Comments on “Bond Supply and Excess Bond Returns” by Robin Greenwood and Dimitri Vayanos

John H. Cochrane*

May 12, 2008

1 Theory

The theory in this paper is, to me, an A+. It is simple enough to be transparent, complex enough to capture interesting phenomena, and generates interesting, intuitive, but unexpected results. Segmented markets, due to imperfect arbitrage, and limited by arbitrageur's risk-bearing capacity are a very important topic, but difficult to model. The theory in this paper gives a clear benchmark quantitative parable that I will carry with me as I continue to think about these issues.

Let me present a simplified version of the model. It has two central ingredients:

- Government-investors+private issuers supply bonds of maturity τ inelastically in the amount $\beta(\tau)$.
- Arbitrageurs with an instantaneous mean-variance objective buy and sell bonds across maturities.

With these ingredients, the central result is:

Proposition 2 An increase in the relative supply of long-term bonds

- Raises [all] bond yields, especially for long maturities.
- Raises [all] bond risk premiums, especially for long maturities.

Initially, this is a counterintuitive result. How can a change in the *relative* supply of bonds of different maturities change the yields of *all* bonds? Suppose the net bond supply looks like this, so arbitrageurs have to borrow short and lend long:

*Graduate School of Business, University of Chicago. Comments given at the conference “Beyond Liquidity: Modeling Frictions in Finance and Macroeconomics,” May 9-10 2008.

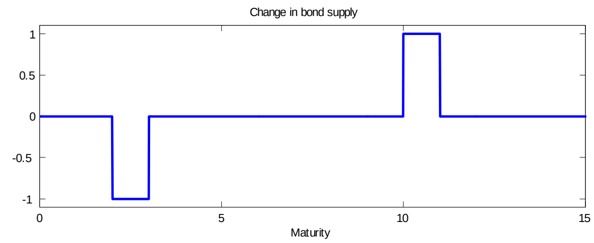

Intuitively, if there are "downward sloping demands" due to "limited risk bearing capacity," we would expect the long yield to rise (price to fall), since arbs are being forced to hold too much of it, and the short rate to fall conversely, something like this:

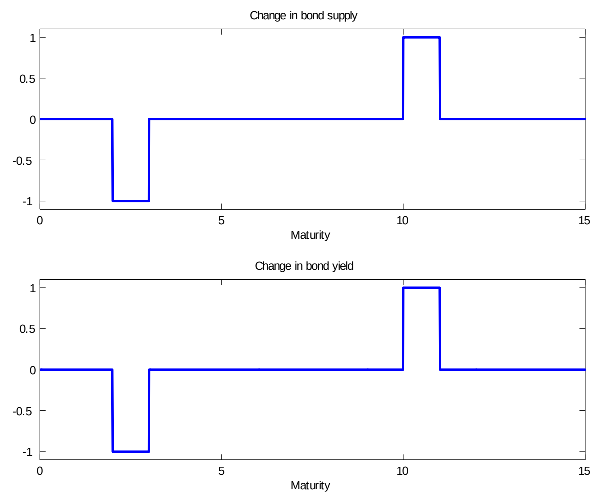

If I require you to eat more apples and fewer oranges, the price of apples will fall and that of oranges will rise – we won't see the price of all fruit rising!

On second thought, 9.9 year bonds are pretty good substitutes for 10 year bonds. Thus, perhaps, we should expect yields to move as in the red line below.

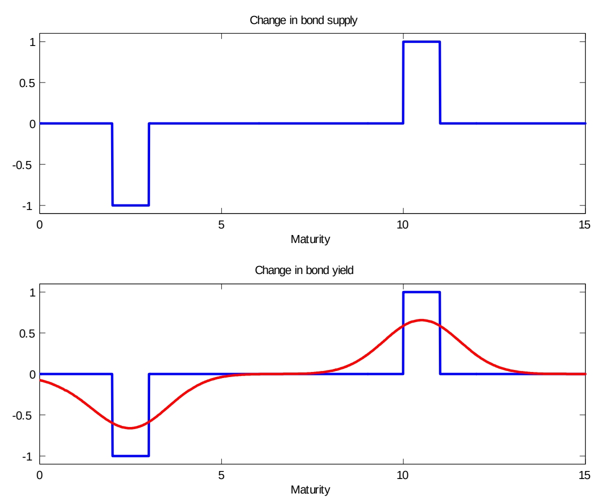

Again, the prediction of the paper in Proposition 2 is that all yields go up! I calculated the response using $\kappa = 0.8$ and the black dashed line plots the resulting bond yields.

