Efficient Predictive Regressions[*]

Jon Faust[†]   Jonathan H. Wright[‡]

[†]Department of Economics, Johns Hopkins University, Baltimore MD 21218, phone: (410) 516-7614, e-mail: faustj@jhu.edu.

[‡]Division of Monetary Affairs, Federal Reserve Board, Washington DC 20551, phone: (202) 452-3605, e-mail: jonathan.h.wright@frb.gov.

Abstract

It is well known that augmenting a standard linear regression model with variables that are correlated with the error term but uncorrelated with the original regressors will increase asymptotic efficiency of the original coefficients. We argue that in the context of predicting excess returns, valid augmenting variables exist and are likely to yield substantial gains in estimation efficiency and, hence, predictive accuracy. The proposed augmenting variables are *ex post* measures of an unforecastable component of excess returns: *ex post* errors from macroeconomic survey forecasts and the surprise components of asset price movements around macroeconomic news announcements. These "surprises" cannot be used directly in forecasting—they are not observed at the time that the forecast is made—but can nonetheless improve forecasting accuracy by reducing parameter estimation uncertainty. We derive formal results about the benefits and limits of this approach and apply it to standard examples of forecasting excess bond and equity returns. We find substantial improvements in out-of-sample forecast accuracy for standard excess bond return regressions; gains for forecasting excess stock returns are much smaller.

KEYWORDS: Excess returns, efficiency, predictive regression, term premiums, seemingly unrelated regression.
JEL Classifications: C22, C53, E17, E43.

1. Introduction

Many empirical papers in finance explore the predictability of excess returns using a simple regression-based approach: estimate a regression for future excess returns based on current predictors and then measure the degree of predictive power of the estimated model. We derive a method to increase efficiency in the estimation step and show that this method can lead to substantial gains in forecastability.

The key idea is familiar from first-year econometrics. If we take any regression and augment it with regressors that are correlated with the error term, but are known not to be correlated with the original regressors in population, we increase asymptotic efficiency of the estimates of the original coefficients without compromising consistency. The augmenting variables are not of direct interest, but *soak up* some residual variance, increasing precision of the estimates of the coefficients that are of interest. This idea is one of many examples of the familiar principle that system estimation imposing correct cross-equation restrictions is more efficient than single equation estimation,[1] but we are not aware systematic treatment of its implications in the forecasting context.

We argue that forecasting excess returns provides an opportunity for gains from this kind of approach. The standard predictive regression is of the form,

$$y_t = \beta' x_{t-h} + \varepsilon_t$$

where y_t is the excess return from $t-h$ to t and the x_{t-h} is observed at $t-h$. In classic

---

[1] As other examples, Hansen (1995) shows how to construct more powerful unit root tests by augmenting the autoregression with additional covariates that are known to be stationary, and Eichenbaum, Hansen and Singleton (1988) show how the power of tests of moment restrictions in a GMM context can be enhanced by imposing other moment restrictions.

applications, the excess return is for government bonds or equity portfolios, and the predictors are term spreads or dividend yields at $t - h$ (e.g., Fama and Bliss (1987), Fama and French (1988), Campbell and Shiller (1988, 1991), Cochrane and Piazzesi (2005)). The variables in such regressions in practice have several distinct properties: excess returns are highly variable and have a large unpredictable component. The predictors are smooth relative to the excess returns and much less variable. In short, we expect large residual variance relative to the variance of the regressors. In finite samples, this is equivalent to saying that the coefficient standard errors will be large.

Our goal is to soak up some of the residual variance to improve the estimates of β. Formally, a valid augmenting variable, w_t must satisfy what we call an *identifying assumption* of being uncorrelated with the regressors, x_{t-h}. If w_t is also correlated with ε_t, then augmenting the estimated regression with w_t will increase the asymptotic efficiency of the estimate of β.

We consider two sets of augmenting variables: The *ex-post* errors in published macroeconomic forecasts, and the surprise components of macroeconomic news announcements that are released between $t - h$ and t. Whether or not these measures satisfy our identifying assumptions is an important question. We derive results about the implications of small violations of the assumption of unpredictability and present tests of this assumption. In the end, however, we argue that out of sample testing presents a stringent test of the empirical benefits of the approach.

Applying our approach to the prediction of excess returns, we show that the efficiency gain in pseudo out-of-sample prediction exercises can be quite substantial for excess bond returns—a 5 to 30 percent reduction in root mean square prediction

error. Gains for forecasting excess stock returns are only found in some cases, and even then are very modest. In no case do we find that the proposed method substantially degrades predictive performance.

In section 2, we describe the econometric methodology in a stylized model and section 3 lays out the specific inference procedures that we apply. Sections 4 and 5, then, report the application to excess bond and equity returns, respectively. Section 6 concludes.

2. Methodology

2.1 The baseline case

This section illustrates the formal logic of our idea in a simplified framework. Begin with a system of equations for the scalar variable to be forecasted, y, the $(p \times 1)$ vector of predictors, x, and a $(k \times 1)$ vector of augmenting variables, w:

$$y_t = \beta' x_{t-1} + \varepsilon_t, \tag{1}$$

$$x_t = A x_{t-1} + v_t \tag{2}$$

$$w_t = \Gamma x_{t-1} + u_t \tag{3}$$

$t = 1, \ldots, T$. Assume that the process for x is stationary and that x_0 is known and fixed. The shock vector, $\zeta_t = (\varepsilon_t, v'_t, u'_t)'$, is iid with $2+\delta$ finite moments for some $\delta > 0$, and $E(\zeta_t | x_{t-1}) = 0$. While the ζs are never directly observed, the other variables with t subscripts are in the forecaster's information set at time t and after, but not before.

In our baseline case, w_t is not correlated with x_{t-1}, which along with what

3

has come before, requires $\Gamma = 0$. Note that ε_t and v_t can be correlated, but ε_t is independent of x_s, $s < t$. Thus, x is *predetermined* but not *strictly exogenous* in (1).

We consider two estimators of the key forecasting parameter, β, in (1): the usual OLS estimator, call it $\hat{\beta}$, and the seemingly unrelated regressions (SUR) estimator based on (1)-(3) *imposing* $\Gamma = 0$. Call this estimator $\tilde{\beta}$. Of course, without the restriction that $\Gamma = 0$, the regressors in each equation of this system are identical, and the SUR and OLS estimates of β coincide.

When we impose $\Gamma = 0$, the SUR estimator of β coincides with a different OLS estimator: the estimator of β in a version of (1) augmented with w_t:

$$y_t = \beta' x_{t-1} + \phi' w_t + \xi_t \tag{4}$$

We codify this result in,

Theorem 1. $\tilde{\beta}$, the pseudo-Gaussian maximum likelihood estimator of β in equations (1)-(3) imposing the restriction that $\Gamma = 0$ is the same as the OLS estimator of β in equation (4).

The proofs of the Theorems are collected in Appendix 1. Theorem 2 compares the asymptotic distributions of $\hat{\beta}$ and $\tilde{\beta}$. Define Ω to be the variance covariance matrix of ζ and partition this matrix conformably with $(\varepsilon_t, v'_t, u'_t)'$ as $\begin{pmatrix} \Omega_{\varepsilon\varepsilon} & \Omega'_{v\varepsilon} & \Omega'_{w\varepsilon} \\ \Omega_{v\varepsilon} & \Omega_{vv} & \Omega'_{uv} \\ \Omega_{u\varepsilon} & \Omega_{uv} & \Omega_{uu} \end{pmatrix}$.
Note that in (4), $\phi = \Omega_{uu}^{-1}\Omega_{u\varepsilon}$ and $\xi_t = \varepsilon_t - \phi' w_t$.

