature within one year and *RSL* is a measure of the dollar amount of liabilities that mature or reprice within a year. A high income gap therefore corresponds to low exposure to long-term fixed-rate assets. More specifically, we construct the income gap using variables from schedule HC-H of form FR Y-9C, which focuses on the interest sensitivity of the balance sheet. The variable *RSA* is provided directly (item bhck3197). The variable *RSL* has four elements: long-term debt that reprices within one year (item bhck3298), long-term debt that matures within one year (bhck3409), variable-rate preferred stock (bhck3408), and interest-bearing deposit liabilities that reprice or mature within one year (bhck3296), such as certificates of deposits. Empirically, the latter is by far the most important determinant of liability-side sensitivity to interest rates. All of these items are available quarterly from 1986 to 2014. We scale these variables by total assets and report summary statistics in Internet Appendix Table IA.I.[11] On average, the variable *RSL* (interest rate-sensitive liabilities) consists mostly of variable-rate deposits that mature or reprice within a year. Long-term debt typically has a fixed rate. Gomez et al. (Forthcoming) validate this measure in the cross-section of banks: they document that when the federal funds rate rises, banks with a larger income gap generate stronger earnings and contract their lending less than other banks.

[11] The Internet Appendix is available in the online version of the article on *The Journal of Finance* website.

Our primary forecasting variable for bond risk premia is the average income gap, which we compute across all banks with more than $1bn in consolidated assets. This variable is available quarterly from 1986 to 2014. Figure 1 shows the time-series evolution of the average income gap over this period (thick dark line). The average income gap exhibits procyclical variation. The income gap peaks during expansions, and banks accumulate interest rate risk—lower values of the gap—ahead of recessions. We favor this simple variable for most of our analysis because (i) it captures the forces of our theory, (ii) it has a transparent construction, and (iii) it reflects how market participants measure interest rate sensitivity in practice. The remainder of this section discusses benefits and limitations of this measure.

Measurement issues. A first dimension is the treatment of deposits. In the BHC data, the item corresponding to short-term deposit liabilities (bck3296) does not include transaction or savings deposits.[12] Interest rates on these "core" deposits, while having zero contractual maturity, are known to adjust sluggishly to changes in short-term market rates (Hannan and Berger (1991), Neumark and Sharpe (1992)). Therefore, despite their short maturity, it is natural to exclude these deposits from our measure, as they do not induce direct cash-flow changes when interest rates change. However, if these core deposits adjust slightly to changes in the federal funds rate, the average income gap will overestimate the real income gap.[13] To investigate the role of deposits, we follow English, den Heuvel, and Zakrajsek (2012), and alternatively, assume that all noninterest-bearing deposits have short maturity. This change results in a lower mean for the average income gap: 0% versus 12% in our baseline. However, this modified income gap exhibits a correlation of 91% with our baseline measure.

A second dimension is that we do not observe holdings of interest rate derivatives. If banks hedge their interest rate risk exposure through derivatives, the income gap may overestimate banks' exposure to interest rate risk. To address this concern, we exploit the fact that, since 1995, banks report on form FR Y-9C the notional amounts of the interest derivatives they contract.[14] We compute the average income gap for banks that never report any notional amounts of interest rate derivatives and report its time-series evolution in Internet Appendix Figure IA.1 (dark dashed line). The time-series correlation of this series with the time series for the average income gap computed across all banks is 93%.

The lack of data on interest rate derivatives also explains why we do not use the *aggregate* income gap, that is, the *asset-weighted* average income gap,

[12] See http://www.federalreserve.gov/apps/mdrm/data-dictionary.

[13] More recently, Drechsler, Savov, and Schnabl (2016) investigate the price and quantity response of deposits to changes in the federal funds rate and find a somewhat larger elasticity of deposits to interest rates.

[14] Banks report seven types of derivative contracts: futures (bhck8693), forwards (bhck8697), written options that are exchange-traded (bhck8701), purchased options that are exchange-traded (bhck8705), written options traded over the counter (bhck8709), purchased options traded over the counter (bhck8713), and swaps (bhck3450).

as our main forecasting variable. Indeed, since large banks hold significant positions in interest rate derivatives, their income gap likely suffers from substantial measurement error. Given the fat-tailed distribution of banks' assets, this bank-level measurement error would translate into significant aggregate measurement error. Internet Appendix Figure IA.1 plots the time-series evolution of the asset-weighted average income gap (orange line), as well as the average income gap computed across the 10 largest banks (blue line). These two series are almost identical—the top 10 banks are so large that they account for most of the variation in the asset-weighted average gap. Any mismeasurement in the gap for some of these 10 banks would significantly garble our forecasting variable.

Despite these limitations, our income gap measure represents a significant contribution to the intermediary asset pricing literature. In this literature, financial intermediaries' risk exposures are typically summarized by their leverage (Adrian, Etula, and Muir (2014), Adrian, Moench, and Shin (2016), He, Kelly, and Manela (2017)). This approach fails to account for the differential exposure of different assets and liabilities to aggregate sources of risk. In contrast, using banks' average income gap allows for some risk-weighting of assets and liabilities.

Income gap and exposure g_t. In the model that we develop in Section II, we show that the relevant measure of banks' exposure, g_t, can be constructed from our basic income gap measure as

$$g_t = 1 - Income\ Gap_t \times \frac{A_t}{E_t}, \tag{2}$$

where E_t is equity value at date t. There are several reasons to favor the standard income gap measure over g_t in our empirical analysis. First, it lies between -1 and 1, it is defined for banks with negative equity, and its distribution has fewer outliers. Second, this measure corresponds to that used in Gomez et al. (Forthcoming), who show that the income gap forecasts banks' net income reaction to changes in interest rates. Importantly, $-g_t$ has a correlation of 94% with the baseline average income gap. Internet Appendix Figure IA.2 depicts the four versions of the income gap (including or excluding deposits, scaling by assets or equity). We standardize the measures, so they have zero mean and unit standard deviation. When considering the quantitative properties of the model in Section III, we define g_t using equation (2) and include deposits. Thus constructed, g_t has a mean of 1.00 and a standard deviation of 0.41. Banks typically have positive exposure to long-term assets, on average equal to their equity and varying roughly between zero and twice their equity. This exposure constitutes a sizable amount of risk, but much less than under a naive approach that assumes that all of their assets are long term and all of their liabilities are short term.[15]

[15] Echoing the typical bank leverage, this would give rise to an interest risk exposure of around 10, an order of magnitude larger than we observe.

A.2. Bond Prices and Other Time-Series Variables

We are interested in relating banks' exposure to interest rate risk with the price of this risk. A natural way to measure this price is to consider Treasury bond risk premia. Bond return data are constructed from the Gurkaynak, Sack, and Wright (2007) data set of interpolated yield curves. These curves are computed by fitting Treasury transaction prices daily using Svensson (1994)'s extension of Nelson and Siegel (1987). We compute time series of bond prices with maturity of n years, $P_t^{(n)}$, where the yield of these bonds is given by $y_t^{(n)} = -\frac{1}{n}\ln(P_t^{(n)})$. The log-forward rate at time t for contracts between time $t+n-1$ and $t+n$ is $f_t^{(n)} = \ln(P_t^{(n-1)}) - \ln(P_t^{(n)})$. The log holding-period return from buying an n-year bond at time t and selling it a quarter later as an $n-1/4$ year bond is $r_{t\to t+1}^{(n)} = \ln(P_{t+1}^{(n-1/4)}) - \ln(P_t^{(n)})$. Quarterly bond excess returns are then defined as $rx_{t\to t+1}^{(n)} = r_{t\to t+1}^{(n)} - y_t^{(1/4)}$. Our analysis focuses on a one-year return horizon and maturities from two to five years, $rx_{t\to t+4}^{(n)} = \sum_{i=0}^{3} rx_{t+i\to t+i+1}^{(n)}$.

We also use several macroeconomic variables known to forecast bond risk premia. The output gap is the difference between real seasonally adjusted GDP (GDPC96 from the FRED database) and the real potential GDP (GDPPOT from FRED), normalized by real seasonally adjusted GDP (Cooper and Priestley (2009)). Industrial production growth is the one-year growth rate in industrial production (INDPRO in FRED). Inflation is the one-year growth rate of the CPI, taken from the FRED database. Table I presents descriptive statistics for these variables.

B. Income Gap and Excess Bond Returns

B.1. Main Results

We estimate the following linear equation using quarterly data:

$$rx_{t\to t+4}^{(n)} = a^{(n)} + b^{(n)} \times \mathit{Income\ Gap}_t + \epsilon_{t+4}^{(n)}, \quad \text{for } n = 2, 3, 4, \mathit{and}\ 5, \tag{3}$$

where $rx_{t\to t+4}^{(n)}$ is the excess return of a zero-coupon bond of maturity n from quarter t to quarter $t+4$, defined in Section I.A.2, and $\mathit{Income\ Gap}_t$ is the average income gap available at the beginning of quarter t, which corresponds to the average income gap of quarter $t-2$. To account for the overlapping nature of our return variable, we use the reverse regression approach of Hodrick (1992) to compute standard errors for our coefficient estimates. Additionally, we account for potential small-sample bias, such as the Stambaugh (1999) bias, by computing p-values from a parametric bootstrap procedure. More precisely, we first estimate a restricted VAR for quarterly excess returns and the income gap under the null of no return predictability by the income gap.[16] We assume that the joint distribution of innovations in the VAR corresponds to

[16] When we add additional controls to the regression, as in Tables III and V, we allow these other variables to predict returns in the VAR estimation.

Table I

Descriptive Statistics

This table uses quarterly data over the 1986 to 2014 period. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year and the dollar amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. Income Gap is the average bank-level income gap. $y^{(n)}$ is the yield of Gurkaynak, Sack, and Wright (2007) (GSW) zero-coupon bonds of maturity n. $rx^{(n)}$ is the excess one-year return of GSW zero-coupon bonds of maturity n. IP growth is the one-year growth rate in industrial production (INDPRO in FRED). Inflation is the one-year growth rate of the CPI, taken from the FRED database. Output gap corresponds to the difference between real seasonally adjusted GDP (GDPC96 from the FRED database) and real potential GDP (GDPPOT from FRED), normalized by real seasonally adjusted GDP.

