DEMAND AND SUPPLY FOR GOVERNMENT BONDS[†]

Interest Rate Risk in Credit Markets

By Monika Piazzesi and Martin Schneider*

Recent events have stimulated interest in the joint behavior of prices and quantities in credit markets. Data sources such as the Federal Reserve Board's Flow of Funds Accounts (FFA) provide statistics on a rich set of credit market instruments. However, it is challenging to interpret such data using economic models that speak to the allocation of risk across agents, such as households or intermediaries.

On the one hand, an instrument class such as "Treasury bonds" typically contains many different instruments that trade at different prices and have different exposure to interest rate shocks (for example, because of differences in duration). On the other hand, a lot of the price movements in instruments like Treasury bonds and mortgage backed securities are due to common interest rate shocks, making those instruments close substitutes from a portfolio choice perspective.

For understanding how interest rate risk is allocated in the economy, one would thus like to use information on many positions at the same time, rather than, say, focus on one set of instruments only. At the same time, models with many closely substitutable assets are problematic. Instead, it would be desirable to compress position data into simple sets of portfolios, like "long" and "short" bonds, but with some confidence that the risk properties of the original instruments are not lost along the way.

This paper proposes an approach to parsimoniously represent positions in many credit market instruments. We start from the fact, established in the literature on fixed income pricing models, that a small number of factors are sufficient to describe *prices* on a wide variety of these securities. In other words, a pricing model can be used to measure and to parsimoniously represent interest rate risk. It then follows that a small number of bonds are sufficient to describe *quantities* movements on a wide variety of credit market instruments.

We illustrate our approach by applying it to positions of the US household sector. In particular, we use the model to replicate any fixed income security in the FFA with a simple portfolio that contains only three "spanning bonds." These bonds are selected to span various kinds of interest rate risk. While we use the approach to think about interest rate risk and US households' positions, the approach could be useful to thinking about other sources of risk as well as the positions of other sectors (such as financial institutions or foreigners). We thus view the empirical implementation in this paper as a first step.

An important property of our approach is that it can deal with all relevant sources of interest rate risk that have been identified in the term structure literature. For example, at the quarterly frequency we study in this paper, prices are well represented by two factors. Roughly, one factor captures low frequency movements in the level of interest rates, and another factor captures business cycle frequency movements in the spread between long and short bonds. For quantities, our approach uses three bonds so that two portfolio weights capture the exposure to those two factors.[1]

*Piazzesi: Stanford University and NBER. Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94024 (e-mail: piazzesi@stanford.edu); Schneider: Stanford University and NBER, Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94024 (e-mail: schneidr@stanford.edu). We thank Juliane Begenau, Stijn Van Nieuwerburgh, Dimitri Vayanos, and Annette Vissing-Jorgensen for comments.

[1] This is in contrast to other simple approaches of summarizing bond positions, such as Macauley duration, which would capture risk exposure fully only if there is a single factor.

I. Basics

An affine model describes movements in interest rates using an N-dimensional vector of factors f_t, where the first factor is the short term (one period) interest rate $i_t^{(1)}$, so $f_t = (i_t^{(1)}, \ldots)^\top$. The dynamics of the factors are described by a vector-autoregression with Gaussian innovations

$$(1) \quad f_t = \mu + \phi f_{t-1} + \sigma \varepsilon_t, \ \varepsilon_t \sim N(0, I_{N \times N}).$$

The pricing kernel is

$$(2) \quad m_{t+1} = \exp\left(-i_t^{(1)} - \frac{1}{2}\lambda_t^\top \lambda_t - \lambda_t^\top \varepsilon_{t+1}\right),$$

$$\lambda_t = l_0 + l_1 f_t$$

where λ_t is an N-dimensional process that depends on parameters contained in the $N \times 1$ vector l_0 and the $N \times N$ matrix l_1.

