Logistic
Research & Trading Inc.

A Parsimonious Model of Treasury Futures

Joel C. Gibbons Stuart School of Business Illinois Institute of Technology Chicago, Illinois, and Logistic Research & Trading Co. Elmhurst, Illinois, 60126

Preliminary Draft, NOT FOR CITATION OR DISTRIBUTION WITHOUT PERMISSION

November 4, 1996

Abstract

We develop and estimate a model of the Treasury futures contracts in relation to the parameters of the Nelson - Siegel model of the Treasury yield curve, and test the hypothesis that this model of the curve identifies the true, unobservable default-free discount function. We find that the four parameters of the NS model explain most of the variance in contract prices within the sample, which runs from 1982 to the August, 1996. But we also find compelling evidence against the null hypothesis. Specifically, we find that even holding the NS curve constant, the Treasury contracts are highly correlated with yields of individual Treasury notes and bonds.

If we take the social discount function – the function which assigns to every future date the present value of a future dollar – as a starting point for bond valuation, then pricing of call free Treasury bonds and notes is a fairly straightforward application

of the theory. As long as we accept that all the promised cash flows will materialize as promised, the value of the bond is easily computed from the discount function. In theory, the prices of Treasury Strips should exactly trace out the discount function. This truth has enormous practical application because of the immense size of the outstanding Treasury debt, but it does not extend beyond Treasuries to any other security. To model and value any other security requires a model of how it is priced *relative to* Treasuries and to the Treasury yield curve.

The starting point, traditionally, for this type of analysis has been to spread the security off the yield curve and then to attempt to understand the dynamics of the spread. If the *promised* cash flows are certain, as in the case of a call free corporate bond, then the credibility of the promise to pay is the only possible reason why there is any spread. The simple spread calculation is a reasonably good way to quantify the extent of discounting for credit risk. Most so-called "fixed income" securities today, however, offer not a single pre-defined stream of cash flows, but promise instead some bundle of contingent claims. It is strictly speaking impossible to spread such securities off the yield curve, because there are no known cash flows to price. The approach which has been taken to date is a simple, heroic effort to bootstrap the conventional spread methodology. On the basis of some space-age mathematics, we replace the whole confusing distribution of future cash flows with one *certainty equivalent* stream. The pure spread of the security is, then, simply the conventional spread of the certainty equivalent stream. The full model of the asset, however, also includes the analytics which produced that stream from the actual definition of the contingent claims. I will not stop here to dwell on the shortcomings of these models, except to note that they all need a model of how the yield curve will evolve in the future, and the historical record is not at all reassuring about our ability to do that.

There is another approach which starts from an entirely different view of the problem which we are trying to solve. Everyone would agree that the task at hand is not to compute a yield spread from a fixed income security; yield spreads were only a means, not an end. The job is to understand how the value of the security correlates with the Treasury yield curve. The reason why we have not been able to take a more direct approach is that the yield curve has simply been too complicated to deal with. In theory, the curve is a continuum of numbers. We simply don't know how to correlate it with anything; we don't know how to correlate functions to each other. We need a concise parameterization of the curve which both captures the essential features of every observed curve and translates them into a small set of parameters which evolve over time.

Charles Nelson and Andrew Siegel [Nelson and Siegel] have offered one such parameterization which gets by with only four parameters: the Level, Slope, and Curvature of the Treasury spot curve, and a fourth parameter called Tau. The idea behind Tau is to allow the curve to have to distinct regimes, one of which applies out to a point in the proximate future – e.g., three years – and the other which applies from there to infinity. Tau is the date at which the transition occurs from one regime to the other. Thus Tau is measured in years, and typically assumes values between two and seven. The idea of a parameterization of the yield curve is that there is only one possible curve – I mean,

one possible actual curve in reality – which corresponds to any given foursome of parameters [Level, Slope, Curvature, Tau].

The model is invalidated to the extent that we actually observe different yield curves which correspond to a single set of these four parameters, and it is validated to the extent that whatever are the level, slope, curvature, and tau of the curve on a given day, the rest of the curve exactly fits the corresponding model curve. There are sure to be days when the yield curve has the same model representation, but that not all yields are exactly equal, at all maturities. It remains an open question, however, whether the difference between the curves should or would imply significant differences in the pricing of fixed income securities. It this article, we will look at one particular pricing problem – pricing the Treasury contracts which trade on the Chicago Board of Trade – to attempt empirically to determine how good a job the model does. The essence of the test is to estimate how much of the variance of prices, in this cases, prices of Treasury contracts, is explained by the four factors identified by the curve model, and to check whether the actual correlation works in a plausible fashion.

This is not the simplest or most direct test of the Nelson - Siegel model. The most direct test is to see how well it accounts for the prices of Treasury bonds and notes. Tests of that kind have been done and reported elsewhere. While simple and direct, however, this test is not necessarily the most revealing one, for the following reason. Treasury securities are priced on the social discount function, from which it follows that the actual Treasury yield curve at any point in time is a reasonable pricing benchmark for all fixed income securities. It remains an open question, however, if there is a model curve of some sort which actually better represents that social discount function than does a given observed yield curve. Here "better" has two distinct meanings: better in the sense that the modeled curve is a better prediction of future yields, and better in the sense that it better explains the whole universe of bond prices at the present time.

Review of the Nelson - Siegel Model.

The Nelson -Siegel Model (henceforth NS) starts from the premise that the forward curve is a solution of a second order differential equation which has constant coefficients, i.e.

1.
$$f(t) + b * f(t) + c * f(t) = D$$
,

where f(t) is the instantaneous forward rate at the maturity point to. One rationale for such a model is that yields are determined by two pieces of information: an estimate of the rate at which inflation will accelerate in the near term, and an attenuation rate at which this timely forecast reverts to a slow-moving long term forecast. In the general case, the solution of this equation for any given initial level of the curve is given by

2.
$$f(t) = L + S_1 * \exp(-beta_1 * t) + S_2 * \exp(-beta_2 * t)$$
.