Variable	Obs	Mean	*SD*	P25	P50	P75
Income Gap	109	0.128	0.042	0.092	0.122	0.163
$y^{(1)}$	111	0.041	0.026	0.015	0.046	0.059
$y^{(2)}$	111	0.043	0.025	0.018	0.047	0.063
$y^{(3)}$	111	0.046	0.025	0.023	0.048	0.063
$y^{(4)}$	111	0.048	0.024	0.027	0.05	0.066
$y^{(5)}$	111	0.05	0.023	0.03	0.051	0.069
$rx^{(2)}$	108	0.014	0.022	−0.001	0.013	0.027
$rx^{(3)}$	108	0.02	0.033	−0.002	0.021	0.042
$rx^{(4)}$	108	0.026	0.043	−0.008	0.026	0.057
$rx^{(5)}$	108	0.03	0.052	−0.012	0.031	0.069
IP Growth	107	0.021	0.042	0.012	0.029	0.046
Inflation	111	0.028	0.013	0.02	0.028	0.036
Output Gap	111	−0.059	0.02	−0.07	−0.055	−0.047

their empirical distribution. We then draw 5,000 samples from this estimated process to obtain a distribution of reverse regression t-statistics. We report the p-value of our estimated t-statistic relative to this bootstrapped distribution. Both the asymptotic standard error and the p-value are informative: the asymptotic standard error is robust to the specifics of the data-generating process, while the p-value handles finite-sample issues conditional on a parameterized data-generating process.[17]

The estimation of equation (3) is presented in Table II. The average income gap significantly predicts future bond excess returns. For bonds with a two-year maturity, $b^{(2)}$ is equal to -0.23 and is statistically significant with a p-value of 2.3%. This effect is economically significant. A one-standard-deviation increase in the average income gap is associated with about 97 basis points smaller future excess returns of two-year maturity zero-coupon bonds, which represents 44% of the volatility of these bond returns. A one-standard-deviation increase in the average income gap represents a 4.2 percentage point increase in the fraction of net short-term or variable-rate assets, which, given

[17] Internet Appendix Table IA.II reports estimates of equation (3) using Newey-West standard errors allowing for eight-quarter lags. However, this procedure has been found to overreject the null hypothesis in small samples (see, for example, Ang and Bekaert (2006)).

Table II

Income Gap and Bond Excess Returns

This table presents regressions of bond excess returns on banks' income gap. The sample period is 1986 to 2014. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year, and the dollar amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. Income Gap is the average bank-level income gap. $rx^{(n)}$ is the excess one-year return of GSW zero-coupon bonds of maturity n. Standard errors are computed using the reverse regression approach of Hodrick (1992). *, **, and *** indicate statistically different from zero at the 10%, 5%, and 1% level of significance, where p-values are computed using the bootstrap approach described in Section I.B.1.

	(1) $rx^{(2)}$	(2) $rx^{(3)}$	(3) $rx^{(4)}$	(4) $rx^{(5)}$
Income Gap	-0.23^{**}	-0.36^{**}	-0.47^{**}	-0.55^{**}
	(0.10)	(0.15)	(0.21)	(0.27)
Constant	0.04^{**}	0.07^{**}	0.09^{**}	0.10^{**}
	(0.01)	(0.02)	(0.03)	(0.04)
Bootstrapped p-value	0.023	0.024	0.035	0.048
Observations	106	106	106	106
Adjusted R^2	0.172	0.200	0.201	0.189

an average income gap of 12.8%, corresponds to a 32% increase in the average bank's exposure to interest rate risk.

This correlation increases almost linearly with the maturity of the bond. For bonds with a five-year maturity, $b^{(5)}$ is equal to -0.55, so that a one-standard-deviation increase in the average income gap corresponds to a 231 basis point reduction in five-year bond excess returns. This decrease represents about 44% of the volatility of these bonds. The parameters $b^{(3)}$, $b^{(4)}$, and $b^{(5)}$ are all statistically different from zero at the 5% confidence level, while $b^{(5)}$ is statistically different from $b^{(2)}$ at the 1% confidence level. The adjusted R^2s that we obtain from these forecasting regressions with a single forecasting variable are high, ranging from 17% using two-year maturity bonds up to 20% for bonds with a longer maturity.

Figure 1 highlights the strong forecasting power of the average income gap for future bond returns. This figure plots the value of the average income gap available in quarter t and the excess bond returns from quarter t to quarter $t+4$, $rx^{(n)}_{t\to t+4}$, for zero-coupon bonds of maturity n. The figure displays a striking and robust negative correlation between the income gap series and the excess return series throughout the sample period. In summary, we find that (i) a smaller average income gap predicts larger bond risk premia and (ii) this effect is stronger for long-maturity bonds.

In Table III, we augment equation (3) by including macroeconomic variables that forecast bond risk premia: the inflation rate, the growth in industrial production between $t-4$ and t, and the current output gap. Table III shows that the effect of the average income gap on future bond excess returns is unaffected by the inclusion of these variables. The estimated $b^{(n)}$ and the predictive

Table III

Income Gap and Bond Excess Returns: Controlling for Macroeconomic Conditions

This table presents regressions of bond excess returns on banks' income gap. The sample period is 1986 to 2014. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year, and the dollar amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. Income Gap is the average bank-level income gap. $rx^{(n)}$ is the excess one-year return of GSW zero-coupon bonds of maturity n. IP growth is the one-year growth rate in industrial production (INDPRO in FRED). Inflation is the one-year growth rate in the CPI, taken from the FRED database. Output gap corresponds to the difference between real seasonally adjusted GDP (GDPC96 from the FRED database) and real potential GDP (GDPPOT from FRED), normalized by real seasonally adjusted GDP. NBER recession is a dummy equal to 1 for quarters flagged as a recession by NBER. Standard errors are computed using the reverse regression approach of Hodrick (1992). *, **, and *** indicate statistically different from zero at the 10%, 5%, and 1% level of significance, where p-values are computed using the bootstrap approach described in Section I.B.1.

	(1) $rx^{(2)}$	(2) $rx^{(3)}$	(3) $rx^{(4)}$	(4) $rx^{(5)}$
Income Gap	-0.21^{*}	-0.34^{*}	-0.44^{*}	-0.52^{*}
	(0.11)	(0.17)	(0.23)	(0.28)
Inflation	0.49	0.57	0.57	0.54
	(0.33)	(0.52)	(0.70)	(0.86)
IP Growth	−0.12	−0.12	−0.09	−0.04
	(0.07)	(0.13)	(0.20)	(0.26)
Output Gap	0.11	0.06	−0.03	−0.13
	(0.18)	(0.30)	(0.42)	(0.54)
Constant	0.04	0.05	0.07	0.08
	(0.02)	(0.04)	(0.05)	(0.06)
Bootstrapped p-value	0.063	0.060	0.061	0.079
Observations	104	104	104	104
Adjusted R^2	0.228	0.210	0.187	0.162

power of the regressions are similar to those estimated in Table II, albeit less strongly statistically significant.

B.2. Further Analysis

Longer maturities. In Internet Appendix Figure IA.3, we estimate equation (3) for bonds of longer maturities.[18] Panel A of Internet Appendix Figure IA.3 plots the coefficients $b^{(n)}$, for $n = 1 \ldots 10$, as well as their 95% confidence intervals. The coefficients $b^{(n)}$ decrease for maturities from 2 to 10 years until reaching a level of about −0.6. For the longest maturities, the estimates become more imprecise. Panel B of Internet Appendix Figure IA.3 plots the corresponding adjusted R^2 for each of these regressions. The forecasting power of

[18] The original data cover the range of maturities regularly until 10 years, but are more sparse and thus make estimates less reliable thereafter.

Table IV

Income Gap and Bond Excess Returns: In Real Time

This table presents regressions of bond excess returns on banks' income gap. The sample period is 1986 to 2014. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year, and the dollar amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. Income Gap is the average bank-level income gap. Predicted return is computed using the current value of the income gap and the coefficients from a regression of realized excess returns on the income gap using all data available until that point. $rx^{(n)}$ is the excess one-year return of GSW zero-coupon bonds of maturity n. Standard errors are computed using the reverse regression approach of Hodrick (1992). *, **, and *** indicate statistically different from zero at the 10%, 5%, and 1% level of significance, where p-values are computed using the bootstrap approach described in Section I.B.1.

	(1) $rx^{(2)}$	(2) $rx^{(3)}$	(3) $rx^{(4)}$	(4) $rx^{(5)}$
Predicted Excess Return	0.81*	0.72**	0.69*	0.71
	(0.41)	(0.36)	(0.38)	(0.43)
Constant	0.00	0.00	0.00	0.00
	(0.01)	(0.01)	(0.01)	(0.02)
Bootstrapped p-value	0.052	0.050	0.076	0.112
Observations	71	71	71	71
Adjusted R^2	0.083	0.112	0.106	0.097

the income gap is highest for bonds of three- to five-year maturity, after which it decreases with maturity.

Across horizons. We investigate the predictive power of the income gap at various horizons. While one year is the standard horizon considered in the literature predicting bond returns, banks might make risk management decisions for different horizons. We confirm that our results are robust across horizons. In Internet Appendix Table IA.III, we replicate our baseline regression at the one-quarter horizon. The estimated coefficients are about a quarter of the annual estimates and therefore of similar economic significance. The p-values range from 2.1% to 7.6%. We also construct one-month returns using the Fama constant maturity portfolios obtained from CRSP. These portfolios are formed each month from bonds of maturity ranging in a one-year interval. Internet Appendix Table IA.IV reports these estimates. The results are again consistent with our baseline.

Real-time prediction. The predictive power found in the full sample may not be observable to economic agents in real time. To examine whether this is a concern for our analysis, we construct a real-time version of our predictor. At each date t, we estimate a regression of bond excess returns using all data available up to that point. We use the estimated coefficients of this regression together with the gap at date t to construct a real-time predictor of returns between t and $t + 1$. We start the estimation after eight years of data are available. Table IV reports results of regressions of bond excess returns on this real-time predictor. As the sample period grows large, the coefficient estimate should approach one, whereas in a finite sample, the limited amount of data should generate measurement error and bias the estimate toward zero. However, despite the short sample period used in our case, we report coefficient estimates that are away from zero, ranging from 0.69 for four-year bonds to 0.81 for two-year bonds. The coefficients for maturities of two and four years are significant at the 10% level and those for three-year bonds are significant at the 5% level. The adjusted R^2s range from 8.3% to 11.2%. Thus, while more moderate than the full-sample estimates, these results indicate that the income gap has significant predictive power in real time.

B.3. Relation with Yield-Based Predictors

We now turn to an alternative, more indirect, approach to studying how much of the variation in bond risk premia is captured by the income gap. Of course, we can never fully characterize these expected returns because the set of potential predictors is arbitrarily large. However, in a large family of models—including the one that we present in Section II—spanning holds for most parameter combinations: yields at date t capture all of the information necessary to predict bond excess returns. Predictability by yields therefore is a useful benchmark to consider.

We first examine whether the income gap captures *additional* information about bond risk premia relative to the yield curve. We augment equation (3) by including three and five principal components (PCs) of yields with maturity from 1 to 10 years from the Gurkaynak, Sack, and Wright (2007) data. Table V presents the regression estimates. The average income gap appears to have significant forecasting power for bond excess returns, even after controlling for yields. However, as crucially emphasized by Bauer and Hamilton (2017), conventional statistics are misspecified to test the spanning hypothesis. We therefore also use Bauer and Hamilton (2017)'s bootstrap procedure with three and five PCs to test whether the average income gap is a spanned factor.[19] The bottom row of Table V reports p-values for this test. The results strongly reject the spanning hypothesis.