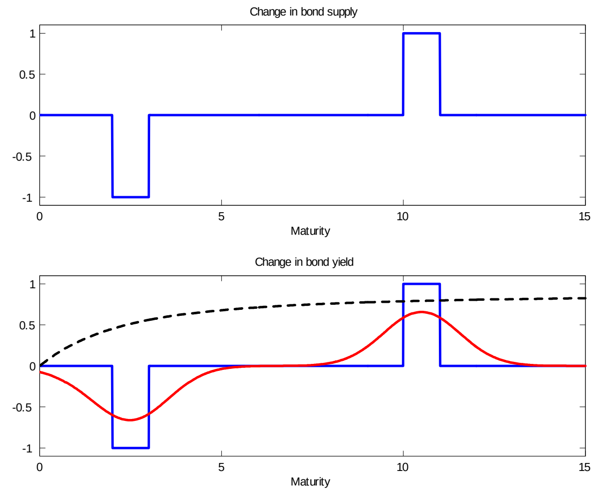

So, why do we get the black line rather than the red line?

I left out a crucial element of the model in my above summary. Actually, the simplified model consists of

- Government-investors supply maturity τ inelastically $\beta(\tau)$.
- Arbitrageurs have instantaneous mean-variance with risk aversion a, and buy and sell bonds across maturities.
- The short rate follows $dr = \kappa(\bar{r} - r)dt + \sigma_r dz$

Now, keep in mind that arbitrageurs are limited by *risk-bearing capacity.* In the paper, the short rate follows AR(1). With this specification, *risks are perfectly correlated across maturities.* The "limit to arbitrage" is only and exactly how much overall exposure to the single interest rate shock the arbitrageur can bear. The position I graphed above (hold long bonds, short short bonds) means bearing positive interest rate risk since long bonds have greater duration than short bonds. Arbitrageurs try to diversify that across the entire term structure, driving up all yields. There are no apples and oranges, there is only fruit; oranges being fruitier, if I require you to eat fewer apples and more oranges, the price of all fruit declines.

My mistaken intuition in the first case implicitly assumed risks were independent across maturities, and in the second case implicitly assumed that bond returns' correlation declines as a function of maturity. The key to proposition 2 is the specification of a one-factor model in which all bond returns are perfectly correlated.

This is a very important result to my thinking. In lots of places we think we see somewhat segmented markets, connected by arbitrageurs or other institutions who are limited by their capacity to bear risk. I learn that the *correlation* between the risks they are trying to connect is absolutely crucial, and often ignored, in "downward sloping demand" thinking. Demand curves in all of economics come from the presence of substitutes, and substitutes in finance depend on the correlation of returns. Of course the model is extreme: multifactor models, or models that allow variation in liquidity and other spreads will introduce other kinds of risks, and bonds will not be perfect substitutes any more. But the extreme nature of the model makes the central point starkly clear.

The central analytical result in the simplified version of the model I have sketched is:

- Equilibrium bond yields follow a Vasicek model with market price of risk

$$\lambda = \frac{a}{\kappa}\int_0^T \beta(\tau)\left(1-e^{-\kappa\tau}\right)d\tau$$

where the "market price of risk" λ means

$$E_t dR - rdt = \lambda cov(dR, dr)$$

You can see in this formula that adding one more long-maturity β, with one less short maturity β, raises λ, and hence raises yields and expected returns across the term structure, since the factor $1-e^{-\kappa\tau}$ downweights the contribution of long-term bonds.

Similarly, you can see that there *are* maturity changes that have *no* effects on the yield curve. That is a nice case for reinforcing the message of the paper. If the government changes the supply of different maturities $\beta(\tau)$ in such a way that $\int_0^T \beta(\tau)\left(1-e^{-\kappa\tau}\right)d\tau$ is unaffected, there is no change in yields. What about "downward sloping demands?" Again, the bonds are perfectly correlated, so in this case the arbitrageurs bear no overall risk.

The actual model in the paper adds an elastic supply curve. Letting $s^{(\tau)}$ denote the net supply at maturity τ, the assumption is

$$s^{(\tau)} = \beta(\tau) - \alpha(\tau)y^{(\tau)}$$

This is obviously very artificial, especially given the above discussion. This is nothing at all like the "supply curve" of the arbitrageurs, since the supply of 9.000 year bonds has nothing at all to do with the yields of 9.001 year bonds. Judge it not for realism however; this assumption delivers a Vasicek-like model with a time-varying risk premium that is also a function of the short rate. That is a beautiful result and a nice extension.

Incidentally, in this extended model, a period of low interest rates such as the current one is essentially just a subsidy to the "arbitrageurs" – let's call them "banks" – with the special ability to borrow short and lend long. It's an interesting and provocative view of monetary policy and bond risk premiums.