The f's and betas are related to the parameters of the differential equation by

3.
$$beta_1 + beta_2 = -b, and$$
$$beta_1 * beta_2 = c.$$

The coefficients L, S_1 , and S_2 are determined by the initial conditions: f(0) and the initial growth rate, df(0)/dt. One of the conditions is easy to state: the overnight rate f(0) is equal to the sum, $L + S_1 + S_2$. The initial slope and curvature of the forward curve are also simple linear combinations of L, S_1 , and S_2 , from which we can solve for the L and the S's as functions of the initial conditions.

In general, the betas can be real or complex numbers, but if they are complex, then they are complex conjugate to each other. We can dismiss this possibility at the outset, because complex betas give rise a situation in which the forward curve oscillates within a fixed range forever. Since that kind of curve is not observed in reality, we can safely assume that if this model is to be of any value, the betas must be real numbers. They can be positive or negative, or if c is negative, one of each. Again, however, practical experience provides a guide. If the yield curve is going to flatten out, both of the tau's have to be positive. Otherwise the curve rises forever at an accelerating rate. There is a special case of this model which arises when $b^2 = 4 * c$. In that case the solution is

4.
$$f(t) = L + S * exp(-t / tau) + C * t * exp(-t / tau).$$

NS attempted to fit both equations 2 and 4 to actual forward curves. They concluded that the general solution, equation 2, has too many parameters and overfits the curve, and that the more parsimonious equation 4 is actually a better representation of empirical curves.

The parameters have an interesting interpretation in terms of the yield curve. Assuming that tau is positive, so that the curve eventually levels off, L is equal to the asymptotic forward rate, which is the interest rate "at infinity;" the rate which the curve is leveling off at. S is equal to the difference between the instantaneous spot rate, f(0), and the interest rate at infinity. I.e.

$$L = f(\infty),$$
5.
$$S = f(0) - f(\infty)$$

The exact interpretation of C is not as intuitive, but C is related to how sharply the curve bows up or down relative to a simple exponential curve. One of the complications of interpreting C is that C depends upon the exponential parameter tau.

Tau admits of a simple intuitive interpretation also. It modifies the time scale itself, by converting nominal time, t, into absolute time, t / tau. If tau is small, the curve very quickly (i.e. "quickly" in terms of nominal time) approaches its asymptotic value L. To take a single numerical example, if tau = .03, at the one year maturity point the curve

would already have settled in to the yield which would be reached at thirty-three years, if tau had been equal to 1. Such a curve would appear to be nearly flat, except perhaps for a gap at the front end. The size of the gap, if there is one, would essentially equal S. Conversely, if tau is high, the curve is determined by S and C over the range of maturities for which we actually have data, e.g. out to thirty years. If tau = 30, for instance, the whole curve we which we would actually observe would look like just the first year of a curve for which tau = 1.

One important implication of the difference between nominal time and absolute time scales is that when we try to fit this model to a yield curve, the value of tau will determine whether we are actually able to estimate all the parameters. If tau is small, we probably would not be able to estimate S and C with much accuracy, because only short term yields depend upon them to a significant degree. Conversely, if tau is large, it may not be possible to estimate L with much precision because over the range of available data the curve does not level off sufficiently to identify the asymptotic forward rate. ¹

If we accept for the time being the thesis that the NS parameterization describes the true, unobserved social discount function, then its parameters would explain the pricing of Treasury futures. Before we proceed to the empirical implementation of this idea, however, we have to stop and consider the sources of measurement error.

Fitting Errors of the Model.

The NS Model allows for two sorts of fitting errors which we need to account for as much as possible. They are errors which intervene between the prices of actual Treasury securities and the true social discount function, and errors which arise between the NS curve and the true discount function. A curve model is useful to the extent that the second type of error is small, and our empirical tests will shed some direct light on that matter. One such test is to ask how well the model prices actual Treasury securities, but this is actually a joint test of the absence of both types of fitting errors. Poor performance by the model curve can be the result of either poor performance of the model or of noisy Treasury yields.

What sorts of error does the bond market make in pricing Treasury securities? One kind of error is undoubtedly short run supply / demand imbalance at a single maturity point. There is the potential for an excess supply of a given maturity around the issue date of notes of that maturity. On the demand side, various kinds on institutions are observed to have narrow maturity preferences, and thus when they are active in the market they will tend to move yields at those maturities relative to the rest of the curve.

_

 $^{^1}$ This heuristic insight can be made more precise in terms of the standard errors associated with maximum likelihood estimates of the coefficient parameters L, S, and C. The variance - covariance matrix of the sample estimates is equal to the negative of the expected value of the INVERSE of the Hessian matrix of the log-likelihood function. I.e. to get standard errors of the parameter estimates, one has to invert the matrix of second derivative of the log-likelihood function. As tau approaches either 0 or ∞ the Hessian matrix becomes singular.

One of the great attractions of a parsimonious model of the yield curve is precisely that it is comparatively insensitive to these sorts of security—specific forces. It is appropriate to filter them out because they primarily reflect decisions which are independent of the valuation issue which lies at the heart of a yield curve model. If the Treasury, or some city or corporation, needs money, it will borrow, without asking or much caring whether rates are perhaps out of line with rational expectations. They just need the money, and they are convinced that they can afford to pay the rate which is demanded of them.

There is no way of knowing how tightly observed Treasury yields should fit the true social discount function, or conversely of knowing how large a price impact these sorts of exogenous decisions commonly have. Consequently, we can not use the goodness of fit of the NS model as a basis of judging whether it is the true social discount function. Because of the addition of separate price observations – in this case, observations of the Treasury futures contracts – we do have another, independent test.

Our null hypothesis is that the NS model exactly coincides with the true discount function. The test of that hypothesis rests on the proposition that the true social discount function is the definitive pricing basis for all default-free discount factors, and therefore the definitive pricing basis for all fixed income securities. What this means in practice is that the true curve acts like a sufficient statistic. Given the curve, the prices of fixed income securities, or in our case, the pricing of Treasury futures contracts, should be independent of the prices of individual Treasury securities. The price of the Ten Year Note Contract, for instance, should be independent of the actual yield on ten year notes, holding constant the social discount function. The reason is that the basis between the actual note and the theoretical price dictated by the yield curve only reflects a pricing anomaly in the ten year note; the true curve cannot by definition have pricing anomalies.