Given these results, it is natural to ask how the information captured by banks' income gap relates to the information contained in yields. We compare the predictive power of the various forecasting variables. In our sample, the first three PCs predict bond returns with an R^2 around 5%, whereas five PCs achieve an R^2 around 20%. This latter value is of similar magnitude to what we obtain with the income gap. In Internet Appendix Figure IA.4, we examine the evidence visually to better understand the relationship between various risk premium forecasts. Specifically, we plot forecasts of five-year Treasury bond excess returns using four methods: the income gap (thick line), Cochrane and Piazzesi (2005) (dotted line), and three and five PCs of yields (dashed line and solid line, respectively). All four forecasts exhibit broadly similar cyclical

[19] This procedure is similar to the bootstrap that we described above in Section I.B.1, except that the data-generating process for the PCs of yields automatically generates return dynamics.

Table V

Income Gap and Bond Excess Returns: Testing the Spanning Hypothesis

This table presents regressions of bond excess returns on banks' income gap. The sample period is 1986 to 2014. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and is defined as the difference between the amount of assets that reprice or mature within one year, and the amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. Income Gap is the average bank-level income gap. $rx^{(n)}$ is the excess one-year return of GSW zero-coupon bonds of maturity n. PC1 to PC5 are the principal components of GSW yields of maturity 1 to 10 years, divided by 100. Standard errors are computed using the reverse regression approach of Hodrick (1992). *, **, and *** indicate statistically different from zero at the 10%, 5%, and 1% level of significance, where p-values are computed using the bootstrap approach described in Section I.B.1. The last row reports the p-value of a test of the spanning hypothesis using the parametric bootstrap of Bauer and Hamilton (2017).

	Three Principal Components				Five Principal Components			
	(1) $rx^{(2)}$	(2) $rx^{(3)}$	(3) $rx^{(4)}$	(4) $rx^{(5)}$	(5) $rx^{(2)}$	(6) $rx^{(3)}$	(7) $rx^{(4)}$	(8) $rx^{(5)}$
Income Gap	−0.39***	−0.60***	−0.77***	−0.89***	−0.36***	−0.55***	−0.71***	−0.83***
	(0.11)	(0.18)	(0.25)	(0.31)	(0.11)	(0.18)	(0.25)	(0.31)
PC1	0.09	0.12	0.15	0.16	0.09	0.12	0.14	0.15
	(0.03)	(0.05)	(0.06)	(0.08)	(0.03)	(0.05)	(0.06)	(0.08)
PC2	0.17	0.34	0.53	0.72	0.17	0.34	0.52	0.72
	(0.17)	(0.28)	(0.38)	(0.47)	(0.17)	(0.28)	(0.38)	(0.47)
PC3	−0.46	−1.05	−1.61	−2.11	−0.38	−0.93	−1.46	−1.94
	(0.82)	(1.34)	(1.81)	(2.24)	(0.82)	(1.33)	(1.80)	(2.23)
PC4					−1.95	−2.45	−2.47	−2.07
					(2.15)	(3.39)	(4.51)	(5.49)
PC5					−4.27	−7.14	−10.06	−12.68
					(3.39)	(5.57)	(7.65)	(9.57)
Constant	0.03	0.06	0.08	0.09	0.05	0.08	0.11	0.13
	(0.02)	(0.03)	(0.05)	(0.06)	(0.03)	(0.04)	(0.05)	(0.07)
Spanning p-value	0.003	0.003	0.007	0.015	0.005	0.009	0.012	0.020
Observations	106	106	106	106	106	106	106	106
Adjusted R^2	0.417	0.428	0.426	0.414	0.442	0.451	0.447	0.433

variation. The measure based on three PCs is remarkably similar in this sample (a similarity is also present in terms of predictive R^2) and appears to depart significantly from the income gap. However, going up to five PCs brings yield forecasts much closer to the average income gap. The main difference between the two measures is the smoother pattern of the average income gap, reflecting the sticky nature of balance sheet quantities relative to asset prices. Echoing our regression results, it seems difficult to argue that one measure is much more informative than the other for forecasting bond excess returns. These results further support both the importance of these additional PCs, advocated, for example, by Duffee (2011) and Adrian, Crump, and Moench (2013), and the economic relevance of the income gap.

II. Theoretical Framework and Predictions

We provide a simple framework to interpret the empirical results presented in Section I.B. In particular, we consider a setting that abstracts from many relevant activities and risks that banks take on, such as credit risk, to focus solely on interest rate risk.

A. Model and Predictions

Assets. We assume that there are two main assets on banks' balance sheets. Short-term risk-free assets provide an instantaneous rate of return r_t, while long-lived assets provide a stream of payments $\theta e^{-\theta\tau}dt$ at each date $\tau \geq t$, like a console bond. The parameter θ controls the maturity of long-lived assets; the promised coupons add up to one, and their average maturity is $1/\theta$.

These two types of assets represent the saving and borrowing instruments available to the economy (productive assets, loans, corporate bonds, deposits, commercial paper, etc.). This separation into two categories allows us to consider differences between short-term and long-term fixed-rate instruments. The model also allows us to consider variable-rate assets. The latter instruments are equivalent to rolling over short-term assets, and thus, can be counted together with the short-term assets in our model.

Agents can also trade zero-coupon Treasury bonds of all maturities.[20] Since bonds of all maturities trade, the long-lived asset is redundant: a portfolio that consists of $\theta e^{-\theta\tau}$ bonds of each maturity τ replicates a unit long position in the long-lived asset. We denote by $P_t^{(\tau)}$ the price of the zero-coupon bond with maturity τ. We define the yield on this bond as $y_t^{(\tau)} = -\log(P_t^{(\tau)})/\tau$. Importantly, these Treasury bonds need not constitute a large part of banks' balance sheets. We nonetheless include them and use them for measurement because, as will become clear below, they are a simple instrument to measure the price of interest rate risk.

Banks. In each period, there is a continuum of banks indexed by i. Denote by $E_{i,t}$ the initial net worth of bank i at date t, by $X_{i,t}^{(\tau)}$ the bank's net dollar position in bonds of maturity τ and by $x_{i,t}^{(\tau)} = X_{i,t}^{\tau}/E_{i,t}$ the bank's position in bonds relative to its net worth. We drop the index i for aggregate quantities. A bank's net worth evolves according to

$$dE_{i,t} = \int_0^\infty X_{i,t}^{(\tau)} \frac{dP_t^{(\tau)}}{P_t^{(\tau)}} d\tau + \left(E_{i,t} - \int_0^\infty X_{i,t}^{(\tau)} d\tau\right) r_t dt. \tag{4}$$

Banks select their net holdings $X_t^{(\tau)}$ so as to maximize the instantaneous mean-variance criterion

$$\max_{\left\{X_{i,t}^{(\tau)}\right\}_\tau} \mathbb{E}(dE_{i,t}) - \frac{\gamma}{2E_{i,t}} \mathrm{var}(dE_{i,t}), \tag{5}$$

[20] All quantities are real. It would be straightforward to include an exogenous process for inflation in the model.

where γ is a risk aversion coefficient.[21] This objective can be rationalized in a setting where banks form overlapping generations, living for an infinitesimal interval dt, and maximize expected utility of final wealth as in Greenwood and Vayanos (2014).[22]

With this objective function, we capture the risk management decisions of banks without taking a particular stance on their origin. One interpretation of the risk aversion parameter γ is that it comes from the actual risk aversion of the bank's manager, or from her career concerns. Another interpretation is that γ is the Lagrange multiplier on a no-default condition for the bank or a regulatory risk constraint like value-at-risk limits.[23] Irrespective of its origin, risk aversion by banks is key in our theoretical framework. The fundamental underlying force for our results to hold is that banks trade off expected profits and risk in a stable way over time. This assumption is plausible as banks and regulators often explicitly express their concerns over interest rate risk. As an illustration, Bank of America states in its 2016 annual report that "Our overall goal is to manage interest rate risk so that movements in interest rates do not significantly adversely affect earnings and capital." In its Supervisory Insights, the Federal Deposit Insurance Corporation (FDIC (2005)) expresses the view that "Interest rate risk is fundamental to the business of banking." The presence of such a risk management motive is also supported by evidence coming from the cross-section of banks; see, for example, Drechsler, Savov, and Schnabl (2018), Kirti (2020), Rampini, Viswanathan, and Vuillemey (Forthcoming), and Vuillemey (2019). In Section IV.B, we investigate this risk management motive directly in the data. Our risk management theory implies that banks suffer when they hold significant balance sheet exposures and interest rates increase. Using banks' equity returns, we do not find empirical support for this conjecture. This result constitutes a challenge to the model and interpretation that we introduce in this section.

Equilibrium yield curve. Rather than completely specifying the model, we derive relations between the short rate, banks' investment decisions, and the yield curve that must hold in the equilibrium of any economy in which banks make risk management decisions as specified above. The relationships derived this way are the banking counterpart of household Euler equations for bonds of various maturities.

First, note that, scaled by their equity, all banks solve the same problem. Therefore, the optimal holdings per dollar of equity, $x_{i,t}^{(\tau)} = \frac{X_{i,t}^{(\tau)}}{E_{i,t}}$, are constant across banks. Let g_t be the net amount of long-term assets held by banks, divided by their equity. This quantity maps into holdings of the form

$$\forall \tau > 0, \quad x_{i,t}^{(\tau)} = g_t \theta e^{-\theta\tau}. \tag{6}$$

[21] Note that given the redundancy of the long-lived asset and zero-coupon bonds, banks simply maximize their holdings of the bonds without loss of generality.

[22] An alternative foundation would be to assume that banks are long-lived and their myopia comes from log utility.

[23] He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014), and Adrian and Shin (2013) are examples of more complete models of banks' risk appetite.

We study equilibria for which the joint dynamics of the short rate and the net position are of the form

$$dg_t = -\kappa_g(g_t - \bar{g})dt + \sigma_g dW_{g,t}, \tag{7}$$
$$dr_t = -\kappa_r(r_t - \bar{r})dt - \kappa_{g\to r}(g_t - \bar{g})dt + \sigma_r dW_{r,t}.$$

These simple processes capture key properties of the dynamics we observe empirically. We assume that $\kappa_g, \kappa_r > 0$, that is, both the exposure of banks to long-term assets and the short rate exhibit mean reversion. The term in $\kappa_{g\to r}$ allows the exposure g_t to predict future changes in the short rate. However, this is not a causal statement: we simply entertain the possibility that, in the data, there is a relationship between changes in the short rate dr_t and banks' exposure g_t. Additionally, note that the insights we derive hereafter hold for a larger family of specifications, for example, including other determinants of the short-rate dynamics.[24]

As in Greenwood and Vayanos (2014), we guess that yields are linear in the variables we specified, the short rate r_t, and the net exposure to long-lived assets g_t,

$$-\log\left(P_t^{(\tau)}\right) = y_t^{(\tau)} = A_r(\tau)r_t + A_g(\tau)g_t + C(\tau), \tag{8}$$

where $A_r(\tau)$ (respectively, $A_g(\tau)$) is the exposure of the yields of bonds with maturity τ to the short-term rate r_t (respectively, to the net exposure to long-lived assets g_t). These coefficients are an endogenous outcome of the model that we compute in equilibrium. Plugging in the law of motions of r_t and g_t, we obtain an expression for the expected bond returns that we denote by $\mu_t^{(\tau)}$.