2 Quibbles

Part of a discussant's job is to complain, which I will do quickly. Taking a beautiful theoretical model, simplified down to show one particular and striking effect, to actual historical analysis is difficult. It's all too easy to ask for more – I don't have to do the hard work. Still, It's worth stating the wish list.

1. The paper describes a maturity-lengthening exercise as follows: "An increase in the relative supply of long term bonds corresponds to a decrease of $\beta(\tau)$ for small τ and increase in $\beta(\tau)$ for large τ holding $\int_0^T \beta(\tau)d\tau$ constant" But in such an operation, the government must keep constant the *market value* of debt $\int_0^T P(\tau)\beta(\tau)d\tau$ not the *face value* $\int_0^T \beta(\tau)d\tau$, since it can only sell long-term bonds for their market values. However, this consideration means even larger face values of long-term debt, so I believe it reinforces rather than dulls the paper's conclusions.

2. Changes in bond supply are unexpected, probability zero, one-time-only, comparative statics exercises. What if If agents know such a change in supply might come? It would introduce a second "factor", a second source of risk in holding bonds. Without considering that effect, we can't really compare comparative statics to time series.

3. Private supply? States and local governments? Foreign supply? Ricardo? The assumption that bond net supply is either completely inelastic, or a mechanical function of one yield only, is obviously artificial, especially when we consider changes in supply that last decades. More supply of long term bonds means more taxes to be paid in on the maturity date, so people should offset by demanding more bonds of the same maturity. (This paper implicitly assumes a very non-Ricardian fiscal regime, something I am sympathetic to, but other readers may not appreciate as much.) Ultimately, once Modigliani and Miller are done, isn't the supply of various maturities related to the durability of the capital stock?

4. I have to echo the usual discussant whining for generality, without losing clarity and tractability. In particular, the assumption that bond arbitrageurs have only instantaneous mean-variance objectives seems strained. What if they are long-lived? What if we have a multifactor model for the short rate – do we not get something like the red waves in the effects of bond supply shifts? The paper's assumption that risk aversion rises after losses comes from thin air – can't that be endogenized with long-lived traders?

5. Why *does* the government choose the maturity structure of the debt? This is an honest question. There is a small (but nonzero) literature asking this question both theoretically and empirically. (Long term debt has the advantage that a Lucas-Stokey state-contingent default requires less inflation.) In the model, the short rate is infinitely elastic, while all other maturities are inelastic.

6. Measurement. Shouldn't we include state, local and agency bonds? Corporate bonds? The monetary base is actually a significant part of Federal debt in the 1970s, and might be included in "short term debt."

7. Dimitri and Robin go to all the trouble to assemble the MBX database, but then use only rough measures (fraction of debt over 10 years, average maturity) as right hand variables. Why? The model predicts the premium exactly. In my simple example, we have, as above, $\lambda = \frac{a}{\kappa} \int_0^T \beta(\tau)\left(1 - e^{-\kappa\tau}\right) d\tau$. The mbx data give $\beta(\tau)$. Why not compute λ exactly?

3 Empirical work

The paper focuses on regressions. I am less convinced by these, but I also learned an important econometric lesson, which bears on return-forecasting regressions more widely

3.1 Yield spreads

The first question is, *are* yield spreads larger when there is more long term debt? The paper says yes, and

> Figure 3 ... shows a strong positive correlation

Well, here's Figure 3. We clearly have different ideas of what "strong" means!

Figure 3
Term spread and government debt maturity

Scatterplot of term spread against debt maturity. Debt maturity is measured as the fraction of principal and coupon payments due in ten years or longer. The term spread is the difference in log yield between the 20-year bond and the one-year bond, and is measured in December each year. Yield on the 20-year bond is from Ibbotson; yield on the one-year bond is from Fama and Bliss.

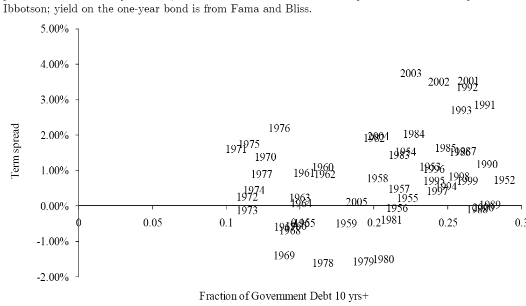

In the plot, the major evidence is in the data points 1991, 1993 and 2001-2003 – without these data points the slope is completely absent. A time series plot is a better tool for understanding this correlation.