If the NS Model correctly identifies the true curve, therefore, it embodies all pricing information germane to the Treasury Futures, and the prices or yield of actual Treasury securities will contain no further information about where the contract should be priced. In formal terms, if y(m) denotes the actual yield on a Treasury security and $y_N(m)$ denotes the yield at that maturity implied by the NS model, the partial derivative of any other price or yield with respect to y(m), holding $y_N(m)$ fixed, will be zero.

Empirical Implementation: Specification.

The program we have laid out consists, roughly, of three distinct steps. The first is to estimate the parameters of the NS model over a sample of yield curves. I will defer a discussion of this step to the next section, which deals with data sources and treatment of the data. The second step starts from the assumption that the NS model captures the true yield curve parameters, by using this data to build an empirical model of the Treasury futures contracts. The third step is to implement the test of sufficiency outlined above. In practice, we do that by adding Treasury cash yields to the contract models. The coefficients of the cash yields are our test statistics. If they are all zero we can safely

conclude that the NS model has passed this test. If a disturbing number are significantly different from zero, we can reject the hypothesis that the NS curve identifies the true yield curve. The exact cutoff for what rates as disturbing is not precise, because obviously if we test enough cash yields, some of these will have high t-ratios. We need a joint test, in the spirit of a F-test of the joint significance of several coefficients.

The Price model of each Treasury contract has three parts, which correspond to the accepted analysis of determinants of futures (see [Burkhardt et. Al.]). One component is called "carry," and arises from the difference between the current yield on the cheapest Treasury to deliver, on the one hand, and the yield on cash balances to be available on the delivery data of the contract. Other things being equal, the price of a contract is discounted relative to the price of cheapest-to-deliver. The owner of futures gives up the yield difference between the cheapest-to-deliver and the yield, generally a lower yield, on cash balances until he actually takes delivery of the bond. Carry is a common factor in all forward contracting, because the contract long earns interest on the cash balances he holds until the point of delivery, but sacrifices any current flows of valuables from the underlying commodity. In the case of Treasury futures, that flow of services is the accrual of coupon on whichever bond or note will actually be delivered. Our proxy for carry uses the yield difference between the Treasury cash yield at an appropriate maturity point – our proxy for the current yield of the cheapest-to-deliver – and the yield of the year bill. The yield on cash is close to some sort of term CD or commercial paper rate and not the year bill, but in practice these yields tend to be close. This yield spread is generally close to the accrual rate of carry, but for safety we have allowed for a constant spread, A, between the simple yield spread and the true rate at which carry accrues. The parameter A is simply added to the model as a parameter to estimate. This yield spread is the instantaneous rate of carry, carry over the term is therefore this spread times the length of the term. In practice, we assume that delivery will be made at the latest possible date, which is the last business day of the next contract month. For the Treasury futures contract at maturity m, then

6. Carry(m) =
$$(y(m) - y(1) + A) * term + u_1$$
, where

7. Term = length of the period from today to the end of the next contract month.

u₁ is a residual term which accounts for all errors in modeling the cost of carry, but in practice they are sure to be small. As defined, the coefficient of Carry will be negative, because the contract long has to be compensated for giving up y(m) and settling for y(1) − A instead. The coefficient of Carry is actually dictated by the model. When the contract and Carry are measured in the right dollar units, the coefficient of Carry has to be −1.

The real complication associated with Treasury futures contracts arises from the fact that the long does not know which note or bond will be the delivered at the end of the contract delivery month. The contract short retains a sheaf of valuable delivery options, which derive their value when the cheapest bond to delivery changes. The underlying option is the option which the short holds to deliver any one of a number of

difference bonds, whose prices do not move be exactly equal dollar amounts. One immediate consequence of this optionality of the contracts is that the basis depends on the volatility of the yield curve over the remaining period until delivery. This is captured by an explanatory factor Volfactor, defined in terms of two factors, Term and the Implied Volatility of the bond contract.

The carry and embedded options effects explain what amounts to an effective spread of the contract off the actual Treasury yield curve. The volatility factor is quite explicitly a spread, and is analogous to any other option spread for a callable bond. The only difference is that the call option embedded in a callable bond expires at the maturity of the bond. The option features of the contract expire much sooner, when the contract settles. The option embedded in Treasury contracts is accordingly much less valuable than a currently exercisable call option would be, because of the limited window of exercise, but in concept the two options are basically the same.

The curve dependence of a Treasury contract can be modeled directly by relating the contract price to the fitted curve. If the contract was a single, known Treasury bond, i.e. if the set of deliverable bonds had only one member, the curve model would take the very simple form of a Taylor expansion of price in terms of yield to maturity, viz.

9.
$$P(ytm) = P(0) - Dur * ytm + .5*Conv * ytm^2 + ...$$

Our model makes three modifications. First, we simply stop the expansion after the second order term and consign the rest of the expansion to the residual term. In so doing, we accept that the theoretical "residuals" are functions of the explanatory variable, yield, but since the residuals are small in magnitude, the resulting bias is insignificant. It follows that we should reject as misspecified any implementation which does not have a high R² statistic. It turns out that this is no problem; empirical R²'s are on the order of .995.

A second modification is that we want to use not "yield to maturity," but the NS model, as our explanatory variable. It is possible to compute the implied yield to maturity of a par bond at any maturity point directly from the NS parameters, but the computation is elaborate and results in a very complex, nonlinear model. It is, furthermore, quite unnecessary. The model we are estimating explains the price of a contract directly in terms of the NS model of the curve. We need not use the actual yield to maturity on the right hand side of equation 9; we can use any statistic of the NS model which is mathematically equivalent to the yield to maturity. Because it results in the simplest model, we choose to use the instantaneous forward rate corresponding to the duration of the cheapest to deliver bond. i.e.

- 10. m = Duration of the Cheapest to Deliver Bond,
- 11. f(m) = L + S * exp(-m / tau) + C * m/tau * exp(-m / tau), and

12.
$$P = P(0) + b_1 * f(m) + b_2 * f(m)^2 + u$$
.

The reason why m is the duration of the cheapest-to-deliver, rather than the maturity, is that the forward rate model is mathematically equivalent to a model of the spot rate with the same m (see Nelson and Siegel). The maturity parameter in this model is the maturity of a spot rate, i.e. the term to maturity of a zero coupon bond. In that case, however, it is also the duration of the instrument – a zero coupon bond – which is used to identify the spot rate.