Theorem 2. $T^{1/2}(\hat{\beta} - \beta) \to_d N(0, \sigma^2 \Sigma_{xx}^{-1})$ and $T^{1/2}(\tilde{\beta} - \beta) \to_d N(0, \sigma^2(1-\lambda)\Sigma_{xx}^{-1})$ as $T \to \infty$, where $\Sigma_{xx} = E(x_t x'_t)$, $\sigma^2 = \Omega_{\varepsilon\varepsilon}$ and $\lambda = \sigma^{-2}\Omega'_{u\varepsilon}\Omega_{uu}^{-1}\Omega_{u\varepsilon}$.

Theorem 2 implies that both estimators are consistent, regardless of the correlation between the elements of w_t and ε_t—that is, even if w_t is worthless, augmenting the

4

regression does not affect consistency. If w_t and ε_t are correlated, however, then $\lambda \neq 0$, and $\tilde{\beta}$ is asymptotically more efficient than $\hat{\beta}$. The stronger is the association, the greater is the efficiency gain. Note that λ is the population R^2 in a regression of ε_t on w_t.

Now consider forecast accuracy. The researcher observes x_T and wishes to forecast y_{T+1}. A conventional forecast is

$$\hat{y}_{T+1|T} = \hat{\beta}' x_T,$$

with mean square prediction error $E((y_{T+1} - \hat{y}_{T+1|T})^2 | x_T) = \sigma^2 + T^{-1}\sigma^2 x_T' \Sigma_{xx}^{-1} x_T + o(T^{-1})$ as $T \to \infty$. The alternative predictor proposed in this paper is

$$\tilde{y}_{T+1|T} = \tilde{\beta}' x_T,$$

which has mean square prediction error $E((y_{T+1} - \tilde{y}_{T+1|T})^2 | x_T) = \sigma^2 + T^{-1}\sigma^2(1 - \lambda)x_T' \Sigma_{xx}^{-1} x_T + o(T^{-1})$. Once again, the alternative MSE is smaller if $\lambda \neq 0$.

Note that the augmenting regressors are not used directly in forming the alternative forecast; they are used only at the estimation stage to improve the precision of the estimate of the projection coefficient of y_t on x_{t-1}. The augmented regression soaks up a component of the error term that is uncorrelated with x_{t-1}, thereby reducing the error variance and giving more precise parameter estimates and, hence, better forecasts.

This idea is related to the recent work of Campbell and Yogo (2006), who consider a system consisting of equations (1) and (2) alone, with a scalar x_t. They note that if A is known, then v_t (the innovations to x_t) are observed and that one can then obtain a more powerful test of the hypothesis $\beta = \beta_0$ by subtracting the

component of ε_t that is correlated with v_t. More precisely, if $\beta_0 = 0$ and the errors are Gaussian with known variance-covariance matrix, they show that the optimal test reduces to the conventional t-statistic in a regression of $y_t - \frac{\Omega_{v\varepsilon}}{\Omega_{vv}} v_t$ on $x_{t-1}.^2$ When the error variance-covariance matrix is not known, this suggests augmenting the regression of y_t on x_{t-1} with the additional regressor v_t . Of course the parameter A is generally not known, but Campbell and Yogo model it as being local-to-unity allowing superconsistent estimation. They show how to use Bonferroni methods in conjunction with a superconsistent estimate of A to improve inference on β.

Our setup is different. The A parameter need only be $T^{1/2}$-consistently estimable. In this case, the system estimation of (1) and (2) as in Campbell and Yogo, would reduce to OLS. We posit an additional measured variable in (3), which allows us to form a more efficient system estimator.

There are two clear ways one could go awry in applying this approach: using a large number of augmenting variables that are of limited value (small λ), and using augmenting variables that are, in fact, predictable. In the next two subsections, we provide some results about these two problems.

2.2 Many augmenting variables: the Large k Case

Asymptotically, adding augmenting variables cannot cause an efficiency loss, even if they are uncorrelated with ε_t. However, in a finite sample, choosing a number of augmenting variables, k, that is large relative to T will reduce efficiency. To illustrate

---

[2]This test is optimal in the sense that it is the uniformly most powerful (UMP) test in the system given by equations (1) and (2), conditional on the ancillary statistic $\Sigma_{t=1}^T x_{t-1}^2$. Campbell and Yogo (2006) derive the conditional UMP test for the general hypothesis $\beta = \beta_0$. It only reduces to this t-statistic when $\beta_0 = 0$.

this intuition formally, we consider an alternative asymptotic nesting of our baseline model in which the number of extra regressors goes to infinity at the same rate as the sample size so that k/T approaches a fixed, positive limit.

As usual with such asymptotic nestings, it is not that we are taking a position about what we would do if someone gave us arbitrarily large samples. Rather, we hope this alternative asymptotic theory will give a better guide as to the finite-sample distributions of $\hat{\beta}$ and $\tilde{\beta}$ in moderate sample sizes when k is fairly large relative to T.

We find that there still may be efficiency gains, but there may be losses, and the trade-off turns on how strongly the extra regressors are correlated with ε_t.

Theorem 3. Take the assumptions of the baseline model with the alterations that $k = \alpha T$ and ζ_t is Gaussian such that $\Omega_{v\varepsilon} = 0$ and $\Omega_{uv} = 0$. Then

$$T^{1/2}(\hat{\beta} - \beta) \rightarrow_d N(0, \sigma^2 \Sigma_{xx}^{-1})$$

and

$$T^{1/2}(\tilde{\beta} - \beta) \rightarrow_d N(0, \sigma^2 \tfrac{1-\lambda}{1-\alpha} \Sigma_{xx}^{-1})$$

as $T \rightarrow \infty$.

Under this asymptotic formulation, the efficiency of $\tilde{\beta}$ relative to $\hat{\beta}$ falls as k/T rises. As before, efficiency rises with λ.

2.2 Small Violations of the Identifying Assumption

The assumption that w_t and x_{t-1} are uncorrelated ($\Gamma = 0$) is central: the gains in efficiency come exclusively from imposing what we call our *identifying assumption*. If

this assumption is not satisfied, then $\tilde{\beta}$ will not be consistent. Of course, it will often be difficult in practice to rule out small violations of the assumption. To illustrate the implications of such violations we consider a variation on the baseline model in which Γ is not exactly zero, but instead is local-to-zero: $\Gamma = T^{-1/2}G$. The limiting distributions of $\hat{\beta}$ and $\tilde{\beta}$ are provided in,

Theorem 4. Take the assumptions of the baseline model with the alteration that $\Gamma = GT^{-1/2}$.

$$T^{1/2}(\hat{\beta} - \beta) \to_d N(0, \sigma^2\Sigma_{xx}^{-1})$$

$$T^{1/2}(\tilde{\beta} - \beta) \to_d N(-G'\phi, \sigma^2(1-\lambda)\Sigma_{xx}^{-1})$$

$$E((y_{T+1} - \hat{y}_{T+1|T})^2|x_T) = \sigma^2 + T^{-1}\sigma^2 x_T'\Sigma_{xx}^{-1}x_T + o(T^{-1})$$

and

$$E((y_{T+1} - \tilde{y}_{T+1|T})^2|x_T) = \sigma^2 + T^{-1}x_T'\{G'\phi\phi'G + \sigma^2(1-\lambda)\Sigma_{xx}^{-1}\}x_T + o(T^{-1})$$

as $T \to \infty$.