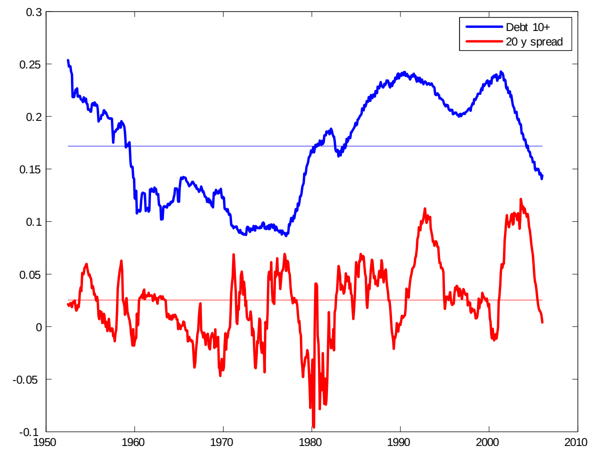

Here again, you can see that most of the evidence for correlation is in the last three wiggles -essentially three data points.

But what about Table II, which presents a regression of these data?

$$y_t^{(20)} - y_t^{(1)} = a + 0.077(t = 3.677)\frac{D_t^{10+}}{D_t} + \varepsilon_t;\ R^2 = 0.097$$

The t statistic of 3.677 seems awfully good. A hint comes in the footnote: "t statistics.. follow Newey-West (1987) allowing up to [?] six lags." Yes, the residuals of such a regression are serially correlated, and we have to make some correction. To investigate, here is a plot of the term spread (red), the fitted value of the regression (blue) and the residuals (black)

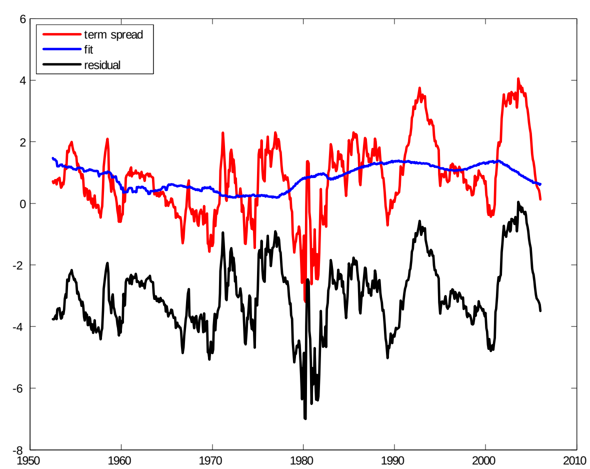

You can see in the blue line that much of the regression evidence comes from the very low-frequency difference between 70s and 90s – counting peaks and troughs there really are about three data points. More to the current point, it's clear that these residuals have autocorrelation that extends much further than a 6 month moving average that Newey -West might hope to pick up.

To make an alternative correlation, I use the parameteric alternative to the "non-parametric" Newey-West. If the residuals follow an AR(1) with correlation ρ, and calculating the required $\sum_{j=-\infty}^{\infty} E\left(\varepsilon_t \varepsilon_{t-j}\right)$ term by $\rho^j \sigma_\varepsilon^2$, then you correct t statistics with a factor $\sqrt{\frac{1+\rho}{1-\rho}}$. I estimate $\hat{\rho} = 0.9543$ giving an actual t-statistic of $t = 1.399$. (Even this calculation is optimistic, since the estimated autocorrelation coefficient is biased down.) "Nonparametric" corrections, based on asymptotic distribution theory and lots of promises that one would increase window size with samples, can be very misleading when there is a lot of serial correlation to correct for.

Of course the standard discussant comments also apply: we're regressing price on quantity, and interepreting the result as a demand curve. That requires us to assume that demand is rock-stable, over 50 years, and all variation comes from exogenous supply shifters. A "slow adjustment" will give "preferred habitats" that last a few years, but not 50.

3.2 Return forecasts and a new econometric pitfall

The paper also asks whether expected returns are higher when there is more long-term debt. Again, the paper says

> Figure 4 shows a strong positive correlation [p. 17]

Again, looking at Figure 4, we have different ideas of what "strong" means

Figure 4
Bond returns and government debt maturity

Scatterplot of excess bond returns against debt maturity. Debt maturity is measured as the fraction of principal and coupon payments due in ten years or longer. In Panel A, excess bond returns are the difference between the one-year log return of the 20-year bond and the log yield of the one-year bond. In Panel B, excess bond returns are the three-year cumulation of one-year excess returns. Returns are measured in December each year.