There is one final modification, which goes directly to the heart of the delivery option embedded in a Treasury futures contract. The duration parameter "m" is itself a function of the level and shape of the yield curve. The set of deliverable bonds contains bonds which have a range of durations. As interest rates change, the prices of the deliverable bonds therefore change at different rates. The bond which lags is liable to become cheapest to deliver. When rates rise, for instance, long duration deliverables fall faster in price than the other deliverables do, and they thus tend to become cheapest to deliver. Thus, the variable m is a function of interest rates. We will simply take a linear approximation as our model of this relationship. In order to allow for effects of curve shape and level, however, we use three points on the curve.

13.
$$m = a_0 + a_1 * y_5 + a_2 * y_{10} + a_3 * y_{30}$$
.

where yT represents yield to maturity at maturity T. The behavior of the cheapest to deliver implies that m is an increasing function of the yield of the cheapest to deliver, which requires at the least the sum $a_1 + a_2 + a_3$ must be positive. Some of the a's may be negative, though. The actual pattern of signs depends on how the factor m depends on the shape of the curve. Is the coefficient of y5 is negative, for instance, it follows that m increases when the curve steepens. If the coefficient of y30 is also negative, m also increases when the curve becomes more curved.

This the complete model, expressing the price of a Treasury contract in terms of carry, implied volatility of interest rates, the four parameters of the NS model, and three Treasury yields. In theory, the yields are themselves functions of the NS parameters, but it would complicate the model needlessly to make this substitution. We will not be able to ignore it, however, when we use this model to estimate the partial derivatives of contract price P with respect to the NS parameters. At that point, we will have to remember to differentiate the y's too.

Empirical Implementation: Data.

There is not a lot of choice when it comes to building a database for estimating this model.² Data on the Treasury contracts first began to emerge with the initiation of trading of the Bond contract ("Bonds") in 1977. Contracts on ten year notes ("Notes"), five year notes ("Five Years") and two year notes ("Two Years") have appeared at approximately five year intervals since then. The Two Year contract is still a fledgling with little liquidity at this time. Implied volatilities are estimated from bond options which first began to trade in the early 1980's. Treasury yield curves, on the other hand, have been around for many decades. Prior to February, 1985, all long Treasury bonds were callable – which means of course that they were priced at some spread off the true yield curve – but since then all newly issued bonds have been free of embedded options. Even before that time, the value of the call option in a new issue long bond was small because the first call date was twenty-five years in the future. Thus, potentially, we have a full set of data which starts for Bonds and Notes in the early 1980's, and for Five Years in the late 1980's.

We do not have a full census of the universe of Treasury data, i.e. daily contracts and yield curves from the early 1980's. We do have two drawings from the universe: month end data covering the whole period, and daily data which starts at the beginning of 1990. Our database has an important limitation; our yield data is limited to the new issue, or "current," Treasuries. For the monthly database we have yields at nine maturity points: 1, 2, 3, 4, 5, 7, 10, 20, and 30 years.³ For the daily database, we have four maturities: 2, 5, 10, and 30 years. Current Treasuries are not necessarily representative of the whole Treasury universe, because the currents are not perfect substitutes for seasoned bonds. The current bonds are in demand for hedging other fixed income positions, e.g. corporate bonds or mortgages, and depending on this demand they can be priced well below yields on comparable seasoned Treasuries. As a new Treasury seasons, its tradable float declines (in the parlance of the bond market, the bonds are "put away"). At that time the yield rises to offset the cost of diminishing liquidity. We thus expect current Treasuries to be priced at some negative spread to the whole Treasury universe. This fact need not invalidate this empirical study, however. To the extent that the true yield curve really is characterized by just a few parameters, the yields of current and seasoned bonds will be very highly correlated even if they are not precisely equal. Creation of the NS Model Curves.

The necessary first step to estimating the model is to obtain parameter estimates of the NS model at every data point, or in other words, to convert observations of individual yields into observations on the parameters. The full curve model has four parameters, including the factor called tau. It is an open question whether it is better—"better" in terms of statistical properties of robustness, consistency, and the like—to estimate the full model or to simply fix tau at 3 years and estimate a restricted, three parameter model. Since in the sample of monthly data we have nine points alone the

_

²We note in passing that the number of serious attempts to model the Treasury contracts could easily exceed the number of months of available data, although that is probably not the case today.

³There has been no current twenty year bond since 1989. After that time, we use the nearest seasoned long bond.

curve, we were able on the monthly data to try both approaches and to compare the results. In applying the model to daily data, since we had only four yield points we concluded that it would be impractical to try to estimate all four curve parameters. Our implementation of the model on daily data, therefore, is for the restricted, three parameter model.

Before we could fit the NS model to yields, we had to derive data on forward rates from the given yields to maturity. Since we have only a few observations from each yield curve, we could not be too sophisticated about this. We used the simple boot-strap method to construct a series of forward yields which span the observed yields. The way this works can be illustrated by taking two adjacent maturities, m_1 and m_2 , and corresponding yields, $y(m_1)$ and $y(m_2)$. Since these are the yields of par bonds, they must also be equal to the coupons rates. The forward yield from m_1 to m_2 is simply the coupon rate, f, of a par bond which will be issued at m_1 , to mature at m_2 . The forward yield, f, is the one which would leave an investor indifferent between the constant coupon stream $y(m_2)$ out to m_2 and the two-part stream which gives $y(m_1)$ out to m_1 and f thereafter.

The curve model applies to exact, instantaneous forward rates, but the forwards we obtain are quoted for a maturity interval. As a reasonable approximation, we assumed that the forward yield between two maturities, m_1 and m_2 , is equal to the instantaneous rate at the midpoint, $(m_1 + m_2)/2$. This converted our yield data into equivalent data on forward rates. For our sample of monthly data, the maturity points quoted above, 1, 2, 3, 4, 5, 7, 10, 20, and 30 years, give rise to forward rates at the maturity points .5, 1.5, 2.5, 3.5, 4.5, 6, 8.5, 15, and 25 years. The final step was to fit the NS model to this data. We used simple, unweighted least squares with nine observations, to estimate the unrestricted, four-parameter model and the restricted, three parameter version of the model. For the sample of daily data, the yields are maturities 2, 5, 10, and 30 years corresponds to forward rates at 1, 3.5, 7.5, and 20 years. As explained above, we estimated only the restricted model on this data set. The yield curve at time t is either a four-tuple (L(t), S(t), C(t), Tau(t)) or a triple (L_R(t), S_R(t), C_R(t)), where the subscript R denotes the restricted model.