When $G \neq 0$, $\tilde{\beta}$ is biased and inconsistent, which will tend to degrade forecasting performance. It still may be true, however, that the variance of $\tilde{\beta}$ is smaller than that of $\hat{\beta}$. Thus, we have the sort of bias-variance trade-off that often arises in forecasting. Theorem 4 suggests that the balance can go either way.

Both theorems 3 and 4 set out trade-offs, and one might seek to optimize these trade-offs to choose a number of augmenting variables, k, taking account of modest violations of the identifying assumption, $G \neq 0$. Following up this idea would require a much richer framework, and we do not pursue these lines. In the end, we believe that the case for using the augmented regression would be hard to make if we do not have strong reason to believe that w_t and x_{t-1} are very nearly uncorrelated and so

we henceforth revert to our baseline model in which it is assumed that $\Gamma = 0$.

2.4 A Monte Carlo Simulation

We illustrate the potential gains or losses from using augmenting variables using a Monte Carlo simulation. The design of the experiment is given by the baseline model with $p = 1$ where v_t and ε_t are iid standard normal random variables with correlation $\delta = \Omega_{v\varepsilon}$ and the sample size is T. The value of A is set at 0.9 to capture the high degree of persistence that is common in the explanatory variables in the returns prediction exercises we take up below. Without loss of generality, the true value of β is normalized to zero. The k augmenting variables in w_t are iid $N(0, I_k)$. Table 1 shows the simulated mean square error of $\tilde{\beta}$ relative to that of $\hat{\beta}$ for various values of k, δ, λ and T. Entries less than one mean that the augmented estimator, $\tilde{\beta}$, is doing better than the OLS estimator in a mean-square error sense.

As expected, the higher is λ (the population R^2 in a regression of ε_t on w_t), the better the augmented estimator fares, consistent with Theorem 2. Meanwhile, other things being equal, the augmented estimator works less well as k increases relative to T, consistent with Theorem 3. The augmented estimator can even give a larger mean square prediction error than the ordinary estimator if k is very large relative to T, or if λ very small.

Perhaps the most notable finding here, however, is that with λ as small as 10 percent, the augmented estimator gives a notable improvement in mean square prediction error in all the cases considered except for those in which $k/T > 0.1$.

2.5 The General Model

9

The baseline model is quite stylized. The general regression that we consider in this paper is the h-period-ahead forecasting regression

$$y_t = \beta' x_{t-h} + \varepsilon_t \tag{5}$$

where the p regressors are stationary and predetermined (perhaps including an intercept) and we allow for heteroskedasticity and serial correlation in the errors, with the latter being the norm in the case of overlapping forecasts, $h > 1$. Let $\hat{\beta}$ denote the OLS estimator of β in (5). Let $\tilde{\beta}$ denote the OLS estimator of β in the augmented regression

$$y_t = \beta' x_{t-h} + \phi' w_t + \varepsilon_t \tag{6}$$

where w_t is a $k \times 1$ vector of additional regressors such that

$$w_t = \Gamma x_{t-h} + u_t \tag{7}$$

where $\Gamma = 0$ and $E(\zeta_t | x_{t-h}) = 0$ where $\zeta_t = (\varepsilon_t, v'_t, u'_t)'$ which of course means that $E(w_t | x_{t-h}) = 0$. Given that $\Gamma = 0$, both $\hat{\beta}$ and $\tilde{\beta}$ will be consistent. Motivated by the results in the baseline model, we expect that $\tilde{\beta}$ should be more efficient so long as the extra regressors are relevant ($\phi \neq 0$), though $\tilde{\beta}$ is no longer the system maximum likelihood estimator. We deliberately do not pursue these gains: any further gains above the ones we focus on require stronger assumptions, and we seek to focus on the gains flowing only from existence of an unpredictable, but *ex post* measurable component.[3]

---

[3]For example, in our baseline model, we assumed that x_t was a VAR(1). If instead it were a vector autoregression of order greater than 1, then the regressors in (1) and (2) would no longer be

10

3. Our specific inference procedures

There are three natural hypotheses of interest: 1) Is the identifying assumption that the ws are unpredictable satisfied ($H_1 : \Gamma = 0$?) 2) Are the augmenting variables irrelevant ($H_2 : \phi = 0$)? 3) Is the relative root mean square prediction error of the augmented approach smaller or greater than one ($H_3 : RRMSPE = 1$ against the one-sided alternative $RRMSPE < 1$)? As we define it, RRMSPEs less than one indicate that the augmented approach has smaller errors. In the best case for the approach, we fail to reject the first hypothesis, and reject the latter two. Unfortunately, all of these inferences are likely to be complicated by well-known problems arising from persistent variables.

It is straightforward to test the identifying assumption that $\Gamma = 0$ and that the augmenting variables are worthless ($\phi = 0$): simply run the relevant regression-based Wald test the parameters of interest are zero. In the present case, however, conventional asymptotics for evaluating marginal significance are likely to provide a poor approximation. In our applications—and in most predictive regressions in finance—the predictors, the xs, are highly persistent. Meanwhile, the overlapping nature of the long-horizon returns we will be predicting implies that the appropriate augmenting variables will also be overlapping. Hence, they will be persistent as well. Regressing persistent variables on each other poses severe challenges to small-sample inference: the relationship between $\{w_t\}$ and $\{x_{t-h}\}$ is akin to a spurious regression

---

identical (because of different lag lengths), the SUR estimator of (1) and (2) would not be the same as equation-by-equation OLS, and Theorem 1 would not go through. In this case, taking account of the different lag lengths would admit some further efficiency gain. In general, while the system maximum likelihood estimator should be more efficient than $\hat{\beta}$ or $\tilde{\beta}$, it requires stronger assumptions to be made.

(see, for example, Hodrick (1992), Goetzmann and Jorion (1993) Elliott and Stock (1994) and Stambaugh (1999)). Thus, we expect that the conventional asymptotic p-values would be misleading and that an appropriate bootstrap would be more reliable for testing the hypothesis that $\Gamma = 0$ and, similarly, for testing the hypothesis that $\phi = 0$.

We base our bootstrap on a resampling scheme in which $\Gamma = 0$ and $\phi = 0$ by construction, but the persistence properties of the variables are preserved. In particular, each re-sample holds the excess returns $\{y_t\}$ and predictors $\{x_{t-h}\}$ fixed, but we re-sample from the ws, making them uncorrelated with both excess returns and the predictors. The details of how we re-sample the ws depend on the character of the ws and are described below.

We use this same resampling scheme to assess the statistical significance of all three inferences. For the first two hypotheses we compare the relevant Wald statistic to the empirical distribution of the test statistic under the resampling scheme. In particular, we compute the Wald statistics using Newey-West heteroskedasticity and autocorrelation consistent variance-covariance matrices.[4]

For assessing the statistical significance of any deviation of RRMPSE from 1, we compare the distribution of the relevant Diebold-Mariano statistic to the empirical distribution, in a one-sided test.[5]

---

[4]The truncation lag is h.

[5]The hypothesis that we are effectively testing in this inference procedure is that the augmenting variables are strictly exogenous with respect to $\{y_t\}$ and $\{x_t\}$. That is sufficient, but not necessary, for the augmenting variables to be irrelevant for forecasting. If there is feedback from w_t to future values of x_t, then w_t is not strictly exogenous, but is not necessarily of any help in prediction.