Panel A. One-year-ahead excess bond returns

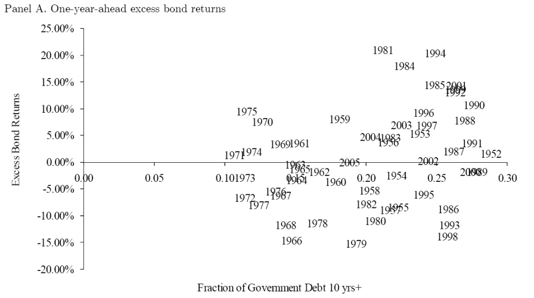

Some of these regressions seem to give even stronger evidence for the effect of debt maturity.

This case reveals to me an underappreciated econometric issue that may cause trouble for return forecasting regressions more broadly. This is a one-factor model, so yield spreads and expected returns move one for one. Actual returns are expected returns plus noise. Thus, compared to the yield regressions, *under the null of the model,* all we have done is to add noise to the left hand variable. How can adding noise to the left hand variable give more precise estimates? Again, correcting for serial correlation of residuals is the key.

Here is a graphical illustration of the situation. (I simulate $x_t = 0.995x_{t-1} + v_t$, $z_t = 0.995z_{t-1} + w_t$, and then $y_t = x_t + z_t$. $y + \varepsilon$ adds an iid error. All shocks are iid normal.)

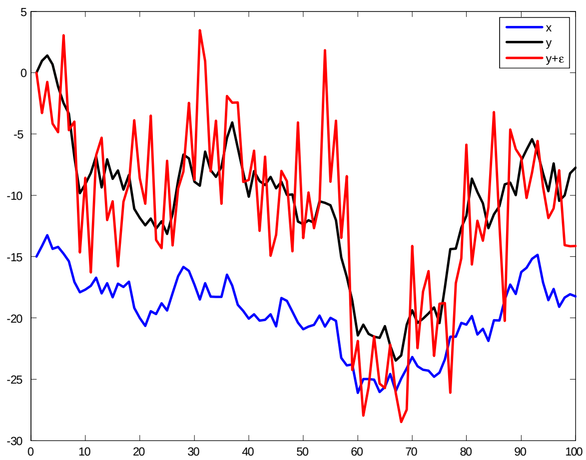

The red $y + \varepsilon$ line is the black y line, with the addition of iid error ε. The regression of y on x leaves strongly serially correlated residuals, as in the regression of yield spreads on debt. When we add noise ε, however, the residuals become less serially correlated. In return-forecasting situations, the ε shocks are much larger than the time-varying expected return y movement, so this effect is much larger.

Examining the properties of the residual in this case, you might conclude that *no* correction for residuals was needed, since the residuals are so close to white noise! If you adopt my AR(1) parametric correction strategy, you will estimate a much lower parameter ρ, and again potentially miss the correlation of residuals. In this example, the residuals are a very long lasting AR(1) (I used a coefficient 0.995 to produce the plot) plus white noise – an ARMA(1,1). Any feasible "nonparametric" correction such as Newey-West would surely miss the long run correlation of residuals. "Parametric" approaches such as my AR(1) are better, far less used, but would also fail in this case. Nobody has ever (at least to my knowledge) fit the required ARMA(1,1) to the residuals in a return-foreasting regression.

This problem is likely to be important in a wide variety of return-forecasting situations, such as forecasting stock returns with dividend yields. In the typical case, we think that expected returns vary slowly over time – like my "y" line in the above figure. We examine a forecasting variable that is also slow-moving, and correlated with expected returns, but is not perfectly correlated with true expected returns. (Part of the residual consists of additional, slow-moving variables left out of the forecasting regression). The residuals in a regression of ex-post returns on x then contain a slow-moving component as well as the serially uncorrelated unexpected return component.

We conventionally compute standard errors and t statistics "under the null" which means assuming returns are not forecastable. In this case the residuals are serially uncorrelated, so ignoring serial correlation is correct for a precise, classical test of the null. The trouble occurs here when we want to use standard errors for measurement rather than classical hypothesis testing, or when "the null" is a model (as in this paper) that does not predict constant expected returns. (We conventionally use serial correlation corrections only to handle overlapping data. The serial correlation corrections we're talking about here go past the overlap horizon.)

In sum, I learned from both regressions that there is a rather gaping hole in standard econometric practice. Asymptotic "nonparametric" corrections don't work, simple models (AR(1)) are potentially misleading for return regressions, since we expect a mixture of a long-lasting and white noise component, and my back-of-the-envelope procedure is clearly also subject to small sample problems. I hope to hear from time-series econometricians of a good solution for this problem (or to be educated where it already exists and I don't know about it.)