Estimates of the Model on Monthly Data.

In this section I will rather briefly summarize the findings of one representative model of each of the three main contracts, bonds, notes, and five year notes. *The* model described above is not really a single, fixed model; it comes in several alternative versions. One difference is which contract we are modeling. The three maturity points give rise to very different parameter estimates. Another difference is whether we restrict the Tau factor to equal 3.0, as explained above, and a third difference is whether we use the simple linear specification outlined above, or use a log-linear specification. Our conclusion after considering all alternatives is that the log-linear, unrestricted model is as good as any version. We will accordingly focus attention on that specification here, leaving the full range of alternatives to an appendix. The exact specification for the bond contract is the following, where Contract(t) denotes the settlement price of the contract at t and yC denotes the Treasury yield corresponding to the notional maturity point of the

contract, i.e. five, ten, or thirty years. As complicated as this expression is, there are only ten parameters to be estimated: there are five b's, four a's, and the spread parameter A which appears in the model of carry. We estimated the parameters by means of the nonlinear least squares routine contained in the RATS statistical package. A table of parameter estimates follows.

$$\begin{split} 14. \qquad & Ln[Contract(t)-b_1\left[yC(t)-r1(t)\right] term(t)-b_2 term(t)-b_3 \ volfactor(t)] \ = \ b_0 + \\ & b_4 * \left\{L(t)+S(t) \exp[(-a_0-a_1\,y5(t)-a_2\,y10(t)-a_3\,y30(t)) \ / \ tau(t)\right] + C(t) * (-a_0-a_1\,y5(t)-a_2\,y10(t)-a_3\,y30(t)) \ / \ tau(t) \right\} + b_5 * \left\{L(t)+S(t) \exp[(-a_0-a_1\,y5(t)-a_2\,y10(t)-a_3\,y30(t)) \ / \ tau(t)\right\} + b_5 * \left\{L(t)+S(t) \exp[(-a_0-a_1\,y5(t)-a_2\,y10(t)-a_3\,y30(t)) \ / \ tau(t)\right\} + C(t) * (-a_0-a_1\,y5(t)-a_2\,y10(t)-a_3\,y30(t)) * \exp[(-a_0-a_1\,y5(t)-a_2\,y10(t)-a_3\,y30(t)) \ / \ tau(t)\right] / \ tau(t) + u(t). \end{split}$$

The summary statistics are generally very favorable for all three contracts. Most importantly, the R² statistics are extremely high, as was mentioned previously. We will have to test directly whether individual Treasury yields are correlated with the regression residuals, which is the most direct test of the thesis of this model, but in any case, the model does not seem to lack much in terms of explanatory power with or without the addition of individual yields. Despite the high proportion of variance explained, the residual variances are still surprisingly high. In the case of the bond contract, the standard error of estimate is greater than 43 / 32nds.⁴ To put that in a profit and loss perspective, it implies that *if the model on the right hand side really equals fair value*, one third of the time a trader would have had an opportunity to make more than one full point in two months by trading the basis. Since the contract must converge to the yield curve at the time of delivery, and since in our sample the contract is two months from delivery, this interpretation of the standard error of estimate follows.

Table 1.

	Bond	Note	Five Yr. Note
$\mathbf{b_0}$	193.07	152.11	126.47
	3.81	1.88	1.22
$\mathbf{b_1}$	4.97	3.61	53
	.71	.65	0.5

_

⁴The contract between the high proportion of explained variance, on the one hand, and the large size of the residual variance is actually a sobering reminder of how volatile Treasury yields have been over this epoch. The range of prices of the Bond contract runs from about 70 to 122. Yields have ranged from more than 14%, at the start of the estimation period (1982), to 5 3/4% in 1993.

b ₂	-8.92	-6.16	0.76
	1.92	1.27	1.09
b ₃	-331	-143	13
	.06	.036	.064
b ₄	-14.51	-7.37	-2.84
	.76	.37	.32
b ₅	.375	.120	052
	.037	.018	.022
\mathbf{a}_0	2.782	1.86	24
	1.46	.65	.33
$\mathbf{a_1}$	2.027	0.014	04
	.79	.32	.12
\mathbf{a}_2	-3.947	40	16
	1.42	.56	.25
a ₃	2.28	0.54	0.45
	.79	.29	.14
S.E.E.	1.34	0.70	0.29
\mathbb{R}^2	. 994	.996	.997
Estimation	10 / 82 – 8 / 96	10 / 82 – 8 / 96	5 / 88 – 8 / 96
Period d.f.	154	154	87
D.W.	1.34	1.31	1.50

Notes: t-ratios appear beneath the corresponding parameter estimates.

There is one observation per month, which is based on the front contract. Each contract, therefore appears three times in the data set, with three, two, and finally one month to remaining to deliver. All interest rates are quoted in per cent, e.g. 5.7%. Term to delivery is quoted in fractions of one year, i.e. term can be either .25, .167, or .083.

All parameters of the Bond model are significant at all conventional levels.

We tested directly for convergence by estimating equation 14 with weights on the residuals. Convergence implies that the variance of model residuals is an increasing function of the term to delivery. We estimated equation 14 in a form which allows the residual variance to be proportional to term. Results of the version are reported in the appendix. In general, there is little to choose between, but the R² statistics of the weighted model are slightly higher than those reported above, consistent with the hypothesis of convergence. Interestingly, the Durbin – Watson statistics of the weighted model indicate that the residuals are less positively autocorrelated, which may be

evidence that a lot of the autocorrelation comes from pairs of successive observations corresponding to the months which are three and two months from the delivery date. It is quite possible that the comparatively large residual errors at these times make it more difficult for the market to recognize and correct mispricing of the contracts.