We assess whether our concerns with conventional asymptotics are warranted and whether our bootstrap approach overcomes the problems using a Monte Carlo experiment reported in detail in Appendix 2. Using the 5 percent critical value from the χ^2 distribution, we find that the actual sizes of the Wald tests of $\Phi = 0$ and $\Gamma = 0$ are about 25 percent and 60 percent, respectively. The well-known problems with inference in persistent data manifest themselves quite dramatically in this case. The actual size of the nominal 5 percent bootstrap tests are between 2 and 5 percent. For the test of statistical significance of the RRMSPE, the empirical size of a nominal 5 percent bootstrap test is close to 5 percent. Thus, we conclude our inference approach is fairly well calibrated.

4. Predicting Excess Bond Returns

As noted in the introduction, regression-based excess return prediction is a natural application for our method. We expect these regressions to have very modest predictive power (if any) for excess returns, which are very volatile. These facts suggest that the relevant β may be difficult to estimate precisely. Further, the facts suggest that there will be a great deal of variance in the forecast error that could potentially be *soaked up* if we can find augmenting variables. These intuitions are examined in two widely studied areas: excess bond and equity return prediction.

There are many regressions predicting excess bond returns, but we take as our baseline the recent and influential work of Cochrane and Piazzesi (2005). Their predictions are based on a regression of excess bond returns on the term structure of forward rates.

To describe the regressions, define $P_{n,t}$ to be the price of an n-month zero-

coupon bond in month t; the yield on this bond is $z_{n,t} = -\frac{1}{n}\log(P_{n,t})$, and the 12-month forward rate ending n months hence is $f_{n,t} = \log(P_{n-12,t}) - \log(P_{n,t})$. The return from buying an n-month bond in month $t-12$ and selling it as an $n-12$-month bond in month t is $\log(P_{n-12,t}) - \log(P_{n,t-12})$ and the excess return from holding an n-month bond for 12 months over holding a 12-month bond for that same holding period is

$$rx^n_{t-12,t} = \log(P_{n-12,t}) - \log(P_{n,t-12}) - z_{1,t-12}$$

Cochrane and Piazzesi consider the regression of excess returns on forward rates at the first five annual horizons:

$$
\begin{aligned}
rx^n_{t-12,t} = \quad & \beta_{0,n} + \beta_{1,n} z_{12,t-12} + \beta_{2,n} f_{24,t-12} \\
& + \beta_{3,n} f_{36,t-12} + \beta_{4,n} f_{48,t-12} + \beta_{5,n} f_{60,t-12} + \varepsilon_{t-12,t} \quad (8)
\end{aligned}
$$

for $n = 24, 36, 48, 60$. They also estimate a restricted version of this model in which the coefficients on the forward rates are the same, up to a scaling factor, for each maturity n, i.e. $\beta_{j,n} = \gamma_n \beta_j$. To estimate this restricted model, they first run the regression

$$rx^*_{t-12,t} = \beta_0 + \beta_1 z_{12,t-12} + \beta_2 f_{24,t-12} + \beta_3 f_{36,t-12} + \beta_4 f_{48,t-12} + \beta_5 f_{60,t-12} + \varepsilon^*_{t-12,t} \quad (9)$$

where $rx^*_{t-12,t} = \frac{1}{4}\Sigma^5_{j=2} rx^{12j}_{t-12,t}$, and then regress the excess returns on the fitted values from (9).

We focus on improving the precision of the estimates in equations (8) and (9). Our baseline regressions consist of (8) and (9) estimated using the CRSP Fama-Bliss dataset of monthly zero-coupon bond prices.

We consider several different sets of augmenting variables, each constructed as a sort of *ex post* forecast error from some forecast related to the macroeconomy.

A1. *Ex post* errors from the Survey of Professional Forecasters. Each quarter, in the middle of February, May, August and November, the Survey of Professional Forecasters (SPF) reports analysts' predictions for several variables over the next four quarters. For each SPF back to the beginning of the survey in 1968Q4 we take the median forecast of nominal GDP growth, GDP deflator inflation and the unemployment rate four quarters hence and then take the differences between these forecasts and the actual realized values to form the augmenting variables $\{w_t\}$.[6] Because the survey data are available only at the quarterly frequency, the regressions (8) and (9) are run only using data from each January, April, July and October (i.e. just before the survey) for a total of four observations per year. For example, we use the *ex-post* forecast error from the forecast made in February as the augmenting variable for predicting 12-month excess bond returns starting in January. The timing ensures that only what forecasters learned *after* time $t - 12$ goes into the augmenting variables[7].

A2. Expanded SPF errors. Starting with the 1981Q3 survey, the SPF expanded the set of variables being predicted to include the CPI and some interest rates. Accordingly, for each SPF back to 1981Q3, we take the median predictions of the variables considered in A1 plus CPI inflation, short-term Treasury bill yields

---

[6]For GDP growth and GDP deflator inflation, this is the annualized growth rate from the quarter before the survey to four quarters later. As the actual realized values of the series, we use the first released values from the Federal Reserve Bank of Philadelphia's realtime dataset.

[7]The survey deadline date is a few days before the SPF publication date, but is always in the second month of each quarter.

and long-term Treasury bond yields and again construct realized forecast errors as described above. We then take the first 3 principal components of these 6 realized forecast errors to form the augmenting variables $\{w_t\}$. As with A1, only four observations are used per year.

A3. A News Index of Macro Announcement Surprises. We take the following monthly macroeconomic news announcements: CPI, durable goods orders, housing starts, industrial production, index of leading indicators, nonfarm payrolls, PPI, retail sales and unemployment. For each month and each of these announcement types, we construct the difference between the actual released value and the expected value as found in the MMS survey taken the previous Friday. We form news index as a weighted average of these surprises, giving each type of release a weight equal to the slope coefficient in a regression of the intraday changes in the fourth Eurodollar futures contract[8] from 5 minutes before the announcement until 15 minutes afterwards on the surprise component of the news announcement.[9] This is designed to weight each type of announcement by its market impact. We then cumulate the resulting index over all months from $t - h + 1$ to t, inclusive, to form the augmenting variable w_t where h denotes the horizon of the regression. Our data for these announcements

---

[8]This is the fourth contract in the quarterly cycle and settles to the three-month interest rate about one year hence.

[9]Some announcements come out concurrently. In A3, A4 and A5, the slope coefficients were obtained from a single regression of the intraday change in the fourth Eurodollar futures rate on the surprise components of all of the following announcements: capacity utilization, consumer confidence, CPI (total and core), durable goods orders, the employment cost index, factory orders, the advance release of GDP, hourly earnings, housing starts, initial jobless claims, industrial production, the index of leading indicators, the Michigan survey, NAPM, nonfarm payrolls, new homes sales, personal consumption expenditures, PMI, PPI (total and core), retail sales (total and ex autos) and unemployment. Each surprise was set to zero whenever that particular announcement type did not come out or was missing from our dataset. The regression was run over the period 1982 to 2006.

and, hence, the news index spans 1985:02 to 2006:12.

A4. Expanded News Index. As in A3, except adding the following announcements as well: capacity utilization, core CPI, factory orders, the advance release of GDP[10], new home sales, personal consumption expenditures, core PPI, retail sales excluding autos. Our data for this larger set of announcements go back to 1989:09. The extra variables are not available earlier.

A5. Further Expanded News Index. This is as in A3, except adding the following announcements as well: consumer confidence, initial jobless claims[11], the NAPM index. Our data for this largest set of announcements go back only to 1991:07.