Given this form for yields, we can easily write down banks' first-order condition with respect to their holdings in bonds of maturity τ:

$$\mu_t^{(\tau)} - r_t = A_r(\tau)\lambda_{r,t} + A_g(\tau)\lambda_{g,t}, \tag{9}$$

$$\text{where } \lambda_{j,t} = \gamma\sigma_j^2 \int_0^\infty x_{i,t}^{(\tau)} A_j(\tau)d\tau, \text{ for } j = g, r. \tag{10}$$

This condition is akin to a standard Euler equation. The first line tells us that for a bond of a given maturity, the bank requires a risk premium proportional to the exposures ($A_j(\tau)$) of the bond to the fundamental shocks of the economy. The second line characterizes the amount of compensation requested for bearing each of these risks. This compensation is proportional to the product

[24] Consider more state variables z_t to capture the dynamics of interest rates and the income gap (e.g., inflation, employment, etc.). As long as this joint system follows a continuous-time VAR(1), that is, the vector $\zeta_t = (r_t, g_t, z_t)$ follows

$$d\zeta_t = -K(\zeta_t - \bar{\zeta})dt + \Sigma_\zeta dW_t.$$

Proposition 1 will hold.

of the risk aversion γ, the risk σ_j^2, and the total exposure accumulated through positions in bonds of various maturities.

Plugging the equilibrium portfolio positions into banks' first-order condition, we obtain the equilibrium risk premia. We provide detailed calculations and proofs, and we verify the conjectured form of prices, in Internet Appendix Section IA.I.

PROPOSITION 1: *Consider an equilibrium in which banks' balance sheet positions and the short rate are given by equation (7). The expected excess return $\mu_t^{(\tau)}$ on the τ-maturity bond is proportional to banks' net position in long-term assets g_t,*

$$\mu_t^{(\tau)} - r_t = g_t \times \left(c_r A_r(\tau) + c_g A_g(\tau)\right) = g_t \times \phi(\tau), \tag{11}$$

where c_r and c_g are constants determined in equilibrium and $\phi(\tau) > 0$.

Proposition 1 shows that the risk premium on a bond of maturity τ is positively correlated with banks' net exposure to long-term assets g_t. When banks hold more long-term assets, they stand to lose more if interest rates increase. As a consequence, they are less willing to hold zero-coupon bonds of various maturity as they lose value at the same time. In an equilibrium in which banks do not decide to change their positions in these bonds, the expected return must have adjusted to compensate for this lower willingness to bear risk. Thus, in equilibrium, higher net exposure is correlated with a more significant bond risk premium.

We can further characterize the relationship between bond risk premia and banks' net exposure across maturities.

PROPOSITION 2: *Consider an equilibrium in which equation (7) describes the relationship between banks' balance sheet position and the short rate. The expected excess returns of bonds of longer maturity are more sensitive to the net exposure of banks: $\phi(\tau)$ is strictly increasing in τ.*

Proposition 2 shows that more significant exposure to long-term assets predicts a higher risk premium for bonds with longer maturities. Indeed, longer maturity bonds are riskier—their exposure to variation in interest rates is higher than the exposure of short-maturity bonds. As a result, following an increase in the net exposure to long-term assets, holding risk premia constant, banks are relatively less willing to hold bonds of longer maturity. Thus, as banks' net exposure increases, the equilibrium risk premium on bonds of longer maturity will increase more than the risk premium on shorter maturity bonds.

Our model thus makes two direct testable predictions: (i) a larger average net exposure of banks to long-term assets should predict higher bond risk premia, and (ii) this effect should be stronger for longer maturity bonds. These predictions correspond to the empirical results established in Section I.

B. Additional Considerations

Excess returns versus yields. Our main predictions link bond risk premia $\mu_t^{(\tau)} - r_t$ and banks' balance sheets g_t. Equation (8) suggests other testable implications that link banks' balance sheets and yields directly. However, the sign and magnitude of this relation between banks' net position and yields depend on the joint dynamics of rates and positions. This result is in contrast to the relationship between banks' net position and bond excess returns, the sign of which is unambiguous in our model. The following proposition illustrates this property.

PROPOSITION 3: *Consider an equilibrium in which banks' balance sheet position g_t and the short rate r_t are given by equation (7). Then the exposure of bond prices to the net position g_t, $A_g(\tau)$, is of the same sign as $\gamma\sigma_r^2\frac{1}{\theta+\kappa_r} - \kappa_{g\to r}$. The net position g_t is an unspanned factor if and only if $\kappa_{g\to r} = \gamma\sigma_r^2\frac{1}{\theta+\kappa_r}$.*

In equilibrium, yields depend on not only current risk premia but also expectations of future rates. Additionally, risk management by banks creates a link between risk premia and the banks' balance sheet composition. Periods of large risk holdings by banks correspond to periods of large risk premia and high yields. However, risk management by banks does not constrain the relation between short-rate expectations and banks' balance sheets. If periods of high long-term holdings g_t happen to coincide with periods in which the short rate decreases (positive $\kappa_{g\to r}$), then yields should be lower when banks' net position increases to reflect expectations of future rates. Proposition 3 characterizes which of the two effects dominate. For one particular parameter value, $\kappa_{g\to r} = \gamma\sigma_r^2\frac{1}{\theta+\kappa_r}$, the two effects cancel each other and yields of all maturities do not depend on g_t. The net exposure g_t is then an unspanned factor. Close to this knife-edged case, $\kappa_{g\to r} \approx \gamma\sigma_r^2\frac{1}{\theta+\kappa_r}$, the role of the income gap in explaining yield dynamics is quantitatively limited. This situation echoes our empirical finding that while not outperforming yields, the information contained in the income gap about bond risk premia reflects the information contained in the higher order component of yields (Table V and Internet Appendix Figure IA.4).

Completing the model. Here, we sketch out an economy in which equation (7) describes the joint dynamics of the short rate r_t and banks' net position g_t. To do so, we specify the supply of other assets as well as the behavior of other market participants.

We first assume the existence of an instantaneous risk-free asset that is in perfectly elastic supply at rate r_t, which is now a primitive of the model. We also assume that long-lived assets are in finite supply, while zero-coupon bonds are in zero net supply. Next, we introduce a second group of agents in addition to banks—households. We consider households in an extended sense, that is, we pool them together with nonfinancial firms and the government. Households are endowed with the entire supply of long-lived assets. In addition, at each date t, households borrow from banks an exogenous amount B_t of the long-lived asset and lend to banks an exogenous amount L_t of long-term assets.[25] We define the net imbalance, $-g_t$, as the difference between the ratio of long-term savings to total bank equity, $l_t = L_t/E_t$, and long-term borrowing to total bank equity, $b_t = B_t/E_t$. It follows that g_t is now also a primitive of the model. Then as long as the exogenous laws of motion for r_t and g_t are given by equation (7), we obtain the equilibrium yield curve in equation (8).

The assumption of exogenous changes in households' and firms' portfolios is a simplification of a more complicated decision problem—households' and firms' savings and borrowing decisions. Exogenous shocks to l_t and b_t are meant to capture the fact that factors other than simple risk-return tradeoffs influence those decisions. For instance, changing liquidity needs, incorrect heuristics, or hedging demands can affect those decisions. We come back to potential empirical counterparts of these shocks in Section IV.A.2.

III. Model Estimation

In this section, we take the model presented in Section II to the data. This exercise allows us to consider the ability of our theory to quantitatively rationalize the relationship between banks' balance sheets and expected returns throughout the yield curve. In addition, we obtain an estimate of banks' willingness to take risk, γ, a key parameter of our model, which is also central to many macroeconomic models with intermediaries.

To define the state variable g_t, we use equation (2) and construct the bank-level gap g_{it}, which corresponds to $1 - \textit{Income Gap}_{it} \times \frac{\textit{Assets}_{it}}{\textit{Equity}_{it}}$ (see Section I.A.1). Internet Appendix Table IA.VI confirms that the predictive results of Section I.B hold for this measure as well. To estimate the dynamics of the model's state variables, we discretize the model using one year as the time unit. More specifically, if t is a year, we estimate the following equations, which correspond to the discrete-time version of equation (7):

$$\begin{cases} y_{t+1}^{(1)} - y_t^{(1)} = -\kappa_r y_t^{(1)} - \kappa_{g\to r} g_t + \epsilon_{r,t} \\ g_{t+1} - g_t = -\kappa_g g_t + \epsilon_{g,t}, \end{cases}$$

where g_t is the exposure measure defined above measured in the first quarter of the year. We use a parametric bootstrap to correct the estimates of κ_r, κ_g, and $\kappa_{g\to r}$ for small-sample bias, and σ_r and σ_g are estimated as the empirical standard deviation of ϵ_r and ϵ_g, respectively. We estimate $\phi(\tau) = (c_r A_r(\tau) + c_g A_g(\tau))$ defined in Proposition 1 by regressing the one-year excess returns of bonds of maturity τ on g_t at the yearly frequency. We calibrate $\frac{1}{\theta}$, the time to maturity of the long-term asset, to 10 years.[26] We use these estimated coefficients and the calibrated θ to compute $\hat{I}_r = \frac{1}{\hat{\kappa}_r}(1 - \frac{\theta}{\theta + \hat{\kappa}_r})$ and $\hat{A}_r(\tau) = \frac{1 - e^{-\hat{\kappa}_r \tau}}{\hat{\kappa}_r}$.

[25] We can easily relax the exogeneity assumptions and allow borrowing and lending by households to be price-elastic. This would not change the qualitative predictions that we have derived so far.

[26] To assess the importance of this calibration, below we provide a sensitivity analysis using a range of alternative values for θ.

Table VI
Model Estimation

This table presents the model's parameter estimates. The estimation procedure is described in details in Section III. The 90% CI corresponds to bootstrapped 90% confidence intervals.

Parameters	Estimates	90% CI
κ_r	0.146	[−0.033, 0.295]
κ_g	0.062	[−0.164, 0.263]
κ_{gr}	0.019	[0.012, 0.029]
σ_r	0.0001	[0.0001, 0.0002]
σ_g	0.071	[0.037, 0.093]
$\hat{\phi}(2)$	0.016	[0.008, 0.026]
$\hat{\phi}(3)$	0.030	[0.016, 0.047]
$\hat{\phi}(4)$	0.042	[0.023, 0.064]
$\hat{\phi}(5)$	0.051	[0.026, 0.079]
γ	19.221	[1.767, 75.995]

Finally, the risk aversion coefficient, γ, is estimated by minimizing the squared distance between the average $\phi(\tau)$ across maturities $(\frac{1}{4}\sum_{\tau=2}^{5}\hat{\phi}(\tau))$ and their theoretical counterpart. Table VI presents the coefficient estimates. The estimated risk aversion is about 19. Given banks' optimization problem in equation (5), this risk aversion coefficient corresponds to a relative risk aversion coefficient. The model reveals much larger risk aversion than the typical calibration in macroeconomic models with a financial sector. He and Krishnamurthy (2014) use a relative risk aversion coefficient of two; Brunnermeier and Sannikov (2014) use log-utility. This number is within the range of estimates in Greenwood and Vayanos (2014), who use variation in Treasury supply to identify the *absolute* risk aversion of *all* arbitrageurs in fixed-income markets. Because their method does not observe arbitrageurs directly, they must make assumptions about arbitrageurs' wealth, giving rise to a wide range of plausible estimates. We can obtain point estimates without such assumptions because we measure banks' portfolios directly. If we assume that banks constitute the entire group of arbitrageurs in the market for interest rate risk, their estimates are a third of ours.[27] The literature on quantitative easing interventions reports estimates of risk aversion that are about 85% of what we find here.[28] The other interpretation of this difference is that banks are only a subset of the arbitrageurs in this market. Under this view, banks constitute a

[27] Greenwood and Vayanos (2014) estimate the absolute risk aversion of arbitrageurs to be 57. To obtain an estimate of relative risk aversion, assuming that arbitrageurs are only banks, one needs to multiply this number by the total banking sector's capital as a fraction of GDP. In 2015, U.S. GDP was $18tn, banks' total assets amounted to $17tn, and banks' capital to assets ratio was 11.7%. These numbers imply a relative risk aversion of 6.3.