3.3 Regression horse races and other evidence

The return-forecasting regressions also include competitions with other variables, and in particular the "return forecasting factor $\gamma' f$ that Monika Piazzesi and I found summarizes the information about expected returns in all yields. The paper claims that bond maturity and $\gamma' f$ each survive in the presence of the other, indicating that there are two dimensions to bond expected returns, which we might think of as a "business cycle" premium captured by $\gamma' f$ and a longer-term effect of supply captured by, say, the fraction of debt over 10 years' maturity.

Multiple regression forecasts of this sort need to be evaluated carefully. Monika and I don't (yet!) have an economic explanation for $\gamma' f$, so we'd be just as happy if bond supply drove out $\gamma' f$ in a multiple regression. Now we would have an economic explanation for $\gamma' f$.

Conversely, even if the model is right, we should expect yields to drive out supply variables. This is an instance of the a general Catch-22 in all "macro" models of the yield curve. If you start with a model in which a macro variable x_t is a state variable for yields, you find that bond prices and yields are functions of x_t. But then x_t is revealed by bond yields etc., so x_t can't add any predictive power to observed bond yields! In fact, since yields are better measured, bond prices should drive the underlying macro factors x_t out in any term structure exercise. More generally still, it is a mistake to evaluate "macro models" of asset prices by horse races with "finance models" that use prices as state variables. The latter will always win, even if the former are correct. Macro models explain finance facts, they don't supplant them.

In sum, in this view, it is a mystery that yields do *not* drive out bond supply! (Or, maybe, we should take my econometric doubts above as an indication that they do drive out bond supply, and the catch-22 as an interpretation that this is good news, not bad news.)

On the other hand, there are exceptions to the general Catch-22. Some models or parameter configurations lead to "unspanned state variables" (such as Pierre Collin-Dufrensne's "unspanned stochastic volatility") in which underlying state variables are *not* revealed by bond yields. This may be one of them. If we take seriously the prediction that the market price of risk at time t is $\lambda_t = \frac{a}{\kappa}\int_0^T \beta_t(\tau)\left(1 - e^{-\kappa\tau}\right) d\tau$, (or the more complex version in the model with supply) then I think we have a Vasicek model with time-varying coefficients. In that case there is no stable linear relation between yields and expected returns, and calculation of λ_t can in fact improve forecasts. I'm guessing, but this prediction seems worth working out.

The paper and presentation also discuss specific "experiments" including "operation twist", the removal of coupon caps on long-term bonds, and the treasury repurchase program. There are others, such as the change in UK regulations that forced pension funds to dedicate more closely, and thus hold more long-term nominal debt. These have the advantage of potentially isolating exogenous supply changes, and since the time-frame of the experiences are a few years rather than decades, it is more plausible that investors don't respond, arbitrageurs do not flow into the business attracted by persistent profits, and foreign and corporate supplies do not respond. I might emphasize them over regressions.

4 Where to go: theory to follow facts

In sum, I think the question is fascinating and relevant. Imperfect arbitrage limited by risk-bearing capacity leads to supply-induced variation in yields and risk premia. But I think the "fact," illustrated in my graph of decade-long variation in yields and risk premiums occasioned by decade-long variation in maturity structure, is tenuous. (As above, investigation of particular, shorter-lived "experiments" may be more fruitful.)

What else can we do? Perhaps, rather than search for new facts suggested by the simplest form of the model, we should take the general idea of the model and adapt it to follow better-known facts. There is nothing wrong with theories reverse-engineered to explain known and puzzling facts – most progress is made this way.

Here are a few bond-market facts that come to my mind immediately, that this sort of theory seems well adapted for:

1. The benchmark effect in Japanese Government bonds. (Boudoukh and Whitelaw.) The "actively traded" bond has a lower yield than bonds of almost exactly the same maturity. When markets change the actively traded benchmark, the yield spread follows.

2. The similar on-the-run/off-the-run liquidity spread in US government bonds (LTCM).

3. "Conundrum." Perhaps (despite my own cynical comments in other forums) demand by agents such as UK pension funds really does drive surprisingly low

long-term rates in the past few years.

4. Refcorp. Francis Longstaff shows a substantial spread between Refcorp bonds and equivalent treasuries – although Refcorp bonds are guaranteed by the full faith and credit of the US treasury! On a cashflow basis, they are equivalent to treasuries in all respects. This picture from Longstaff's paper, showing the spread between Refcorp and equivalent treasuries, summarizes the result nicely.