The coefficient b_3 , which is significant in all equations, relates to the pricing of the implicit delivery options in each contract. The higher that volatility of Treasury yields, the more valuable the delivery option is. Since the option belongs to the short position – to the trader who is long the actual Treasury securities and short futures against them – higher volatility enhances his position by diminishing the value of the contract, other things being equal. The estimates presented above imply that to increase implied volatility by one percentage point with one month remaining to expiration lowers the Bond contract (and raises the Basis) by about 3/32nds. With two months to delivery, the effect is about 4.5/32nds, and at three months it is 5.5/32nds.

The b₄ and b₅ parameters are related to the duration and convexity of the contracts (b₄ is actually the negative of duration, as duration is usually defined). The duration estimates are plausible, though not conventional. The effective duration of the Bond contract is usually estimated to be around 10.5 years. Our estimate depends upon the level of the curve, L. At L equal to 8%, our estimate based on the "b" coefficients would be 8.5 years. The duration estimates for the Note and the Five year note are much closer to conventional estimates, although most models produce a duration a little less than 6 years for the Note, while we get 5.5 years⁵. Our duration estimate for the Five year is almost precisely in line; we get 3.7 years where 3.7 years is more usual. It is not so easy to compare our estimates of convexity with conventional ones, because of large structural differences between our model and conventional models. They approach the modeling task from a theoretical starting point, based on direct modeling of the embedded delivery options, whereas we have taken an empirical tack. We will return to this issue in a moment. For the present, we note that at least the Bond and Note contracts appear to have positive convexity, as they should in our model. Our model recognizes the optionality of Treasury contracts directly, by incorporating a (simple linear) model of the duration of the Cheapest to Deliver bond or note. The "b" coefficients therefore are estimates conditional on holding the Cheapest to Deliver constant. They should thus correspond to the duration and convexity of a non-call bond. Considered in this light, the model of the Five Year Note is less successful, because the convexity estimate is negative, implying that we have not fully accounted for the optionality of the contract. In practice, however, the convexity of a non-call five year note is almost zero, so sampling error could very easily change the apparent sign of the parameter. While it is a negative number, our estimate of the convexity of the Five Year Note contract is almost zero.

⁵ Duration is equal to $-b_4 - 2$ b_5 L. For L equal to 8%, this gives Bond duration = 14.5 - 2 * .375 * 8 = 8.5Note duration = 7.4 - 2 * .12 * 8 = 5.5Five Year Note duration = 2.84 + 2 * .052 * 8 = 3.7

The Model of Cheapest to Deliver.

While the "b" coefficients are in units of duration and convexity, they are not the correct estimates implied by our model, because they are estimated conditional on holding constant the Cheapest to Deliver bond. The true, unconditional estimates also involve the "a" coefficients and the NS model. We will limit our discussion here to the case of the Bond contract, but the analysis is easily extended to the other two contracts. The model of Cheapest to deliver is summed up by equation 13:

13.
$$m = a_0 + a_1 * y_5 + a_2 * y_{10} + a_3 * y_{30}$$
,

where m is the duration of the Cheapest to Deliver. The delivery option enters this relationship through the dependence of the y's on the NS parameters. The y's we have used are all yields to maturity, which complicates their relationship to the NS parameters, but one important case is simple. As the NS model is written, the level factor, L, enters all yields linearly; any X basis point shift in L simply shifts all yields by the same X basis points. Bond duration, moreover, is always equal to the sensitivity of price with respect to this factor, i.e. its sensitivity with respect to a parallel shift of the curve. Thus the derivative of m with respect to L is always equal to the sum of the a_1 , a_2 , and a_3 , and this is precisely the derivative we need to adjust the duration of the contract for optionality. How could we be so lucky!

15.
$$dm/dL = a_1 + a_2 + a_3 = .36$$
 years

Stated in words, a one percentage point increase in the level of the curve lengthens the duration of the Cheapest to Deliver bond by .36 years. We can convert this estimate roughly into maturity terms. The deliverable bonds have maturities between fifteen and thirty years, though in practice the range is from fifteen to twenty-five years, because the newest long bonds are almost always very rich. At a yield of 8%, a fifteen year par bond has a duration of 8.6 years, and a twenty-five year par bond has a duration of 10.7 years. According to our estimates, it would require roughly a 600 basis point shift in yields, from say 5% to 11%, to cause the Cheapest to Deliver to extend by two years of duration. These calculations are admittedly crude, because among other things they ignore the convexity of the individual deliverable bonds. The twenty-five year bond that I used, having an 8% coupon and a duration 10.7 years when yields were at 8%, would have a much smaller duration at 11%. Nonetheless, the estimate is not implausible.

We can now obtain an unconditional, mutatis mutandis estimate of effective duration. For simplification, let X denote the forward rate f(m).

16.
$$d \ln(Contract) / dL = b_4 dX/dL + 2 b_5 X dX/dL$$
, where

17.
$$dX / dL = 1 - \{(S - C) \exp(-m / tau) / tau\} dm / dL$$

-

⁶ At a yield of 11%, its duration is about 8.7 years.

 $- \{C \operatorname{m} \exp(-\operatorname{m} / \operatorname{tau}) / (\operatorname{tau}^2)\} \operatorname{dm}/\operatorname{dL}.$

In order to evaluate this expression we would need values of the NS parameters. From them we would calculate the yields which determine m, and would also calculate f(m). These would then give us the duration estimate. Without carrying our these calculations, we can obtain important qualitative information directly from the general formula. Most importantly, when the curve is positive, S is negative, and in this case, the derivative in equation 17 is larger than 1.0. How much larger depends upon dm/dL. The delivery option therefore causes the contract to have a higher effective duration than it would if m did not depend on L. The size of the difference, moreover, increases linearly with the sensitivity factor, dm/dL.

It is interesting to observe, however, that if the curve is inverted, i.e. if S is positive, and C is small in absolute value, the effect could work the other way. It is possible for dX/dL to be less than 1.0 for some inverted curves. In the same vein, we note that when the curve is flat, the delivery options have no effect on the effective duration of the contract. This is the case in which S and C are both zero, or equivalently, tau is zero. In that case, dX/dL is always exactly equal to 1.0, and the unconditional duration of the contract is equal to its conditional duration based on the b's alone.

Curve Effects.