Under these assumptions, the time t price of an n-period bond is

$$P_t^{(n)} = \exp(A_n + B_n^\top f_t),$$

with coefficients A_n and B_n that start at $A_0 = 0$ and $B_0 = e_1$ and solve difference equations

$$A_{n+1} = A_n + B_n^\top(\mu - \sigma l_0) + \frac{1}{2}B_n^\top \sigma_f \sigma_f^\top B_n$$

$$B_{n+1}^\top = B_n^\top(\phi - \sigma l_1) - e_1^\top$$

We estimate the parameters $(\mu, \phi, \sigma, l_0, l_1)$ with quarterly postwar data on nominal zero-coupon bond yields. Monika Piazzesi and Martin Schneider (2009) describe in detail the dataset, estimation procedure, and parameter estimates of the two-factor affine model that we use below as an example. For quarterly data, two factors, $N = 2$, are enough; they explain much of the variation in the data. Here, we choose the short (one quarter) interest rate and the spread between the short rate and a long (five-year) rate.

II. Replicating Zero Coupon Bonds

We now select a number N of long bonds (where N is also the number of factors), zero-coupon bonds with maturity greater than one period. Our goal is to construct a portfolio containing the long bonds and the short bond, such that the return on these $N + 1$ bonds replicates closely the return on any *other* zero-coupon bond with maturity n. This replication argument is exact in continuous time. Below, we derive an approximate replicating portfolio in discrete time from a discrete approximation of the continuous time returns of the replicating portfolio.

We start from the approximate change in price of an n-period bond, which is given by

$$(3) \quad P_{t+1}^{(n-1)} - P_t^{(n)}$$

$$\approx P_t^{(n)}\Big(A_{n-1} - A_n$$

$$+ B_{n-1}^\top(f_{t+1} - f_t)$$

$$+ (B_{n-1} - B_n)^\top f_t$$

$$+ \frac{1}{2}B_{n-1}^\top \sigma \sigma^\top B_{n-1}^\top\Big)$$

$$= P_t^{(n)}\Big(A_{n-1} - A_n + B_{n-1}^\top \mu$$

$$+ B_{n-1}^\top(\phi - I)f_t$$

$$+ (B_{n-1} - B_n)^\top f_t$$

$$+ \frac{1}{2}B_{n-1}^\top \sigma \sigma^\top B_{n-1}^\top\Big)$$

$$+ P_t^{(n)}B_{n-1}^\top \sigma \varepsilon_{t+1}$$

$$=: a_t^{(n)} + b_t^{(n)}\sigma\varepsilon_{t+1}.$$

Conditional on date t, we thus view the change in value of the bond as an affine function in the shocks to the factors $\sigma\varepsilon_{t+1}$. Its distribution is described by $N + 1$ time dependent coefficients: the constant $a_t^{(n)}$ and the loadings $b_t^{(n)}$ on the N shocks. In particular, we can calculate coefficients $(a_t^{(1)}, b_t^{(1)})$ for the short bond, and we can arrange coefficients for the N long bonds in a vector $\hat{a}_t$ and a matrix $\hat{b}_t$.

Now consider a replicating portfolio that contains θ_1 units of the short bond and $\hat{\theta}_i$ units of the ith long bond. The change in value of this portfolio is also an affine function in the factor shocks, and we can set it equal to the change in value of any n-period bond:

$$(4) \quad (\theta_1 \ \hat{\theta}') \begin{pmatrix} a_t^{(1)} & b_t^{(1)} \\ \hat{a}_t & \hat{b}_t \end{pmatrix}\begin{pmatrix} 1 \\ \sigma\varepsilon_{t+1} \end{pmatrix}$$

$$= (a_t^{(n)} \ b_t^{(n)})\begin{pmatrix} 1 \\ \sigma\varepsilon_{t+1} \end{pmatrix}.$$

Since the $(N + 1) \times (N + 1)$ matrix of coefficients on the left-hand side is invertible for a nondegenerate term structure model, we can select the portfolio $(\theta_1 \; \hat{\theta}^\top)$ to make the conditional distribution of the value change in the portfolio equal to that of the bond. This approximate replication argument works as in the continuous time case, where we match the drift and diffusion of the n-period bond price with a replicating portfolio. By solving for the portfolio $(\theta_1, \hat{\theta})$ in equation (4), we are matching the conditional mean and volatility of the change in the n-period bond price.