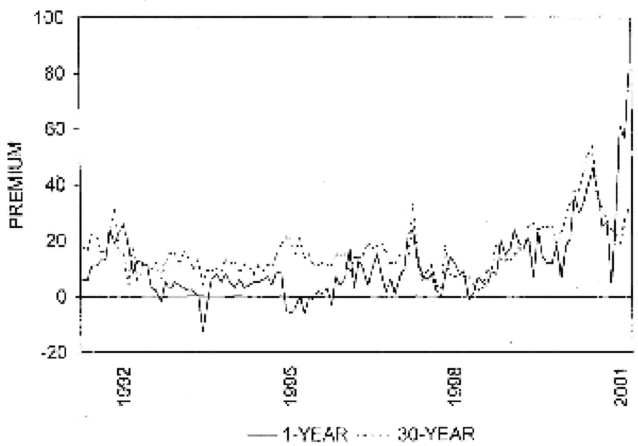

5. Supply and *spreads.* Arvind Krishnamurthy and Annette Vissing-Jorgensen show that the overall supply of government debt correlates with corporate-treasury *spreads* at the same maturity. I interpret their story as a "money demand" or "convenience yield" for treasury debt.

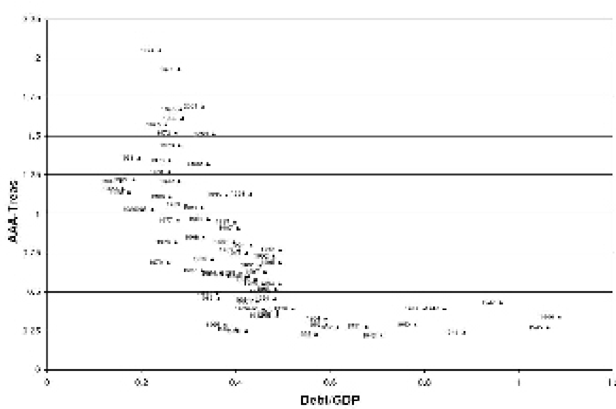

Now, most of these observations are demands for specific *bonds* – specific issues, or issues with specific characteristics unrelated to cashflow – rather than demands for specific *maturities* (except maybe the conundrum) as modeled in this paper. They are related to "liquidity" or "transactions" nature of bonds rather than their maturity. And the risk that arbitrageurs face is the risk that "liquidity" spreads will widen, not risks in any of the conventional term structure factors (level, slope, and curvature).

Still, it seems to me that this gorgeous theory can address these kinds of facts that are also suggestive of "limited arbitrage." We will need to specify characteristics that generate somewhat "preferred habitats" for investors, and we will predict that "liquidity factors" common across lots of bonds (and other assets!), factors that force arbitrageurs with well-diversified portfolios to take on undiversifiable risks, will drive the resulting premiums. (Of course, I am to some extent preaching to the choir here, as Dimitri and Robin have and continue to attack these phenonmena in other papers.)

Theory

Simplify to fixed supply ($\alpha(\tau) = 0$)

net supply $\beta(\tau)$

$$dr_t = \kappa(\bar{r} - r_t)dt + \sigma dB_t$$

(Implicitly, net supply infinite at this price)

Arbs

$$\max E_t dW_t - \frac{a}{2} Var_t dW_t$$

Proposition 1

$$p_t^{(\tau)} = -A(\tau)r_t - C(\tau)$$

$$A(\tau) = \frac{1 - e^{-\kappa\tau}}{\kappa}$$

$$C(\tau) = \kappa\bar{r}^* \int_0^\tau A(u)du - \frac{\sigma_r^2}{2} \int_0^\tau A(u)^2 du$$