A6. Alternative Aggregation of News Index. We take the first three principal components of the monthly surprises that go into the construction of the index in A5. We then cumulate these principal components over all months from $t - h + 1$ to t, inclusive, to form the augmenting variables w_t.

The cumulative economic news indexes in A3, A4 and A5 are plotted in Figure 1. If the SPF and MMS forecasts are efficient forecasts of upcoming releases and quarterly macroeconomic data respectively, then all the variables A1–A6 must be orthogonal to x_{t-h}. Evidence on the efficiency of SPF survey forecasts is mixed. Froot (1989) and Romer and Romer (2000) report evidence against the efficiency of survey forecasts, but Thomas (1999), Mehra (2002) and Ang, Bekaert and Wei (2007) report more favorable evidence. Much of the discrepancy appears to relate to the sample period

---

[10]This is the one quarterly release that we consider. The monthly advance GDP surprise series is set to zero in all months for which there was no advance GDP announcement.

[11]This is the one weekly release that we consider. All surprises within a given month are cumulated to form the monthly claims surprise series.

17

considered. In the 1970s and early 1980s the surveys appear to have had poor success in forecasting some variables, notably inflation, but have been more successful subsequently. Evidence on the efficiency of MMS survey forecasts (forecasts for a specific news release taken the previous Friday) is more uniformly favorable (see, for example, Balduzzi, Elton and Green (2001)).

In any case, forecast efficiency is a sufficient but not necessary condition for our identifying assumption. If the expectations are efficient but for an expectational error that is orthogonal to x_{t-h}, then w_t and x_{t-h} will still be uncorrelated and the augmented regression will still give a consistent estimate of the population projection coefficient of y_t onto x_{t-h}. Our identifying assumption will only fail if the expectational error is correlated with x_{t-h}, which seems unlikely to us, but which we test below.

With the sets of augmenting variables defined, we can now complete our description of the method for re-sampling from the augmenting variables for the bootstrap. Remember that we hold the excess returns and predictor variables fixed in the bootstrap samples, but re-sample from the augmenting variables so as to generate augmenting variables that are of no value in increasing efficiency ($\phi = 0$) and are not correlated with the predictors ($\Gamma = 0$). Owing to differences in the SPF versus news index variables, the re-sampling methods are slightly different. The SPF data are for overlapping forecast periods and to preserve this structure insofar as possible, we draw blocks of four-quarter-ahead SPF forecast errors with a block length of 3 years. For the news index, we have underlying monthly surprises, which are not overlapping and are arguably uncorrelated. Thus, we re-sample randomly from the monthly sur-

18

prise indexes, and then, as with the actual data, we aggregate these monthly surprise data between $t - h$ and t to form augmenting variables in each bootstrap sample.

4.1 Results

The results from estimating equation (8) using the augmenting regressors A1-A6, as well as the results of the baseline regression over the same sample periods for $n = 24, 36, 48$ and 60 are shown in Table 2. The corresponding results for the estimation of equation (9) are also shown in Table 2. The Table reports estimates of β and asymptotic standard errors from the regressions, both baseline and augmented.

Consistent with the theory, the standard errors on the elements of β in the augmented regressions are typically—though not always—lower than in the baseline regressions. Often they are often substantially lower. This is a preliminary indication that the inclusion of additional regressors may be improving efficiency.

Cochrane and Piazzesi emphasized a "tent" shape in the coefficients on the forward rates whereby the shortest and longest term forward rates have negative coefficients while the intermediate term forward rates have positive coefficients. In some cases, the inclusion of the additional regressors indicates a more pronounced tent shape than is found in the OLS regressions.

Table 2 also gives the Wald test statistics testing the hypothesis that the coefficients on the augmenting regressors are jointly equal to zero ($\phi = 0$ in (6)) which would imply that the augmenting variables are not correlated with the forecast error. For augmenting variables A1–A5, this hypothesis is rejected at conventional significance levels, typically at levels between 0.001 and 0.05. For variables A6 (panel 2-6), the hypothesis is not rejected at the 10 percent level.

19

Finally Table 2 gives p-values for testing our identifying restriction $\Gamma = 0$. This assumption is never rejected at conventional significance levels using the bootstrap p-values. It is true, however, that for variables A1, the result is borderline with a p-value of about 0.12.

4.2 Pseudo-Out-of-Sample Forecasting

Perhaps the most stringent test of the practical usefulness of the proposed approach to prediction can be obtained in a standard pseudo-out-of-sample recursive prediction exercise. Starting half-way through each of the respective estimation periods, we estimate equation (8) using the data that would have been available in that month, and then construct predictions of excess bond returns over the subsequent year using benchmark and augmented regressions; we repeat this estimation and prediction exercise in each subsequent month through the end of the sample. We compute the out-of-sample root mean square prediction error for each of the augmented regressions, relative to the out-of-sample root mean square prediction error from the corresponding baseline regression model. The results are reported in Table 3 for both the unrestricted, (8), and restricted (9), models. A relative root mean square prediction error (RRMSPE) below one means that the augmented regression is giving better out-of-sample predictions of excess bond returns than the baseline.

As can be seen from the Table, the relative root mean square prediction errors are mostly between 0.7 and 0.95, implying about a 5-30 percent reduction in root mean square prediction error. The best results obtain with the SPF regressors (additional regressors A1 and A2) and the news index that combines the largest number of announcements (additional regressors A5). The weakest results are found when using

the three principal components of the surprises as additional regressors (A6).

The improvement in root mean square prediction error is significant at the 5 percent level for additional regressors A1, A2 and A5 and the p-values are between 5 and 20 percent for A3 and A4. Using the principal components of the surprises as additional variables (A6), the improvement in root mean square prediction error is not close to being significant at any conventional level.

5. Predicting Excess Stock Returns

The second predictive regression that we consider is the prediction of excess stock returns, following authors such as Fama and French (1988), Campbell and Shiller (1988) and Ang and Bekaert (2007). We use the following notation: the return on the CRSP value-weighted portfolio from month $t-1$ to month t is $R_{t-1,t} = \log(\frac{P_t + D_t}{P_{t-1}})$ where P_t denotes the price and D_t denotes the dividend in month t. We define three-month excess stock returns as $rx_{t-3,t}^S = \Sigma_{j=0}^2 R_{t-3+j,t-3+j+1} - z_{1,t-3}$, where $z_{1,t-3}$ is the three-month Fama-Bliss risk-free rate. Define the log dividend-price ratio in month t as $dp_t = \log(\Sigma_{j=0}^{11} D_{t-j}/P_t)$. Summing the dividends over the past year deals with the seasonality in dividends. Finally define the stochastically detrended short-term interest rate as $\tilde{r}_t^{RF} = r_t^{RF} - \frac{1}{4}(r_t^{RF} + r_{t-3}^{RF} + r_{t-6}^{RF} + r_{t-9}^{RF})$ where r_t^{RF} is the three-month Fama-Bliss riskfree rate. We consider predicting excess stock returns with the following predictors: dp_t alone, $\tilde{r}_t^{RF}$ alone, and dp_t and $\tilde{r}_t^{RF}$ together in the following three regressions:

$$rx_{t-3,t}^S = \alpha_{1,0} + \alpha_{1,1}dp_{t-3} + \varepsilon_{1t} \tag{10}$$

$$rx_{t-3,t}^S = \alpha_{2,0} + \alpha_{2,1}\tilde{r}_{t-3}^{RF} + \varepsilon_{2t} \tag{11}$$

and

$$rx_{t-3,t}^S = \alpha_{3,0} + \alpha_{3,1}dp_{t-3} + \alpha_{3,2}\tilde{r}_{t-3}^{RF} + \varepsilon_{3t} \tag{12}$$

As in section 3, we take these regressions as the baseline prediction equations for excess stock returns, and our focus is on improving the precision of the estimates in equations (10), (11) and (12). The horizon here is $h = 3$. We estimate these baseline regressions and augmented regressions including the augmenting regressors A1-A6, where these news surprise measures are now cumulated over 3-month rather than 12-month horizons.