This summarizes the relationship between the effective, or unconditional duration of a contract and its conditional duration. The model of Cheapest to Deliver also exhibits yield curve dependencies. One common property of the models, for all contracts, is the duration of the Cheapest to Deliver is an increasing function of the slope of the curve. In all models, the coefficient of the bond yield, y30, is positive and the coefficient of the ten year note, y10, is negative. Thus other things being equal, all of these contracts lengthen out when the curve steepens and shorten up when it flattens or inverts. For the bond contract, the most dramatic curve effect arises when the curve "bows out," in the sense that y10 rises relative to y5 and y30. This causes the Cheapest to Deliver to shorten up, and thus causes the Bond contract itself to shorten up.

This concludes out discussion of the behavior of the models themselves. We turn now to the test, which we outlined above, of the hypothesis that the NS curve embodies all relevant information about the True yield curve.

Test of the NS Yield Curve Model.

We can test, at least in a preliminary way, the hypothesis that the NS Model identifies the true, unobserved Treasury yield curve. As explained previously, the basis for the test is the observation that if the NS Model identifies the true curve, all fixed income instruments will be priced off the curve generated by the model, and that any pricing errors will be security-specific. The test is to add Treasury yields to the model developed in the preceding sections. Since under the null hypothesis any mispricings of the Treasury futures contracts are specific to the contracts, they should not be correlated

to yields. Thus the NS Model should serve as a sufficient statistic which embodies all relevant information about the Treasury curve.

While the premise of this test is quite clear and compelling, the actual implementation requires judgment. First of all, we only have eight key Treasury yields in the monthly data, and four yields in the daily data. While in principle we could simply add all the yield data we have to each model, it is obvious in advance that the resulting collinearity would make the test statistics vacuous. None of the Treasury yields would be significant, and we would be unable to reject the null hypothesis. If the test is to be worthwhile, we have to select a few key Treasury yields on which to base a test. Since we have already introduced the yields at five, ten, and thirty years directly into the model, these maturity points do not seem to be an appropriate choice. We have chosen to use three maturity points which seemingly lie in between the maturity points of the contracts. Thus, we use a test based on adding three yields, y3, y7, and y20, to the model. where yk denotes either y3, y7 or y20.

Representative empirical results of this test are summarized in Table 2. They are absolutely overwhelming. Each of the three Treasury yields is highly correlated with all three Treasury contracts in most cases. A representative set of test statistics appears in the

17.
$$\operatorname{Ln}(\operatorname{Contract}(t) - b_1 [\operatorname{yC}(t) - \operatorname{r1}(t)] \operatorname{term}(t) - b_2 \operatorname{term}(t) - b_3 \operatorname{volfactor}(t)) = b_0 +$$

$$\begin{array}{l} b_4 * \left\{L(t) + S(t) \exp[(-a_0 - a_1 y 5(t) - a_2 y 10(t) - a_3 y 30(t)) / tau(t)] + C(t) * (-a_0 - a_1 y 5(t) - a_2 y 10(t) - a_3 y 30(t)) / tau(t)] + C(t) * (-a_0 - a_1 y 5(t) - a_2 y 10(t) - a_3 y 30(t)) / tau(t)] / tau(t) \right\} \\ + b_5 * \left\{L(t) + S(t) \exp[(-a_0 - a_1 y 5(t) - a_2 y 10(t) - a_3 y 30(t)) / tau(t)] + C(t) * (-a_0 - a_1 y 5(t) - a_2 y 10(t) - a_3 y 30(t)) * \exp[(-a_0 - a_1 y 5(t) - a_2 y 10(t) - a_3 y 30(t)) / tau(t)] / tau(t) \right\}^2 \\ + b_6 y k(t) + u(t), \end{array}$$

following table; a fuller set appears in the appendix. The table reports coefficients, standard errors, and t-ratios of these three yields across six different models. All models have the form of equation 17, which is the same as equation 14 with the addition of a linear term which contains one of the three Treasury yields we are using for this test. We report the results of six versions of this model, i.e. two versions for each of the three Treasury contracts. One version corresponds to the model reported in Table 1. The other version is functionally identical to this, but we have weighted the residuals to allow for heteroscedasticity as a function of the term to delivery; the weighting factor is simply one over the square root of "term." As was mentioned above, the weighted equations have slightly higher R² statistics than do the unweighted ones, but in reality all the R²'s are exceptionally close to unity.

All of the test statistics are highly significant. The smallest of the t-ratios is almost 6 in absolute value. All of the estimated coefficients have the appropriate – negative – sign. Holding the NS model fixed, the partial correlation between each yield, on the one hand, and each contract is negative. It follows also that the residuals from the

three contract models in equation 14 are cross correlated, because of course they are all correlated with the three Treasury yields. This also violates the null hypothesis, though it is not really independent evidence. The null hypothesis explained the residuals of any curve model as security—specific or contract—specific mispricing which results from short term supply / demand imbalance. The high degree or positive cross correlation between model residuals implies, quite obviously, that the residuals are not specific to a single contract. They cut across all of them. The fact of cross correlation implies that the model in equation 14 omits important systematic factors, systematic in the sense that they contribute to the pricing of a variety of fixed income objects. We have to speculate whether these omitted factors would also have a significant effect on pricing of other securities, e.g. mortgage securities, corporate bonds, or fixed income swaps.

We have included standard errors for the unweighted equations in Table 2, in order to give some indication of how much the fit has improved by the addition of the Treasury yields. Using the NS model only, the standard error of the bond model was about 1 1/3 points, which is a huge mispricing. Depending on which Treasury yield one adds to the equation, this is reduced to about .9 points or less. While a mispricing of this magnitude still seems to be very large, it is obviously a great improvement. For reference, given the effective duration of about 10 years for the bond contract, a pricing error of 9/10 of a point amounts to about a 90 basis point spread over the Treasury curve.

The statistical results which are summarized in the table provide a source of confirmation of the test methodology, by way of the comparative power of the different maturity effects. The twenty year bond has the highest partial correlation with the bond contract, the seven year note has the highest partial correlation with the note contract, and the three year note has the highest partial correlation with the five year note. Since we wish to interpret our results as implying that the contracts are correlated with individual Treasuries, even after holding the NS curve fixed, it is gratifying that the in each case the most urgent addition to the model – the largest partial correlation – is one which is the most closely related to the actual deliverable set of bonds or notes.