[28] D'Amico and King (2013) find supply effects two and a half times those in Greenwood and Vayanos (2014), and Hamilton and Wu (2012) report an absolute risk aversion twice as large.

sizable part of this group, between a third and 85% depending on the supply-response estimates.

We confirm the model's ability to fit risk premium dynamics across maturities. Internet Appendix Figure IA.5 highlights the model's goodness of fit by comparing the empirical estimates $\hat{\phi}(\tau)$—the coefficient estimates obtained when regressing bond excess returns of maturity τ on g—with their model-implied counterpart. For the two-year bond, the model slightly overestimates the sensitivity of bond risk premia to the income gap; at all other horizons, the model-implied estimates and the empirical estimates are very close. Internet Appendix Figure IA.6 investigates robustness relative to our calibration for θ. Our baseline estimation uses an average time to maturity for the long-term asset of 10 years ($\theta = 0.1$). Internet Appendix Figure IA.6 reestimates our model using different values of $1/\theta$ ranging from two to 50. For high values of θ ($\approx$ 50), the risk aversion coefficient is estimated at 10, while a time to maturity of five years leads to a relative risk aversion for banks of 25.

IV. Interpretation

The model developed in Section II provides a simple interpretation of our results. In this section, we first present a collection of empirical observations that support this view. We then discuss several challenges to this interpretation.

A. Supporting Evidence

A.1. The Income Gap or Other Balance Sheet Quantities?

Our theory relates the quantity of interest rate risk borne by banks to the market price of this risk. Empirically, the predictive power of the income gap may come from specific features of banks' average balance sheets that happen to correlate with bond risk premia for reasons unrelated to banks risk management of interest rate risk. To the extent that this is the case, we should observe similar or higher predictability using dimensions of banks' balance sheets that do not focus on the net exposure to interest rate risk. In what follows, we consider two particular dimensions.

The income gap and its components. Our first tests separately estimate the ability of the asset and liability components of the income gap to forecast bonds' excess returns. Internet Appendix Figure IA.7 shows these two components: "Nonexposed assets" correspond to the average bank-level ratio of assets that reprice or mature within one year normalized by total consolidated assets (in blue); "Nonexposed liabilities" is minus the average bank-level ratio of liabilities that reprice or mature within one year normalized by total consolidated assets (in red). If the forecasting power of the income gap comes only from, say, its asset side (the blue line), then our interpretation cannot be valid—under our theory, only banks' total portfolio exposure should forecast bonds' excess returns. Since the liability side of the gap varies significantly in the time series, such a result would invalidate our interpretation.

Table VII

Asset and Liability Risk Exposure and Bond Excess Returns

This table presents regressions of bond excess returns on banks' asset and liability risk exposure. The sample period is 1986 to 2014. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year, and the dollar amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. Income Gap is the average bank-level income gap. $rx^{(n)}$ is the excess one-year return of GSW zero-coupon bonds of maturity n. Nonexposed Assets is the average bank-level ratio of assets that reprice or mature within one year normalized by total consolidated assets. - Nonexposed liabilities is minus the average bank-level ratio of liabilities that reprice or mature within one year normalized by total consolidated assets. Standard errors are computed using the reverse regression approach of Hodrick (1992). *, **, and *** indicate statistically different from zero at the 10%, 5%, and 1% level of significance,where p-values are computed using the bootstrap approach described in Section I.B.1.

	(1) $rx^{(5)}$	(2) $rx^{(5)}$	(3) $rx^{(5)}$	(4) $rx^{(5)}$
Income Gap	-0.55^{**}			
	(0.27)			
Nonexposed Assets		-0.10		-0.50^{*}
		(0.18)		(0.27)
- Nonexposed Liabilities			-0.36	-0.78^{**}
			(0.23)	(0.34)
Constant	0.10^{**}	0.07	-0.08	0.01
	(0.04)	(0.08)	(0.07)	(0.08)
Observations	106	106	106	106
Adjusted R^2	0.189	0.001	0.081	0.233

To implement this test, we simply replace the *Income Gap* in equation (3) by its two components "Nonexposed assets" and "Nonexposed liabilities." Table VII reports the results. For brevity, we only show the estimated coefficients when the dependent variable is the excess return on five-year bonds. Column (1) replicates the results of column (4) (Table II). Columns (2) and (3) show that taken individually, neither of the two components of the average income gap robustly forecasts future bond excess returns. The estimated coefficients have low statistical significance and are small in magnitude. In column (4), we include the two components in the regression simultaneously. Both coefficients become statistically significant and of a magnitude close to that of the income gap alone. Thus, consistent with our interpretation, only the overall exposure of the average financial intermediary explains bond risk premia.

Interest rate risk versus liquidity risk. We next examine whether balance sheet aggregates focusing on liquidity risk predict bond returns. We consider the liquidity mismatch index (LMI) of Bai, Krishnamurthy, and Weymuller (2018), the bank liquidity creation index of Berger and Bouwman (2009), which we equal-weight or value-weight by total gross assets across banks (BB), and the Basel Committee's liquidity coverage ratio (LCR) as constructed by Choi and Choi (2016).[29]

The measures of liquidity risk behave differently from banks' average income gap. Internet Appendix Figure IA.8 plots these four measures jointly with the income gap, standardized to have unit standard deviation. We flip the sign of the BB measure so that low values correspond to high liquidity risk, like LMI and LCR. The time-series behavior of liquidity risk differs from the income gap in at least three ways. First, while the average income gap evolves smoothly around the crisis, LMI experiences a sharp drop and then rebounds right after the financial crisis.[30] Second, the liquidity measures all exhibit strong growth in the postcrisis period. Third, in the earlier part of the sample, these measures show a slow secular increase in liquidity risk, while the average income gap shows substantial cyclical variation.

The measures of liquidity risk do not predict bond returns. Internet Appendix Table IA.VII reports our baseline predictive regressions, replacing the income gap by the value-weighted BB index, which features an extended sample and echoes the behavior of the other measures in the late part of the sample. None of the coefficients is statistically significant, and the adjusted R^2s are all well below 1%. Interestingly, these results also dispel the idea that it would be enough to exhibit somewhat of a downward trend to capture bond risk premia.[31]

A.2. Demand Shocks and Bond Risk Premia

We consider three measures of the "demand" for savings and borrowing instruments and examine whether they forecast excess bond returns: (i) the aggregate demand for ARMs, (ii) the aggregate demand for deposits, and (iii) the aggregate supply of government bonds. These measures correspond to the supply/demand shocks of long-lived assets in the model of Section II. In particular, none of these shocks should have significant forecasting power for bond risk premia above and beyond the average income gap. Indeed, as per our model, a sufficient statistic for bond risk premia is the net interest rate risk held by banks, as captured by the income gap. Empirically, we consider three such observable shifts in quantities.

First, households' choice of fixed-rate mortgages versus ARMs depends on multiple factors, which can change over time.[32] Using the Monthly Interest

[29] We thank the authors of this work, who graciously shared their data with us.

[30] Bai, Krishnamurthy, and Weymuller (2018) argue that this feature reflects the superiority of their measure, which also leads to better ability to capture how cross-sectional differences in liquidity risk are related to bank lending.

[31] We reproduced this analysis for the equal-weighted measure and continued to find no significant coefficients and all R^2s below 1%. The shorter sample measures, LMI and LCR, do not perform better, with no coefficients distinct from zero and low R^2s of about 2% even in the short sample.

[32] This choice involves a risk-return trade-off and households may use simple imprecise heuristics to make decisions (Koijen, Hemert, and Nieuwerburgh (2009)). This choice also partly reflects

Rate Survey, we compute the quarterly ratio of ARM issuance to total mortgage issuance. To the extent that shifts in household demand are the source of some of this variation, as in our model, an increase in the share of ARMs in total mortgage issuance forces banks to hold more ARMs. Everything else equal, banks' average income gap should decrease. Internet Appendix Figure IA.9, Panel A, shows that there is indeed a positive correlation (59%) between the share of ARMs in mortgage issuance and the average income gap, at least until 2006.[33]

We also consider the average quarterly bank-level ratio of noninterest-bearing deposits normalized by total consolidated assets. When households increase their relative demand for noninterest-bearing deposits, banks end up in equilibrium with more interest-rate-sensitive liabilities. Thus, everything else equal, their income gap increases. There are several time-varying determinants of the demand for noninterest-bearing deposits. Depositors have a choice between many stores of wealth, which, beyond a standard risk-return trade-off, will be determined by liquidity considerations (Tobin (1956), Baumol (1952)) or demand for safety (Krishnamurthy and Vissing-Jorgensen (2012)).[34] Internet Appendix Figure IA.9, Panel B, plots the time-series evolution of noninterest-bearing deposits and highlights its positive correlation (64%) with the average income gap.

Finally, we consider the aggregate supply of government bonds. We use the maturity-weighted supply of Treasuries, normalized by GDP, as in Greenwood and Vayanos (2014). By varying the supply of long-term bonds in the economy, the government may shift the availability of interest rate risk. For instance, to fund an expansionary fiscal policy, the government will increase the Treasury supply, and, in equilibrium, banks' income gap should decrease. Internet Appendix Figure IA.9, Panel C, plots the time-series evolution of the maturity-weighted Treasury supply measure. Given the low-frequency fluctuations in Treasury supply, this series does not exhibit much correlation with the average income gap.

We also investigate whether these measured fluctuations in the supply/demand for interest rate risk predict bond risk premia. We start by replacing the average income gap by each of the "demand" factors in equation (3). We then add the average income gap to the forecasting regression. The estimated coefficients are presented in Table VIII. Columns (1) and (5) show that neither the share of ARMs nor the maturity-weighted Treasury supply measure of Greenwood and Vayanos (2014) significantly correlate with future bond excess returns. Thus, unsurprisingly, columns (2) and (6) of Table VIII show that the forecasting power of the average income gap is not affected by the

the desire of households to manage their liquidity, which may depend on aggregate factors (Chen, Michaux, and Roussanov (2013)).

[33] Of course, this unconditional positive correlation does not have to be present, since other shifts in the demand for other components of banks' balance sheets may force them to adjust their income gap in an opposite direction.

[34] For instance, the fraction of noninterest-bearing deposits exhibits a correlation of 46% with the HP-filtered monetary aggregate M1.