Table 4 shows the coefficient estimates and Wald tests from the estimation of (10), (11) and (12) both with and without the additional regressors. Exactly the same information is shown as for the corresponding excess bond return prediction equations. With additional regressors A1 and A2 (SPF forecast errors), the standard errors in the augmented regressions are a bit lower than those in the baseline regressions. The Wald test testing the hypothesis that the additional variables are relevant is significant in (10), (11) or (12), using the bootstrap p-values. On the other hand, using additional regressors A3-A6 (announcement surprise measures), the standard errors in the augmented regressions are little changed from those in the baseline regressions and the Wald test for the joint relevance of these extra regressors is not significant. This is quite different from what we found in the excess bond return prediction equations, but is perhaps not surprising since many authors have found that a considerably

greater fraction of bond price movements can be explained by macroeconomic news announcements than is the case for stock prices (see e.g. Andersen, Bollerslev, Diebold and Vega (2007)).

For each of the augmented regressions in Table 4, we recursively compute the out-of-sample mean square prediction error relative to that from the baseline regression starting half-way through the sample period. The results are reported in Table 5. Again, the results are different depending on whether one uses the SPF forecast errors or announcement surprise measures as additional variables. With the SPF forecast errors, the RRMSPE is a bit below 1, showing some improvement in root mean square prediction error. Using our bootstrap test, the improvement is statistically significant, at least at the 10 percent level, in all cases. The improvement is more modest than was obtained in forecasting excess bond returns with additional regressors A1 and A2. Meanwhile, with additional regressors A3-A6, the RRMSPE is around 1 in all cases and there are no cases in which the augmented regression gives a significant improvement in out-of-sample forecasting performance, as one might expect given that the announcement surprise measures do not have a significant association with quarterly excess stock returns.

6. Conclusions

Researchers using a regressor x_{t-h} to forecast excess returns, y_t, conventionally regress the excess returns on x_{t-h} and use the resulting coefficients for forecasting. However, if there exists a variable w_t that is correlated with the regression error, but not with x_{t-h}, then a more efficient approach to estimating the coefficient on x_{t-h} in the forecasting regression is to augment the regression with w_t. This may in turn enable

23

better forecasts to be constructed, because the coefficient on x_{t-h} is more precisely estimated, even though w_t is not observed at the time the forecast is made, and so cannot be directly used in prediction.

In this paper, we demonstrate the merits of augmenting the estimation model for predictive regressions with *ex post* measures of any unpredictable component of variable being forecasted. This method is most likely to yield advantages in cases such as forecasting excess returns where there is a large unforecastable component and precision of the coefficient estimates is likely to be a major issue. In the excess returns context, we argue that *ex-post* SPF survey forecast errors and the surprise components of macroeconomic news announcements may satisfy the required conditions for augmenting variables.

We demonstrate the merits of the approach using canonical predictive regressions for excess bond and equity returns. The gains are quite pronounced in our extension of the Cochrane and Piazzesi (2005) study of excess bond returns. We find little, if any, gains in conventional equity returns regressions. Our goal in the empirical work was to show the benefits in well-known cases. We suspect further gains may be found in other cases and using other augmenting variables. For example, other possible augmenting variables include data revisions (Koenig et al. (2003)) and oil supply shocks (as identified, for example, by Kilian (2006)), though, at least for the oil supply shocks, we think that the case that these are uncorrelated with the predictors is relatively weak.

The approach could of course easily be misused, say, by searching over a large set of potential augmenting variables for those that give the greatest reduction in

(in sample) standard errors. While out-of-sample tests can provide some protection against the sort of false inference this could promote, we believe that, in practice, our approach should only be entertained if the researcher has a strong belief that the identifying assumption is satisfied. In any case, we think our results suggest that the approach deserves serious consideration.

Appendix 1: Proof of Theorems
Proof of Theorem 1.

Let $\Delta = \Omega^{-1}$ be partitioned conformably with $(\varepsilon_t, v_t', u_t')'$ as $\begin{pmatrix} \Delta_{11} & \Delta_{21}' & \Delta_{31}' \\ \Delta_{21} & \Delta_{22} & \Delta_{32}' \\ \Delta_{31} & \Delta_{32} & \Delta_{33} \end{pmatrix}$.

The submatrices $\Delta_{11}, \Delta_{21}, \Delta_{22}, \Delta_{31}, \Delta_{32}$ and Δ_{33} are $1 \times 1, p \times 1, p \times p, k \times 1, k \times p$ and $k \times k$, respectively. The log-likelihood of $\{y_t, x_t, w_t\}_{t=1}^T$ is given by

$$-\frac{T(p+k-1)\ln(2\pi)}{2} - \frac{T\ln(|\Omega|)}{2} - \frac{1}{2}\Sigma_{t=1}^T \begin{pmatrix} y_t - \beta'x_{t-1} \\ x_t - (I_p \otimes x_{t-1}')a \\ w_t \end{pmatrix}' \Omega^{-1} \begin{pmatrix} y_t - \beta'x_{t-1} \\ x_t - (I_p \otimes x_{t-1}')a \\ w_t \end{pmatrix}$$

where $a = vec(A')$. The first order conditions for the maximization of the log-likelihood with respect to β and a are

$$\Sigma_{t=1}^T \{\Delta_{11}(y_t - \beta'x_{t-1}) + \Delta_{21}'(x_t - (I_p \otimes x_{t-1}')a) + \Delta_{31}'w_t\}x_{t-1} = 0$$

and

$$\Sigma_{t=1}^T \{\Delta_{21}(y_t - \beta'x_{t-1}) + \Delta_{22}(x_t - (I_p \otimes x_{t-1}')a) + \Delta_{32}'w_t\}'(I_p \otimes x_{t-1}') = 0$$

Defining ϕ and $\tilde{\phi}$ as the solutions to the equations

$$-\Delta_{11}\phi' - \Delta_{21}'\tilde{\phi}' = \Delta_{31}'$$

and

$$-\Delta_{21}\phi' - \Delta_{22}\tilde{\phi}' = \Delta_{32}'$$

25

the first order conditions can be rewritten as

$$\Sigma_{t=1}^{T}\{\Delta_{11}(y_t - \phi' w_t - \beta x_{t-1}) + \Delta_{21}'(x_t - \tilde{\phi}' w_t - (I_p \otimes x_{t-1}')a)\}x_{t-1} = 0$$

and

$$\Sigma_{t=1}^{T}\{\Delta_{21}(y_t - \phi' w_t - \beta x_{t-1}) + \Delta_{22}(x_t - \tilde{\phi}' w_t - (I_p \otimes x_{t-1}')a)\}'(I_p \otimes x_{t-1}') = 0$$

which are satisfied by setting β and a to the OLS estimated coefficients on x_{t-1} in regressions of y_t and x_t on w_t and x_{t-1}, respectively.