Two-factor example.—When stated in terms of units of bonds $(\theta_1, \hat{\theta})$, the replicating portfolio for the n-period zero coupon bond answers the question: how many spanning bonds are equivalent to one n-period bond? Alternatively, we can define portfolio weights that answer the question: how many dollars' worth of spanning bonds are equivalent to one dollar invested in the n-period bonds? The answer to this question can be computed using the units $(\theta_1, \hat{\theta})$ and the prices of spanning bonds. Figure 1 provides the answer computed from the estimated two-factor term structure model. Since the term structure model is stationary, these weights do not depend on calendar time.

In Figure 1, the maturity of the n-period bond to be replicated is measured along the horizontal axis. The three lines are the portfolio weights $\theta_i P_t^{(i)}/P_t^{(n)}$ on the different spanning bonds i; they sum to one for every maturity n. In our two-factor example, $N + 1 = 3$, three bonds are enough to span interest rate risk. As spanning bonds i, we have selected the one-quarter, two-year and ten-year bonds. For simplicity, we refer to these bonds as "short," "middle," and "long," respectively. In this figure, as well as in the figures that follow, the shading of the lines indicates the maturity of the spanning bond, where lighter grays indicate shorter maturities.

The figure shows that the spanning bonds are replicated exactly by portfolio weights of one on themselves. More generally, the replicating portfolios of most other bonds average their neighboring bonds. For example, most of the bonds with maturities in between the one-quarter and two-year bond are generated by simply mixing these two bonds (although there is also a small short position in the long bond). Similarly, most of the bonds with maturities in between the middle


FIGURE 1. WEIGHTS FOR PORTFOLIOS REPLICATING ZERO-COUPON BONDS

Note: The maturity of the bond that is being replicated is measured along the x-axis.

and long bond are generated by mixing those two bonds. For bonds with maturities that are longer than ten years, the replicating portfolio has a portfolio weight larger than one on the long bond, which represents a leveraged position. The needed leverage is achieved by shorting the middle bond.

Quality of the Approximation.—We want the value of the approximating portfolio to be the same as the value of the zero-coupon bond. This approximation is essentially as good as the term structure model itself. For a replicating portfolio defined by (4), the portfolio value $e^{-i_t}\theta_1 + \sum_i \hat{P}_t^{(i)}\hat{\theta}^{(i)}$ differs from the bond value $P_t^{(n)}$ only to the extent that the term structure model does not fit bonds of maturity n. The additional approximation error introduced by the matching procedure is less than 0.0001 basis points.

III. Replicating Nominal Instruments in the US Economy

We now turn to more complicated fixed income instruments. The FFA provides data on *book value* for many different types of nominal instruments. Matthias Doepke and Schneider (2006; DS) sort these instruments into several broad classes, and then use data on interest rates, maturities, and contract structure to construct, for every asset class and every date t, a certain

FIGURE 2. (top) REPLICATING WEIGHTS FOR TREASURY BONDS; (bottom) REPLICATING WEIGHTS FOR MORTGAGES

Note: The shading of each line indicates the maturity of the bond; lighter grays denote shorter maturities.

net payment stream that the holders of the asset expect to receive in the future. Their procedure makes adjustments for credit risk in instruments such as corporate bonds and mortgages. They use these payment streams to restate FFA positions at *market value* and assess the effect of changes in inflation expectations on wealth.