Table VIII

Changing Asset Quantities and Bond Excess Returns

This table presents regressions of bond excess returns on proxies for the demand and supply of interest rate risk. The sample period is 1986 to 2014. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year, and the dollar amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. Income Gap is the average bank-level income gap. $rx^{(5)}$ is the excess one-year return of GSW zero-coupon bonds of five-year maturity. ARM fraction of issuance is the quarterly share of adjustable-rate mortgages in total mortgage issuance, from the Monthly Interest Rate Survey. Non int.-bearing deposits is the average of the quarterly bank-level ratio of noninterest-bearing deposits normalized by total consolidated assets. Mat.-weighted Debt/GDP is the maturity-weighted Treasury supply measure of Greenwood and Vayanos (2014). Standard errors are computed using the reverse regression approach of Hodrick (1992). *, **, and *** indicate statistically different from zero at the 10%, 5%, and 1% level of significance, where p-values are computed using the bootstrap approach described in Section I.B.1.

	(1) $rx^{(5)}$	(2) $rx^{(5)}$	(3) $rx^{(5)}$	(4) $rx^{(5)}$	(5) $rx^{(5)}$	(6) $rx^{(5)}$
Income Gap		−0.83**		−0.51*		−0.60**
		(0.32)		(0.27)		(0.26)
ARM Fraction of Issuance	−0.02	0.09				
	(0.06)	(0.07)				
Non Int.-Bearing Deposits			−0.94	−0.15		
			(0.63)	(0.64)		
Mat.-Weighted Debt/GDP					−0.00	0.01
					(0.01)	(0.01)
Constant	0.04	0.11*	0.15*	0.12	0.04	0.09
	(0.03)	(0.04)	(0.08)	(0.08)	(0.04)	(0.05)
Observations	98	98	106	106	106	106
Adjusted R^2	−0.004	0.269	0.087	0.182	−0.007	0.194

inclusion of these two variables in equation (3). Column (3) shows an R^2 of 8.7% when using the average ratio of noninterest-bearing deposits normalized by consolidated assets to predict returns on five-year bonds. However, column (4) shows that after we control for the average income gap, noninterest-bearing deposits no longer correlate with bond risk premia, while the income gap remains a statistically significant predictor with an economic magnitude similar to our baseline specification.

All of these results are consistent with our interpretation: the net exposure to interest rate risk borne by banks, as measured by the average income gap, appears to better capture variations in expected excess bond returns than quantities of particular types of financial assets.

A.3. Risk-Sharing Evidence

We further exploit information on the income gap at the bank level to study the time-series behavior of the income gap across heterogeneous banks. In our

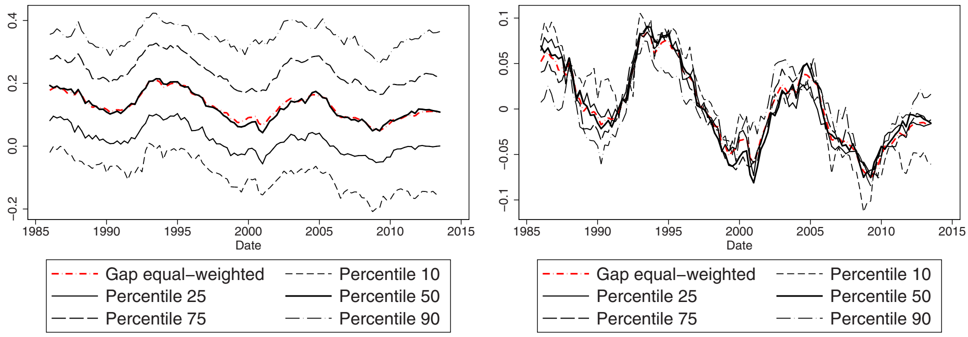

Panel A. Raw Distribution

Panel B. Demeaned

Figure 2. Cross-sectional distribution of banks' income gap. This figure plots the cross-sectional distribution of banks' income gap. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and is defined as the difference between the amount of assets that reprice or mature within one year and the dollar amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. We compute the various percentiles of the income gap for each date on the top panel. The bottom panel presents the demeaned time series (Color figure can be viewed at wileyonlinelibrary.com)

model, the equilibrium risk premium adjusts so that banks are collectively willing to bear the interest rate risk supplied by other agents in the economy. Even if banks face customers with heterogeneous demand, they can use financial markets to share interest rate risk. To the extent that banks have similar risk preferences, they would end up with the same net exposure. Therefore, even across heterogeneous banks, one would expect to find common variation in their income gap, which is close to the average income gap. We find evidence in line with this risk-sharing view using three sources of heterogeneity across banks: size, geography, and leverage.

Panel A of Figure 2 presents the time series of the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th} percentiles of the cross-sectional distribution of the gap each quarter. There is substantial cross-sectional variation in the income gap across banks: the interquartile range is about 20%. However, the entire distribution appears to shift up and down over time, suggesting common variation. Panel B presents the demeaned time series of the various percentiles, which reinforces this point—the series are all strongly positively correlated.

The first dimension of bank heterogeneity that we consider is size.[35] We split banks into 10 groups based on decile of total assets. Figure 3 presents the average income gap for each size group. All series are remarkably similar to the average income gap except for the largest size group, for which we do not capture the income gap accurately, most likely because of this group's use of

[35] For example, Kashyap and Stein (2000) document variation in bank balance sheet composition across the size distribution.

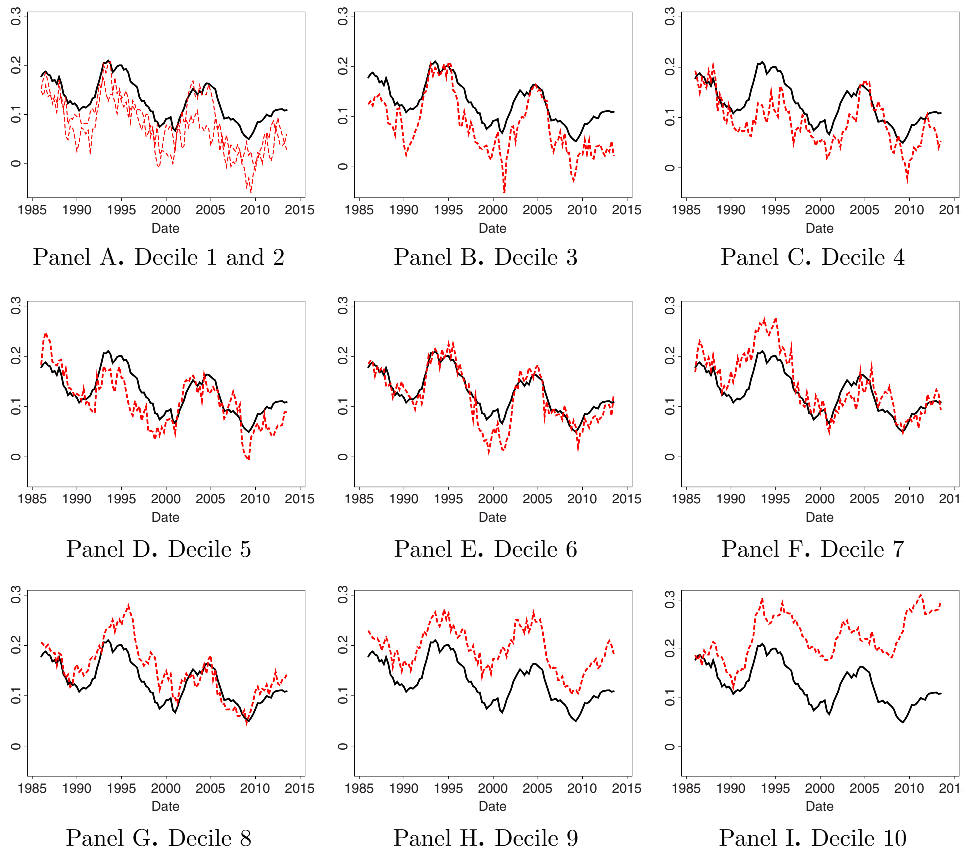

Figure 3. Income gap across bank sizes. The figure plots the time series of banks' income gap for banks of different sizes. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year and the dollar amount of the liabilities that reprice or mature within one year, all scaled by total consolidated assets. The solid black line represents the average income gap, computed across all U.S. bank holding companies with total consolidated assets of $1 bn or more. For each panel, the dashed line corresponds to the average income gap within a decile of total consolidated assets, in increasing order. We represent the first two deciles in the first panel (Color figure can be viewed at wileyonlinelibrary.com)

interest rate derivatives. The correlation with the average income gap is about 85% for each size group except the fourth (72%) and the tenth (18%).

In Figure 4, we repeat this comparison across the nine U.S. Census divisions. Because of heterogeneity in local economic conditions, one would expect banks in different regions to face different demand for interest rate exposure. However, the figure shows that across these nine regions, banks share similar net exposure to interest rates. The local average income gaps all exhibit a strong correlation with the national average income gap. All correlations are

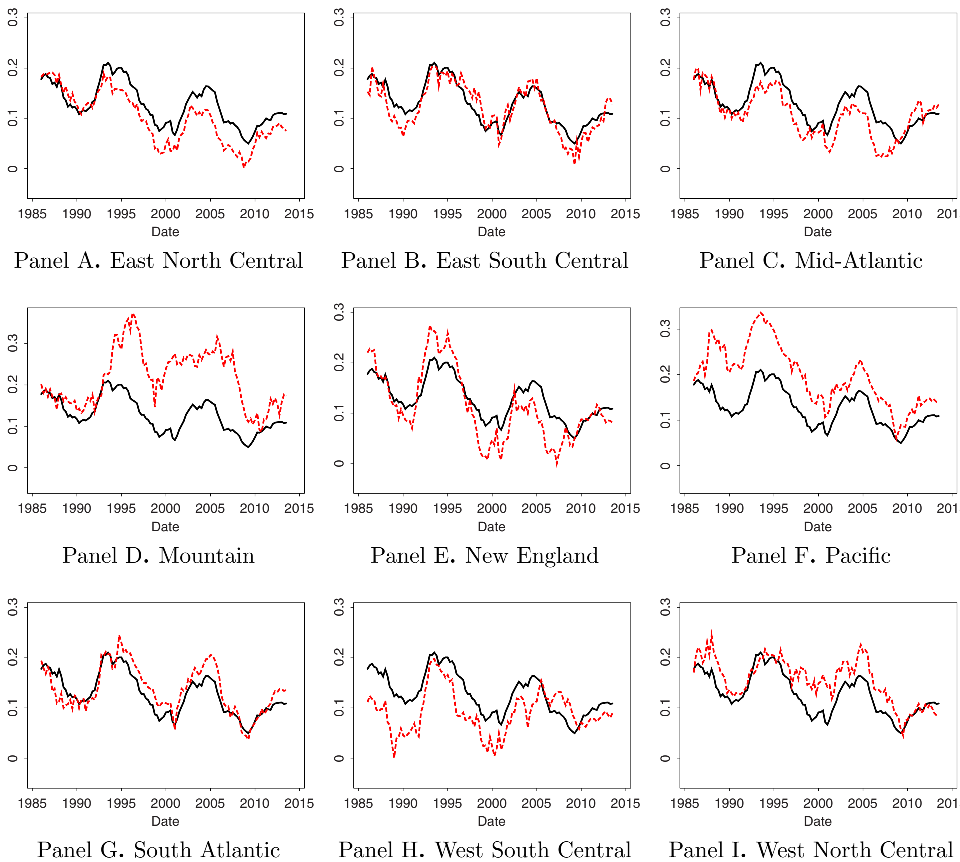

Figure 4. Income gap across U.S. Census divisions. The figure plots the time series of banks' income gap for banks located in different geographical areas. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year and the dollar amount of the liabilities that reprice or mature within one year, all scaled by total consolidated assets. The solid black line represents the average income gap, computed across all U.S. bank holding companies with total consolidated assets of \$1 bn or more. For each panel, the dashed line corresponds to the average income gap within one of the nine divisions of the U.S. Census (Color figure can be viewed at wileyonlinelibrary.com)

between 80% and 90% except for the West South Central division (Panel H, 63%) and the Mountain division (Panel D, 45%).[36]