We can obtain an expression for ϕ as follows. Define $\Delta_A = \begin{pmatrix} \Delta_{11} & \Delta_{21}' \\ \Delta_{21} & \Delta_{22} \end{pmatrix}$, $\Delta_B = \begin{pmatrix} \Delta_{31}' \\ \Delta_{32}' \end{pmatrix}$, $\Omega_A = \begin{pmatrix} \Omega_{\varepsilon\varepsilon} & \Omega_{v\varepsilon}' \\ \Omega_{v\varepsilon} & \Omega_{vv} \end{pmatrix}$ and $\Omega_B = \begin{pmatrix} \Omega_{u\varepsilon}' \\ \Omega_{uv}' \end{pmatrix}$ so that $\Omega^{-1} = \begin{pmatrix} \Delta_A & \Delta_B \\ \Delta_B' & \Delta_{33} \end{pmatrix}$ and $\Omega = \begin{pmatrix} \Omega_A & \Omega_B \\ \Omega_B' & \Omega_{uu} \end{pmatrix}$. From the definition of ϕ and $\tilde{\phi}$, $\begin{pmatrix} \phi' \\ \tilde{\phi}' \end{pmatrix} = -\Delta_A^{-1}\Delta_B$. From the formula for the inverse of a partitioned matrix $\Delta_B = -\Delta_A \Omega_B \Omega_{uu}^{-1}$, so $\begin{pmatrix} \phi' \\ \tilde{\phi}' \end{pmatrix} = \Omega_B \Omega_{uu}^{-1}$ and so $\phi = \Omega_{uu}^{-1}\Omega_{u\varepsilon}$, as required for equation (4). Although we are not aware of any paper where precisely this result has been proven, the idea is not new. For example, similar results were proven by Goldberger (1970) and Goldberger and Olkin (1971).

Proof of Theorem 2.

This follows from the usual formula for the asymptotic distribution of OLS given that $Var(\varepsilon_t) = \sigma^2$ and

$$Var(\xi_t) = Var(\varepsilon_t - \phi' w_t) = \Omega_{\varepsilon e} + \phi'\Omega_{uu}\phi - 2\phi'\Omega_{u\varepsilon} = \Omega_{\varepsilon e} - \Omega_{u\varepsilon}'\Omega_{uu}^{-1}\Omega_{u\varepsilon} =$$
$$\sigma^2 - \Omega_{u\varepsilon}'\Omega_{uu}^{-1}\Omega_{u\varepsilon} = \sigma^2(1 - \lambda)$$

Proof of Theorem 3.

The result for $\hat{\beta}$ follows from the usual formula for the asymptotic distribution of OLS with a fixed number of regressors. To prove the result for $\tilde{\beta}$, write the model in matrix form as $y = X\beta + W\phi + \xi$ where W is a $T \times k$ matrix, the t^{th} row of which is w_t'. Under the stated assumptions, the regressors in this equation are strictly exogenous and so we may condition on them. We can then write

$$\tilde{\beta} = \beta + (X'P_W X)^{-1} X' P_W \xi \tag{A1}$$

where $P_W = I - W(W'W)^{-1}W'$.

Let U be any $n \times m$ matrix with rows $\{u_i\}_{i=1}^n$ that are iid Gaussian with mean zero and variance-covariance matrix Φ, c be any fixed $m \times 1$ vector, and let P be any nonstochastic $n \times n$ projection matrix of rank r. Lemma 2 of Bekker (1994) states that if $r/n \to \psi$ as $n \to \infty$, then $E(n^{-1}U'PUc) \to \psi\Phi c$ and

$$\frac{1}{n^{1/2}}(U'PUc - E(U'PUc)) \to_d N(0, \psi c'\Phi c\Phi + \psi\Phi cc'\Phi) \tag{A2}$$

as $n \to \infty$. Since $Var(n^{-1/2}U'PUa) = O(1)$, $Var(n^{-1}U'PUa) \to 0$, and so, by Chebyshev's Inequality,

$$n^{-1}U'PUc \to_p \psi\Phi c \tag{A3}$$

Now P_W is a projection matrix of rank $T - k$. Letting $U = [X \ \ \xi]$, $n = T$ and

$$\Phi = \begin{pmatrix} \Sigma_{xx} & 0 \\ 0 & Var(\xi_t) \end{pmatrix}$$

From (A3), conditional on W,

$$T^{-1}X'P_W X \to_p (1 - \alpha)\Sigma_{xx}$$

and, from (A2), conditional on W,

$$T^{-1/2}X'P_W\xi \to_d N(0, Var(\xi_t)(1 - \alpha)\Sigma_{xx})$$

and so, from (A1), again conditional on W,

$$T^{1/2}(\tilde{\beta} - \beta) \to_d N(0, \tfrac{1}{1-\alpha}Var(\xi_t)\Sigma_{xx}^{-1})$$

As this limiting distribution does not depend on W, it holds unconditionally as well. As $Var(\xi_t) = \sigma^2(1 - \lambda)$, it follows that

$$T^{1/2}(\tilde{\beta} - \beta) \to_d N(0, \tfrac{\sigma^2(1-\lambda)}{1-\alpha}\Sigma_{xx}^{-1})$$

as required.

Proof of Theorem 4.

The proof for $\hat{\beta}$ is as in Theorem 2. Turning to $\tilde{\beta}$, let $\phi = \Omega_{uu}^{-1}\Omega_{u\varepsilon}$ and $\xi_t = \varepsilon_t - \phi'u_t$. From (1),

$$y_t = \beta'x_{t-1} + \varepsilon_t = \beta'x_{t-1} + \phi'u_t + \xi_t = \beta'x_{t-1} + \phi'(w_t - \Gamma x_{t-1}) + \xi_t = \beta^{*'}x_{t-1} - \phi'w_t + \xi_t$$

where $\beta^* = \beta - \Gamma'\phi$ and $E(\xi_t|x_{t-1}, w_t) = 0$. The estimator $\tilde{\beta}$ is the OLS estimate of the coefficient on x_{t-1} in this regression. Since $\Gamma = GT^{-1/2}$, $T^{1/2}(\beta - \beta^*) = G'\phi$. From the usual formula for the asymptotic distribution of OLS, given that $Var(\xi_t|x_{t-1}, w_{t-1}) = \Omega_{\varepsilon e} - \Omega_{u\varepsilon}'\Omega_{uu}^{-1}\Omega_{u\varepsilon} = \sigma^2(1 - \lambda)$, and

$$T^{-1}\Sigma x_{t-1}w_t' = T^{-1}\Sigma x_{t-1}(x_{t-1}'\Gamma' + u_t') = T^{-3/2}\Sigma x_{t-1}x_{t-1}' + T^{-1}\Sigma x_{t-1}u_t' \to_p 0,$$

it follows that

$$T^{1/2}(\tilde{\beta} - \beta^*) \to_d N(0, \sigma^2(1 - \lambda)\Sigma_{xx}^{-1})$$

Finally, since $T^{1/2}(\tilde{\beta} - \beta) = T^{1/2}(\tilde{\beta} - \beta^*) - T^{1/2}(\beta - \beta^*)$ we have $T^{1/2}(\tilde{\beta} - \beta) \to_d N(-G'\phi, \sigma^2(1 - \lambda)\Sigma_{xx}^{-1})$, as required.