Here we determine, for every broad asset class, a replicating portfolio that consists of spanning bonds. For every asset class and every date *t*, DS provide a certain payment stream, which we can view as a portfolio of zero-coupon bonds. By applying equation (4) to every zero-coupon bond, and then summing up the resulting replicating portfolios across maturities, we obtain a replicating portfolio for the asset class at date *t*. Figure 2 illustrates replicating portfolios for Treasury bonds and mortgages. The top panel shows how the weights on the spanning bonds in the replicating portfolio for Treasury bonds have changed over the postwar period.

The reduction of government debt after the war went along with a shortening of maturities: the weight on the longest bond declined from over 60 percent in the early 1950s to less than 20 percent in 1980. This development has been somewhat reversed since the early 1980s.[2] The bottom panel

[2] The figure shows only the portfolios corresponding to outstanding Treasury *bonds*, not including bills. DS use data from the CRSP Treasury data base to construct a separate series for bills.

shows that the effective maturity composition of mortgages was very stable before the 1980s, with a high weight on long bonds. The changes that appear since the 1980s are driven by the increased use of adjustable rate mortgages.

We do not show replicating portfolios for Treasury bills, municipal bonds, and corporate bonds, since the portfolio weights exhibit few interesting changes over time. All three instruments are represented by essentially constant portfolios of only two bonds: T-bills correspond to about 80 bonds and 20 percent middle bonds, that for corporate bonds corresponds to about 60 percent middle bonds and 40 percent long bonds, and the replicating portfolio for municipal bonds has 70 percent long bonds and 30 percent middle bonds. The final asset class is a mop-up group of short instruments, which we replicate by a short bond.

A. *Replicating Households' Nominal Asset and Liability Positions*

We measure aggregate household holdings in the FFA at date *t*. To derive their positions in spanning bonds, we compute replicating portfolios for households' nominal asset and liability positions in the FFA. One important issue is how to deal with indirect bond positions, such as bonds held in a defined contribution pension plan or bonds held by a mutual fund, the shares in which are owned by the household sector. Here we make use of the calculations in DS, who consolidate investment intermediaries in the FFA to arrive at effective bond positions. The positions of spanning bonds below thus include these households' indirect holdings.

By applying the replicating portfolios for the broad asset classes to FFA household sector positions, we obtain dollar holdings in spanning bonds as percent of GDP. The holdings of short, middle, and long bonds in US households' asset positions over time series are plotted in Figure 3. It is apparent that the early 1980s brought about dramatic changes in US bond portfolios. Until then, the positions in short bonds had been trending slightly upwards, whereas the positions in long and middle bonds had been declining. This pattern was reversed during the 1980s and early 1990s. The bottom right panel in Figure 3 shows total nominal assets. Their time series behavior is dominated by the positions in longer bonds.
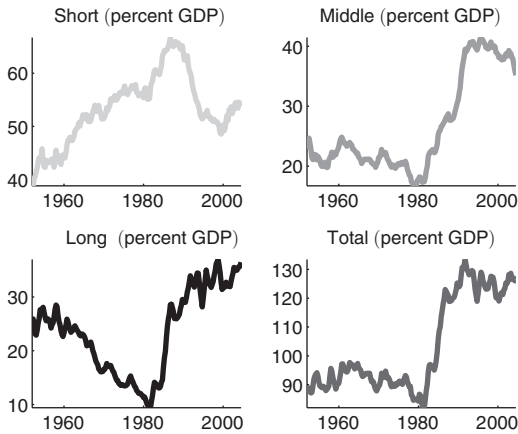
FIGURE 3. US HOUSEHOLD NOMINAL ASSETS
(*as a percent of GDP*)