Finally, we compare banks that vary directly in the composition of their balance sheets. For each bank, we compute the equity-to-assets ratio and deposits-to-assets ratio. Book leverage consists of the ratio of book equity over consolidated assets. The fraction of deposits is the ratio of checking deposits

[36] The latter is the only substantial deviation, likely caused by individual measurement error as the Mountain region has the lowest number of banks in our sample, between 7 and 23 per quarter.

to total assets. For each of those characteristics, we split our sample into two groups. Panel A of Figure 5 presents the average income gap of banks sorted on equity-to-assets. The top group has an average ratio of 10% and the bottom one of 7%. The average income gap for each group does not exhibit any substantial deviation from each other. Panel B compares the two groups based on deposits-to-assets. The average level for the ratio for the two groups is 8% and 17%, a distinction reflecting the likely exposure to interest rate risk of deposits. Over time, the two series also exhibit a strong positive correlation.

B. Challenges to Interpretation

We now discuss several challenges to the interpretation of our results described in Section II. These challenges all relate to a particular assumption that is central to the banking view of bond risk premia. Specifically, our model assumes a risk management objective for banks—when banks hold significant balance sheet exposures and the interest rate increases, banks' value should decrease. This assumption underlies banks' risk management decisions, which, in turn, drive the relationship between the average income gap and excess returns on Treasury bonds. In this section, we use banks' equity returns to test this assumption and its implications.

Our first test investigates the empirical relationship between banks' equity returns and Treasury returns, and how the average income gap affects this relationship. In the simple framework of Section II, when banks' average income gap is low, banks' value should decrease when interest rates increase. Measuring banks' value using their equity returns, we expect the correlation between realized bank returns and realized bond returns to be significantly lower when banks' average income gap is low. We test this hypothesis on quarterly data from 1986 to 2014. We measure bank returns as the excess one-quarter return of the Fama-French industry portfolio for banks. We measure bond returns as the excess one-quarter return of the Fama portfolio of bonds with maturities ranging from 5 to 10 years. We then estimate

$$rx_{banks,t} = b_0 + b_1 r_{Tbond,t} + b_2 Gap_{t-1} + b_3 Gap_{t-1} \times r_{Tbond,t}. \tag{12}$$

Table IX presents the estimation results. Periods of a low-income gap appear to be *negatively*, rather than positively, related to banks' exposure to bond returns (i.e., $b_3 > 0$). This result suggests that when banks' balance sheets exhibit significant exposure, a rise in interest rate increases, rather than decreases, banks' equity value. This result is inconsistent with the risk management motive for banks highlighted in our model. Note, however, that once we control for equity market returns (column (2)), we find that b_0, b_1, and b_2 become insignificant.[37] With this added control, the relation between banks' income gap and

[37] While not a direct consequence of the theoretical model of this paper, controlling for market exposures to understand intermediary risk is a frequent feature of intermediary asset pricing models. See, for instance, Adrian, Etula, and Muir (2014) and He, Kelly, and Manela (2017).

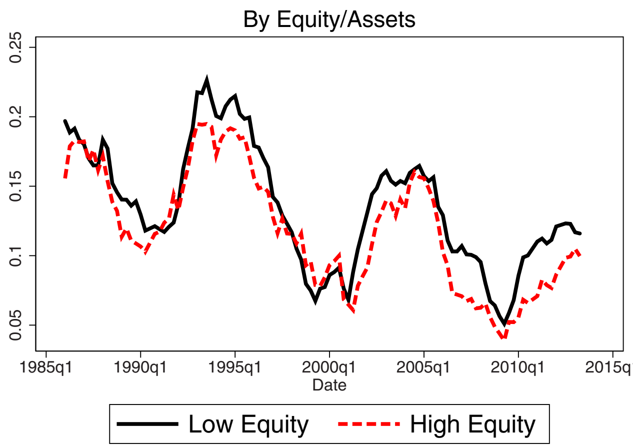

Panel A. Equity/Assets

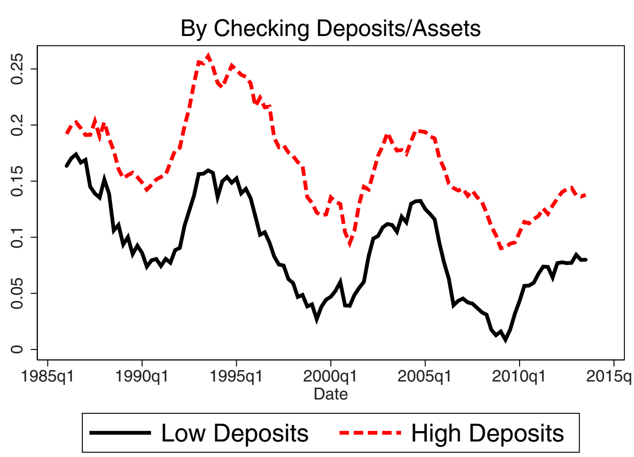

Panel B. Checking Deposits/Assets

Figure 5. Income gap across bank characteristics. This figure plots the time series of banks' income gap for banks with different equity to asset ratios (Panel A) and banks with different checking deposit to asset ratios (Panel B). The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year and the dollar amount of the liabilities that reprice or mature within one year, all scaled by total consolidated assets. Panel A presents the average income gap for banks split into two groups based on the average value of the ratio of book equity to consolidated assets. Panel B presents the average income gap for banks split into two groups based on the average value of the ratio of checking deposits to assets (Color figure can be viewed at wileyonlinelibrary.com)

Table IX

Banks' Stock Return Exposure to Treasuries

This table present regressions of banks' excess stock returns on banks' average income gap. The sample period is 1986 to 2014. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year, and the dollar amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. Income Gap is the average bank-level income gap. rx_{banks} is the excess one-quarter return of the Fama-French industry portfolio for banks. rx_{Tbond} is the excess one-quarter return of the Fama portfolio of bonds with maturities ranging from 5 to 10 years. rx_{mkt} is the excess one-quarter return of the CRSP value-weighted index. Newey-West standard errors with a bandwidth of two years are reported in parentheses. *, **, and *** indicate statistically different from zero at the 10%, 5%, and 1% level of significance.

	(1) $rx_{banks,t}$	(2) $rx_{banks,t}$
Income Gap_{t-1}	0.06	0.10
	(0.26)	(0.19)
$rx_{Tbond,t}$	-4.35^{***}	-0.17
	(1.08)	(0.78)
Income $\text{Gap}_{t-1} \times rx_{Tbond,t}$	28.13^{***}	3.65
	(8.22)	(5.74)
$rx_{mkt,t}$		1.09^{***}
		(0.09)
Constant	0.02	-0.01
	(0.04)	(0.03)
Observations	109	109
Adjusted R^2	0.124	0.620

the covariance between bank returns and bond returns is neither economically nor statistically significant.

We can also construct a more structural test that builds on the particular mean-variance framework used in Section II. In our model, banks aim to limit the volatility of their equity value. This particular specification of banks' risk management objective implies that bond risk premia should be proportional to the expected covariance between banks' equity returns and bond returns. Table X, column (2), tests this prediction empirically. Each quarter, we compute the covariance of excess daily returns from the Fama portfolio of Treasuries with maturities ranging from 5 to 10 years and the Fama-French industry index for banks. We then construct a forecast of this quantity using an AR(1) model and use this forecast to predict bond returns. The coefficient is small and insignificant. Moreover, the regression's adjusted R^2 is negative. Thus, there is no meaningful relationship between bond excess returns and the predicted covariance between daily excess returns on long-term bonds and banks' stock returns. This analysis rejects the mean-variance framework that we use to motivate banks' risk management decisions at the heart of the model in Section II.

Table X

Alternative Risk Measures and Bond Excess Returns

This table presents regressions of bond excess returns on alternative measures of banks' risk. The sample period is 1986 to 2014. The bank-level income gap is computed from the quarterly Consolidated Financial Statements (Files FR Y-9C) and corresponds to the difference between the dollar amount of assets that reprice or mature within one year, and the $ amount of liabilities that reprice or mature within one year, all scaled by total consolidated assets. Income Gap is the average bank-level income gap. $rx^{(n)}$ is the excess one-year return of GSW zero-coupon bonds of maturity n. The conditional variance and covariance forecasts are constructed in two steps: first compute realized values using daily returns for each month, then create a forecast by estimating an AR(1) in the full sample. Standard errors are computed using the reverse regression approach of Hodrick (1992). *, **, and *** indicate statistically different from zero at the 10%, 5%, and 1% level of significance, where p-values are computed using the bootstrap approach described in Section I.B.1.

		$X = \mathrm{Cov}_t(r_{banks}, r_{Tbond})$		$X = \mathrm{Var}_t(r_{Tbond})$		
	(1) $rx^{(5)}$	(2) $rx^{(5)}$	(3) $rx^{(5)}$	(4) $rx^{(5)}$	(5) $rx^{(5)}$	(6) $rx^{(5)}$
Income Gap	-0.55^{**}		-0.68^{**}		-0.59^{**}	-1.04^{*}
	(0.27)		(0.29)		(0.27)	(0.53)
X		-0.01	0.03	-0.03	-0.10	-0.33
		(0.03)	(0.03)	(0.10)	(0.10)	(0.28)
Income Gap $\times X$						2.00
						(1.93)
Constant	0.10^{**}	0.03	0.12^{**}	0.04	0.13^{**}	0.18^{**}
	(0.04)	(0.01)	(0.04)	(0.03)	(0.05)	(0.08)
Observations	106	106	106	106	106	106
Adjusted R^2	0.189	-0.008	0.218	-0.008	0.202	0.211

Another, related, view of banks' objective function is that banks consider risk in a segmented way. In particular, banks may care *specifically* about the volatility created by their interest rate exposure. In this case, banks are reluctant to hold a large *total* quantity of interest rate risk, so that bond risk premia should be proportional to the product of their net quantity of exposed assets and the current variance of long-term bond returns. Table X, column (6) shows that this is not the case empirically. Using the Fama portfolio returns, we construct a quarterly measure of expected variance similar to the measure of covariance described above. The interaction between banks' income gap and bond variance does not significantly forecast bond risk premia: the predictability of bond excess returns arising from banks' income gap is not higher when interest rate risk is larger. This result is not surprising since the variance of bond returns itself does not forecast bond risk premia (Table X, column (4)).[38]

These analyses, together with the evidence of predictability in Section I and the evidence in Section IV.A, constitute an exciting puzzle. On the one hand,

[38] The fact that a strong predictor of returns does not significantly correlate with return volatility and that return volatility does not predict future returns is not unique to our setting. Moreira and Muir (2017), for example, find similar evidence for equities.

banks' balance sheet exposures strongly and robustly forecast bond risk premia (Section I). A natural interpretation of these results relies on banks' risk management—banks try to limit their interest rate exposure and take on significant exposure only when bond risk premia provide appropriate compensation (Section II). This interpretation is consistent with the collection of evidence presented in Section IV.A. On the other hand, this interpretation relies importantly on banks' risk management motive, a motive that remains elusive in the data.