This finding, that each contract is correlated with the Treasuries of nearest maturity, is nonetheless somewhat disturbing. Our hope at the outset of this study was to prove the effectiveness of a parsimonious model of the yield curve, by showing that it could capture most, if not all, of the systematic curve risk which is present in the Treasury futures contracts. From this point of view, even if the four factor NS model had proved to

Table 2.

	у3	<i>y7</i>	y20
Bond		_	
Unweighted	029	05	065
	.002	.003	.003
	-13.4	-18.3	-21.2

SEE	.92	.90	.70
Weighted	-029	051	-066
	.002	.001	.29
	-15.3	-21.0	-22.9
Note		_	
Unweighted	026	06	062
	.0014	.0015	.0019
	-35.7	-38.1	-32.9
SEE	.45	.49	.73
Weighted	027	051	037
	.0012	.0014	.0028
	-22.7	-35.7	-13.3
Five Year Note		-	
Unweighted	025	021	013
	.0013	.0021	.0019
	-18.3	-10.1	-6.9
SEE	.22	.21	.24
Weighted	024	-02	012
	.0015	.002	.002
	-15.6	-9.9	-5.9

Notes: Estimated coefficients are shown with their standard errors and implied t-ratios. All coefficients are highly significant, far exceeding conventional significance levels. For the weighted equations, the standard error of estimate is proportional to the weighting factor, which is the square root of TERM. The standard error of a contract which has two months to delivery are quite comparable to the standard errors of the unweighted equations.

be not entirely successful – even if it is a bit too parsimonious – we would be satisfied to find evidence that another factor is needed. The statistics presented in Table 2 suggest, however, that for theoretical completeness we would need at least three additional factors. While a seven factor model of the curve might be made analytically tractable, it would be hard to describe it as parsimonious.

Conclusion.

We have developed and estimated a set of models of the Treasury futures contracts which attempts to use the NS model as a proxy for the entire yield curve. In some important respects the results are encouraging. First of all, it was possible to define a model which accounts for the optionality of the contracts which comes from the various

delivery options. The way we did this was to take advantage of the fact that the maturity point corresponding to a given yield is a parameter of the NS model. Since the principal delivery option, the so-called quality option, is an option on the maturity of the cheapest to deliver bond, we were able to embed a simple model of the quality option in the NS curve model. Empirically, the resulting option model was successful in capturing the effective duration and convexity of the futures contracts.

Our results are encouraging also because we find that the contracts are very highly correlated with the model. We base this conclusion on the statistics reported in Table 1, which show that all the important factors are highly significant, and that in total our statistical model explains almost all of the sample variance of each contract. From the monthly data reported there, we typically obtain R² statistics in the neighborhood of 99.5%.

Nonetheless, we also found that the part which remains to be explained is large relative to the pricing of futures contracts relative to a measure of fundamental value. The standard deviation of the residuals from the model of the bond contract, for instance, amounts to much more than one point. If our statistical model was exactly equal to the fair value of the bond contract, it would say that the contract is frequently mispriced by one point (i.e. \$1000 per contract) or more by the futures market. It is hard to believe that arbitrage opportunities of this magnitude are available in bond futures, although it is not impossible to believe it.

As a test of the validity of out model, and in effect as a test of the validity of the NS model, we tested the randomness of the model residuals, and found very compelling evidence that they are not random. The residuals from our model of each contract — bonds, ten year notes, and five year notes — are highly correlated with yields of individual Treasuries even though those yields are already reflected in the parameters of the NS model. This led us to reject the proposition that the four parameters of the NS model incorporate all of the systematic behavior of the Treasury yield curve. And on closer inspection, we uncovered some evidence that there is no simple fifth factor which accounts for all the missing information.

We can readily relate these results to a familiar and highly plausible theory of the yield curve, the one which recognizes that both issuers and buyers of bonds have preferred maturity (or duration) niches. The implication of this theory is that "the" bond market is really more like a shopping mall rather than a single retail outlet. For any given player in the market, bonds are not perfect substitutes. They can be induced to depart from their preferred niche when it is sufficiently costly to cling to it. In the buyers' case, this occurs when the yield offered is low enough compared with neighboring maturities. Different points along the curve can then move in different directions, or at different speeds, depending on the momentary balance of supply and demand at that points, and depending on the cross elasticities of supply and demand between different maturity points. The correct number of explanatory factors would necessarily equal the number of (relatively) independent niches.

If the number of niches was small, e.g. two, this theory would have little predictive power. It would account for gross changes in the yield curve, like inversions, but many competing theories can give rise to inversions. The theory of preferred niches is distinctly different from other theories when the number of niches is large. Our empirical results imply that while a model which provides for four niches does a very serviceable job of explaining the pricing of Treasury futures contracts, more niches, perhaps a lot more, are needed to explain all the systematic behavior of the contracts.

Before we throw in the towel on the NS curve model, we would want to replicate this research, but applying it to other fixed income securities. Mortgage-backed bonds are a promising place to start. How much of the systematic behavior of mortgage-backed bonds can be accounted for by the parameters of the NS curve? The hardest part about such a model will be modeling the prepayment option. The rewards of having such a model will more than outweigh the discomforts of developing and estimating it.

Bibliography

- P. Balducci, S. Ranjan Das, S. Foresi, and R. Sundaram. A Simple Approach to Three-factor Affine Term Structure Models. Journal of Fixed Income. December, 1996. Pp. 43 53.
- G. Burghardt, T. Belton, M. Lane, and J. Papa. The Treasury Bond Basis, rev. ed. Irwin Professional Publishing Co.: Burr Ridge, Illinois. 1994.
- P. DeRosa, L. Goodman, and M. Zazzarino. Duration Estimation on Mortgage-backed Securities. Journal of Portfolio Management. Winter, 1996. Pp. 32 38.

- P. Dybvig, J. Ingersoll, and S. Ross. Long Forward and Zero-Coupon Rates Can Never Fall. Journal of Business. 69(1). Pp. 1 26.
- C. Nelson and A. Siegel. Parsimonious Modeling of Yield Curves. The Journal of Business. 60(4), 1987. P. 473 495.