Notes: Each panel indicates nominal asset positions in the maturity indicated. The panel on the bottom right shows total nominal asset positions.
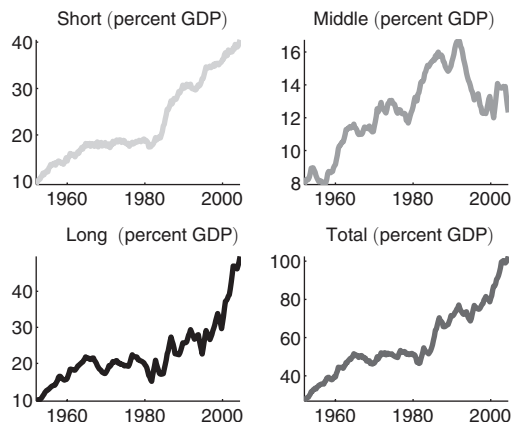

FIGURE 4. US HOUSEHOLD NOMINAL LIABILITIES
(*as a percent of GDP*)

Notes: Each panel indicates nominal liability positions in the maturity indicated. The panel on the bottom right shows total nominal liability positions.

Figure 4 shows nominal liability positions. The bottom right panel shows a dramatic increase in US household debt as a fraction of GDP over time. During the postwar period, the ratio of household debt to GDP has more than tripled. The top and lower left panels indicate that the increase in household debt happened in both short and long term debt instruments. The increase in short term debt reflects a heavier use of consumer loans and adjustable rate mortgages, while the increase in long term debt reflects fixed rate mortgages.

Figure 5 shows the portfolio weights behind the nominal asset and liability positions. The top panel illustrates that the nominal assets increased their weight on short bonds and lowered their weights on middle and long bonds until the 1980s, and then reversed this pattern. The bottom panel for nominal liabilities is dominated by the portfolio weights behind mortgages in Figure 5. Compared to mortgages, the weight on short term bonds is larger because of consumer loans.

IV. Conclusion

The US household sector holds a portfolio consisting of many different fixed income securities. In this paper, we advance an approach to represent this complex portfolio as a much simpler portfolio consisting of a few zero-coupon


FIGURE 5. (top) PORTFOLIO WEIGHTS OF NOMINAL ASSETS.
(bottom) PORTFOLIO WEIGHTS OF NOMINAL LIABILITIES

Note: The shading of each line indicates the maturity of the bond; lighter grays denote shorter maturities.

bonds. To illustrate how the approach works, we use a simple two-factor affine model of the term structure. The approach can also be implemented with more sophisticated models of risk in fixed income security markets. These models may accommodate liquidity risk, special credit risk factors, or changing dynamics.

The approach can also be implemented with models that feature learning about interest rate

risk or, more generally, models that capture subjective beliefs by households (like the ones estimated in Piazzesi and Schneider 2009 based on data on interest rate survey forecasts). The replicating portfolios computed with such models answer how households *perceive* to be bearing interest rate risk. This is an important aspect of understanding households' portfolio choice more generally.

Since factor models work particularly well for fixed income securities, we have focused our attention on nominal positions in this paper. However, our approach also applies to other household positions, such as equity, housing, or claims to labor income (human wealth), which also depend on interest rate risk. To implement the approach with other assets, we need to take a stance on how interest risk affects the prices of these assets. For equity, this can be done using the factor models in the papers by Geert Bekaert and Steven Grenadier (1999), and Martin Lettau and Jessica Wachter (2009), and for human wealth, the model in Hanno Lustig, Stijn Van Nieuwerburgh, and Adrian Verdelhan (2009).

REFERENCES

Bekaert, Geert, and Steven R. Grenadier. 1999. "Stock and Bond Pricing in an Affine Economy." National Bureau of Economic Research Working Paper 7346.

Doepke, Matthias, and Martin Schneider. 2006. "Inflation and the Redistribution of Nominal Wealth." *Journal of Political Economy*, 114(6): 1069–97.

Lettau, Martin, and Jessica A. Wachter. 2009. "The Term Structures of Equity and Interest Rates." Unpublished.

Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan. 2009. "The Wealth-Consumption Ratio." Unpublished.

Piazzesi, Monika, and Martin Schneider. 2009. "Trend and Cycle in Bond Premia." Unpublished.