V. Conclusion

While banks are central intermediaries in the market for interest rate risk, they are notably absent in standard empirical analyses of bond risk premia. Our paper fills this gap in the literature. We show that the banking sector's net exposure to interest rate risk, as measured by banks' average income gap, strongly forecasts future bond excess returns. The economic magnitude of this forecastability is significant: a 4.2 percentage points (as a fraction of their total assets) increase in bank holdings of short-term or variable-rate assets is associated with a 231 basis point decrease in the one-year excess returns of five-year bonds. This relationship is stronger for bonds with longer maturity and survives a battery of robustness checks.

A natural interpretation of these findings is that banks are large marginal investors in the market for interest rate risk. In our term structure model, the price of interest rate risk adjusts so that together banks are willing to bear this interest rate risk and banks' holdings forecast bond risk premia. We document a collection of empirical findings consistent with this interpretation. We show that only the average income gap forecasts bond risk premia, not its liability or asset components. We also show that standard measures of liquidity risk do not forecast bond risk premia, in contrast to our measure of banks' balance sheet exposure. Additionally, isolated shocks to the realized net demand and supply of interest rate risk do not bring additional forecasting power to our income gap measure. Finally, we present evidence consistent with interest rate risk-sharing among heterogeneous banks.

However, this interpretation faces a significant challenge. In particular, the banking view of bond risk premia that we highlight in this paper requires that banks suffer when they hold significant balance sheet exposure and interest rates increase. This risk management motive remains elusive in the data. Solving this apparent puzzle is a challenge we hope to tackle in future research.

Initial submission: January 10, 2018; Accepted: March 21, 2019
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

Adrian, Tobias, Richard K. Crump, and Emanuel Moench, 2013, Pricing the term structure with linear regressions, *Journal of Financial Economics* 110, 110–138.

Adrian, Tobias, Erkko Etula, and Tyler Muir, 2014, Financial intermediaries and the cross-section of asset returns, *Journal of Finance* 69, 2557–2596.

Adrian, Tobias, Emanuel Moench, and Hyun Song Shin, 2016, Leverage asset pricing, Discussion Paper DP11466, CEPR.

Adrian, Tobias, and Hyun Song Shin, 2013, Procyclical leverage and value-at-risk, *Review of Financial Studies* 27, 373–403.

Ang, Andrew, and Geert Bekaert, 2006, Stock return predictability: Is it there? *Review of Financial Studies* 20, 651–707.

Ang, Andrew, and Monika Piazzesi, 2003, A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables, *Journal of Monetary Economics* 50, 745–787.

Bai, Jennie, Arvind Krishnamurthy, and Charles-Henri Weymuller, 2018, Measuring liquidity mismatch in the banking sector, *Journal of Finance* 73, 51–93.

Baker, Malcolm, Robin Greenwood, and Jeffrey Wurgler, 2003, The maturity of debt issues and predictable variation in bond returns, *Journal of Financial Economics* 70, 261–291.

Bauer, Michael D., and James D. Hamilton, 2017, Robust bond risk premia, *Review of Financial Studies* 31, 399–448.

Baumol, William J., 1952, The transactions demand for cash: An inventory theoretic approach, *Quarterly Journal of Economics* 66, 545–556.

Beeler, Jason, and John Y. Campbell, 2012, The long-run risks model and aggregate asset prices: An empirical assessment, *Critical Finance Review* 1, 141–182.

Begenau, Juliane, Monika Piazzesi, and Martin Schneider, 2015, Banks' risk exposures, NBER Working Paper 21334.

Berger, Allen N., and Christa H.S. Bouwman, 2009, Bank liquidity creation, *Review of Financial Studies* 22, 3779–3837.

Brunnermeier, Markus K., and Lasse Heje Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.

Brunnermeier, Markus K., and Yuliy Sannikov, 2014, A macroeconomic model with a financial sector, *American Economic Review* 104, 379–421.

Campbell, John Y., and Robert J. Shiller, 1991, Yield spreads and interest rate movements: A bird's eye view, *Review of Economic Studies* 58, 495–514.

Chen, Hui, Michael Michaux, and Nikolai Roussanov, 2013, Houses as ATMs? Mortgage refinancing and macroeconomic uncertainty, NBER Working Papers 19421.

Choi, Dong, and Hyun-Soo Choi, 2016, The effect of monetary policy on bank wholesale funding, Federal Reserve Bank of New York Staff Report.

Cieslak, Anna, and Pavol Povala, 2015, Expected returns in Treasury bonds, *Review of Financial Studies* 28, 2859–2901.

Cochrane, John H., and Monika Piazzesi, 2005, Bond risk premia, *American Economic Review* 95, 138–160.

Cooper, Ilan, and Richard Priestley, 2009, Time-varying risk premiums and the output gap, *Review of Financial Studies* 22, 2801–2833.

D'Amico, Stefania, and Thomas B. King, 2013, Flow and stock effects of large-scale Treasury purchases: Evidence on the importance of local supply, *Journal of Financial Economics* 108, 425–448.

Di Tella, Sebastian, and Pablo Kurlat, 2017, Why are banks exposed to monetary policy? Technical report, National Bureau of Economic Research.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, 2016, The deposits channel of monetary policy, Technical report, National Bureau of Economic Research.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, 2018, Banking on deposits: Maturity transformation without interest rate risk, Technical report, National Bureau of Economic Research.

Duffee, Gregory R., 2011, Information in (and not in) the term structure, *Review of Financial Studies* 24, 2895–2934.

Duffee, Gregory R., 2018, Expected inflation and other determinants of Treasury yields, *Journal of Finance* 73, 2139–2180.

English, William B., Skander J. Van den Heuvel, and Egon Zakrajsek, 2012, Interest rate risk and bank equity valuations, Finance and Economics Discussion Series 2012-26, Board of Governors of the Federal Reserve System (U.S.).

Fama, Eugene F., and Robert R. Bliss, 1987, The information in long-maturity forward rates, *American Economic Review* 77, 680–692.

FDIC, 2005, A changing rate environment challenges bank interest rate risk management, *FDIC Supervisory Insights* 2.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack, 2011, The financial market effects of the Federal Reserve's large-scale asset purchases, *International Journal of Central Banking* 7, 3–43.

Gomez, Matthieu, Augustin Landier, David Sraer, and David Thesmar, Forthcoming, Banks exposure to interest rate risk and the transmission of monetary policy, *Journal of Monetary Economics*. Forthcoming. https://www.sciencedirect.com/science/article/abs/pii/S0304393220300404.

Greenwood, Robin, and Dimitri Vayanos, 2014, Bond supply and excess bond returns, *Review of Financial Studies* 3, 663–713.

Gurkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, The U.S. Treasury yield curve: 1961 to the present, *Journal of Monetary Economics* 54, 2291–2304.

Gürkaynak, Refet S., and Jonathan H. Wright, 2012, Macroeconomics and the term structure, *Journal of Economic Literature* 50, 331–367.

Hamilton, James D., and Jing Cynthia Wu, 2012, The effectiveness of alternative monetary policy tools in a zero lower bound environment, *Journal of Money, Credit and Banking* 44, 3–46.

Hannan, Timothy H., and Allen N. Berger, 1991, The rigidity of prices: Evidence from the banking industry, *American Economic Review* 81, 938–945.

Hanson, Samuel G., 2014, Mortgage convexity, *Journal of Financial Economics* 113, 270–299.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, *American Economic Review* 103, 732–770.

He, Zhiguo, and Arvind Krishnamurthy, 2014, A macroeconomic framework for quantifying systemic risk, Working Paper 19885, National Bureau of Economic Research.

Hodrick, Robert J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–386.

Kashyap, Anil K., and Jeremy C. Stein, 2000, What do a million observations on banks say about the transmission of monetary policy? *American Economic Review* 90, 407–428.

Kirti, Divya, 2020, Why do bank-dependent firms bear interest-rate risk? *Journal of Financial Intermediation* 41, 100823.

Koijen, Ralph S.J., Otto Van Hemert, and Stijn Van Nieuwerburgh, 2009, Mortgage timing, *Journal of Financial Economics* 93, 292–324.

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

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

Ludvigson, Sydney C., and Serena Ng, 2009, Macro factors in bond risk premia, *Review of Financial Studies* 22, 5027–5067.

Malkhozov, Aytek, Philippe Mueller, Andrea Vedolin, and Gyuri Venter, 2016, Mortgage risk and the yield curve, *Review of Financial Studies* 29, 1220–1253.

Mishkin, Frederic, and Stanly Eakins, 2009, Financial Markets and Institutions, 6th edition (Pearson Prentice Hall, New York).

Moreira, Alan, and Tyler Muir, 2017, Volatility-managed portfolios, *Journal of Finance* 72, 1611–1644.

Nelson, Charles R., and Andrew F. Siegel, 1987, Parsimonious modeling of yield curves, *Journal of Business* 60, 473–489.

Neumark, David, and Steven A. Sharpe, 1992, Market structure and the nature of price rigidity: Evidence from the market for consumer deposits, *Quarterly Journal of Economics* 107, 657–680.

Piazzesi, Monika, 2005, Bond yields and the Federal Reserve, *Journal of Political Economy* 113, 311–344.

Purnanandam, Amiyatosh, 2007, Interest rate derivatives at commercial banks: An empirical investigation, *Journal of Monetary Economics* 54, 1769–1808.

Rampini, Adriano A., S. Viswanathan, and Guillaume Vuillemey, Forthcoming, Risk management in financial institutions, *Journal of Finance*.

Schneider, Andrés, 2017, Risk sharing and the term structure of interest rates, Technical report, UCLA.

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

Svensson, Lars E.O., 1994, Estimating and interpreting forward interest rates: Sweden 1992–1994, Technical report, National Bureau of Economic Research.

Swanson, Eric T., 2011, Let's twist again: A high-frequency event-study analysis of operation twist and its implications for QE2, *Brookings Papers on Economic Activity* 2011, 151–208.

Tobin, James, 1956, The interest-elasticity of transactions demand for cash, *Review of Economics and Statistics* 38, 241–247.

Vuillemey, Guillaume, 2019, Bank interest rate risk management, *Management Science* 65, 5933–5956.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication code.