$$\begin{aligned}
C(\tau) &= \bar{r}^* \int_0^\tau \left(1 - e^{-\kappa u}\right) du - \frac{\sigma^2}{2\kappa^2} \int_0^\tau \left(1 - e^{-\kappa u}\right)^2 du \\
C(\tau) &= \bar{r}^* \int_0^\tau \left(1 - e^{-\kappa u}\right) du - \frac{\sigma^2}{2\kappa^2} \int_0^\tau \left(1 - 2e^{-\kappa u} + e^{-2\kappa u}\right) du \\
C(\tau) &= \bar{r}^* \left[\tau - \frac{1 - e^{-\kappa\tau}}{\kappa}\right] - \frac{\sigma^2}{2\kappa^2} \left(\tau - 2\frac{1 - e^{-\kappa\tau}}{\kappa} + \frac{1 - e^{-2\kappa\tau}}{2\kappa}\right) \\
\left(\frac{1 - e^{-\kappa\tau}}{\kappa}\right)^2 &= \frac{1 - 2e^{-\kappa\tau} + e^{-2\kappa\tau}}{\kappa^2} \\
C(\tau) &= \bar{r}^* \left[\tau - \frac{1 - e^{-\kappa\tau}}{\kappa}\right] - \frac{\sigma^2}{2\kappa^2} \left(\tau + \frac{1 - e^{-2\kappa\tau} - 4 + 4e^{-\kappa\tau}}{2\kappa}\right) \\
C(\tau) &= \bar{r}^* \left[\tau - \frac{1 - e^{-\kappa\tau}}{\kappa}\right] - \frac{\sigma^2}{2\kappa^2} \left(\tau - \frac{e^{-2\kappa\tau} - 2e^{-\kappa\tau} + 1 + 2 - 2e^{-\kappa\tau}}{2\kappa}\right) \\
C(\tau) &= \bar{r}^* \left[\tau - \frac{1 - e^{-\kappa\tau}}{\kappa}\right] - \frac{\sigma^2}{2\kappa^2} \left(\tau - \frac{\kappa}{2}\left(\frac{1 - e^{-\kappa\tau}}{\kappa}\right)^2 - \frac{1 - e^{-\kappa\tau}}{\kappa}\right) \\
C(\tau) &= \left(\bar{r}^* - \frac{\sigma^2}{2\kappa^2}\right) \left[\tau - \left(\frac{1 - e^{-\kappa\tau}}{\kappa}\right)\right] + \frac{\sigma^2}{4\kappa} \left(\frac{1 - e^{-\kappa\tau}}{\kappa}\right)^2
\end{aligned}$$

$$\begin{aligned}
\bar{r}^* &= \bar{r} + \frac{a\sigma_r^2 \int_0^T \beta(\tau)A(\tau)d\tau}{\kappa} \\
\bar{r}^* &= \bar{r} + \frac{a\sigma_r^2}{\kappa^2} \int_0^T \beta(\tau)\left(1 - e^{-\kappa\tau}\right) d\tau
\end{aligned}$$

Standard vasicek

$$\frac{d\Lambda}{\Lambda} = -rdt - \sigma_\Lambda dz$$
$$dr = \kappa(\bar{r} - r)dt + \sigma_r dz$$

$$\begin{aligned} p(\tau, r) &= -C(\tau) - A(\tau)r \\ A(\tau) &= \frac{1}{\kappa}(1 - e^{-\kappa\tau}) \\ C(\tau) &= \frac{\sigma_r^2}{4\kappa}A(\tau)^2 + \left(\bar{r} - \frac{\sigma_r\sigma_\Lambda}{\kappa} - \frac{\sigma_r^2}{2\kappa^2}\right)[\tau - A(\tau)] \end{aligned}$$

Thus, this is the standard vasiceck model, with a distortion to the riskfree rate

$$\bar{r}^* = \bar{r} + \frac{a\sigma_r^2}{\kappa^2}\int_0^T \beta(\tau)\left(1 - e^{-\kappa\tau}\right)d\tau$$

This means it adds a risk premium,

$$\frac{\sigma_r\sigma_\Lambda}{\kappa} = -\frac{a\sigma_r^2}{\kappa^2}\int_0^T \beta(\tau)\left(1 - e^{-\kappa\tau}\right)d\tau$$

$$\sigma_\Lambda = -\frac{a\sigma_r}{\kappa}\int_0^T \beta(\tau)\left(1 - e^{-\kappa\tau}\right)d\tau$$

EG

$$E_t dR - rdt = -E_t\left(\frac{d\Lambda}{\Lambda}\frac{dP}{P}\right) = -E_t\left(-\sigma_\Lambda dz \times \frac{de^{-A(\tau)r - C(\tau)}}{P}\right) = -E_t\left(-\sigma_\Lambda dz \times -A(\tau)\sigma_r dz\right) = -E_t\left(\sigma\right.$$

$$E_t dR - rdt = -\sigma_\Lambda\sigma_r A(\tau)$$

In the usual notation, $\lambda = -\sigma_\Lambda\sigma_r$,

$$\lambda = \frac{a}{\kappa}\int_0^T \beta(\tau)\left(1 - e^{-\kappa\tau}\right)d\tau$$

Quibbles:

1)

Proposition 2: "an increase in the relative supply of long term bonds corresponds to a decrease of $\beta(\tau)$ for small τ and increase in $\beta(\tau)$ for large τ holding $\int_0^T \beta(\tau)d\tau$ constant"

Instead, what if we hold $\int_0^T P(\tau)\beta(\tau)d\tau =$ market value constant $= \int_0^T e^{-C(\tau)-A(\tau)r}\beta(\tau)d\tau$

2) unexpected. If agents know it's coming?