Appendix 2: Monte Carlo Evaluation of Inference Procedures

We rely mainly on bootstrap p-values to assess statistical significance of our results. To investigate the properties of our bootstrap procedure, and of the alternative based on conventional asymptotics, we run a small Monte-Carlo experiment. In each simulation, $\{y_t\}$ and $\{x_{t-h}\}$ were fixed at their values in the sample periods beginning in 1985:02, 1989:09 and 1991:07, respectively. These correspond to the regressions with the three alternative economic news indexes and give sample sizes of T=263, T=208 and T=186, respectively. We draw augmenting variables to mimic key features of the economic news index. In particular, we draw news index values with mean zero and standard deviation 8; this is roughly the standard deviation of each of these economic news indexes. These are then cumulated from time $t - 12$ to time t to form the augmenting variable, making this a Gaussian MA(12) process. The asymptotic and bootstrap p-values for the Wald tests of $\phi = 0$ and $\Gamma = 0$ are then computed. We report the simulated percentage of samples for which these p-values are less than 5 percent. Obviously, a well-calibrated test will give numbers near 5 percent; numbers greater than 5 percent indicate that the true size is greater than the nominal size.

As can be seen from Table A1, the effective size of the asymptotic test of the hypothesis that $\phi = 0$ is elevated, at around 25 percent, while the asymptotic test of the hypothesis that $\Gamma = 0$ shows massive size distortions, with an actual rejection rate of over 60 percent. In contrast, the bootstrap test is roughly correctly sized with the effective size of the test being between 2 and 5 percent. The third block in the table, labelled RRMSPE=1, gives the percentage of times that the bootstrap p-value for testing the hypothesis that RRMSPE=1 in the out-of-sample forecasting experiment is less than 5 percent, using the unrestricted regression (there are no asymptotic p-values for this test). We see that this test has a size that is very close to the nominal level.

References

Andersen, T., T. Bollerslev, F.X. Diebold and C. Vega (2007): Real-Time Price Discovery in Stock, Bond and Foreign Exchange Markets, *Journal of International Economics*, 73, pp.251-277.

Ang, A. and G. Bekaert (2007): Stock Return Predictability: Is it There?, *Review of Financial Studies*, 20, pp.651-707.

Ang, A., G. Bekaert and M. Wei (2007): Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?, *Journal of Monetary Economics*, 54, pp.1163-1212.

Balduzzi, P. Elton, E.J. and T.C. Green (2001): Economic News and Bond Prices: Evidence from the U.S. Treasury Market, *Journal of Financial and Quantitative Analysis*, 36, pp.523-543.

Bekker, P.A. (1994): Alternative Approximations to the Distributions of Instrumental Variables Estimators, *Econometrica*, 62, pp.657-681.

Campbell, J.Y. and R. J. Shiller (1988): The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors, *Review of Financial Studies*, 1, pp.195-228.

Campbell, J.Y. and R. J. Shiller (1991): Yield Spreads and Interest Rate Movements: A Bird's Eye View, *Review of Economic Studies*, 58, pp.495-514.

Campbell, J.Y. and M. Yogo (2006): Efficient Tests of Stock Return Predictability, *Journal of Financial Economics*, 81, pp.27-60.

Cochrane, J.H. and M. Piazzesi (2005): Bond Risk Premia, *American Economic Review*, 95, pp.138-160.

Eichenbaum, M.S., L.P. Hansen and K.J. Singleton (1988): A Time-Series Analysis of Representative Agent Models of Consumption and Leisure Choice under Uncertainty, *Quarterly Journal of Economics*, 103, pp.51-78.

Elliott, G. and J.H. Stock (1994): Inference in Time Series Regression When the Order of Integration of a Regressor is Unknown, *Econometric Theory*, 10, pp.672-700.

Fama, E. and R.R. Bliss (1987): The Information in Long-Maturity Forward Rates, *American Economic Review*, 77, pp.680-692.

Fama, E. and K. French (1988): Dividend Yields and Expected Stock Returns, *Journal of Financial Economics*, 22, pp.3-26.

Froot, K.A. (1989): New Hope for the Expectations Hypothesis of the Term Structure of Interest Rates, *Journal of Finance*, 44, pp.283-305.

Goetzmann, W. and P. Jorion (1993): Testing the Predictive Power of Dividend Yields, *Journal of Finance*, 48, pp.663-679.

Goldberger, A. (1970): Estimation of a Regression Coefficient Matrix Containing a Block of Zeros, Working Paper, University of Wisconsin.

Goldberger, A. and I. Olkin (1971): A Minimum Distance Interpretation of Limited Information Estimation, *Econometrica*, 39, pp.635-639.

Hansen, B.E. (1995): Rethinking the Univariate Approach to Unit Root Testing: Using Covariates to Increase Power, *Econometric Theory*, 11, pp.1148-1171.

Hodrick, R.J. (1992): Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement, *Review of Financial Studies*, 5, pp.357-386.

Kilian L. (2006): Exogenous Oil Price Shocks: How Big are They and How Much do They Matter for the U.S. Economy, Working Paper, University of Michigan.

Koenig, E., S. Dolmas, and J. Piger (2003): The Use and Abuse of 'Real-Time' Data in Economic Forecasting, *Review of Economics and Statistics*, August, 85(3), pp. 618–28.

Mehra, Y.P. (2002): Survey Measures of Expected Inflation: Revisiting the Issues of Predictive Content and Rationality, *Federal Reserve Bank of Richmond Economic Quarterly*, 88, pp.17-36.

Romer, C.D. and D.H. Romer (2000): Federal Reserve Information and the Behavior of Interest Rates, *American Economic Review*, 90, pp.429-457.

Stambaugh, R. (1999): Predictive Regressions, *Journal of Financial Economics*, 54, pp.375-421.

Thomas, L.B. (1999): Survey Measures of Expected U.S. Inflation, *Journal of Economic Perspectives*, 13, pp.125-144.

Table 1: Simulated mean square errors of $\tilde{\beta}$ relative to $\hat{\beta}$

R^2_w	δ	k=1	k=5	k=10	k=20	k=1	k=5	k=10	k=20	k=1	k=5	k=10	k=20
			T=100				T=500				T=1000		
0.00	0	1.01	1.06	1.13	1.25	1.00	1.01	1.02	1.04	1.00	1.00	1.01	1.02
0.01	0	0.99	1.04	1.11	1.23	0.99	0.99	1.00	1.02	0.98	1.00	1.00	1.01
0.1	0	0.90	0.95	1.00	1.11	0.91	0.89	0.91	0.91	0.88	0.92	0.90	0.92
0.3	0	0.76	0.80	0.84	0.93	0.76	0.75	0.76	0.75	0.73	0.78	0.75	0.78
0.5	0	0.66	0.69	0.72	0.80	0.65	0.65	0.66	0.72	0.63	0.68	0.65	0.67
0.00	0.5	1.00	1.05	1.11	1.21	1.00	1.00	1.01	1.03	1.00	1.01	1.01	1.03
0.01	0.5	1.00	1.04	1.09	1.19	0.99	0.99	1.00	1.00	0.99	1.00	1.00	1.01
0.1	0.5	0.92	0.96	0.98	1.08	0.92	0.92	0.92	0.90	0.90	0.93	0.90	0.91
0.3	0.5	0.76	0.80	0.81	0.88	0.78	0.78	0.78	0.74	0.75	0.79	0.75	0.76
0.5	0.5	0.65	0.69	0.69	0.74	0.67	0.68	0.68	0.64	0.64	0.68	0.64	0.66
0.00	0.9	1.00	1.02	1.08	1.16	1.00	1.01	1.01	1.02	1.00	1.01	1.01	1.04
0.01	0.9	1.00	1.01	1.07	1.15	0.99	1.00	1.00	1.00	0.99	1.00	0.99	1.02
0.1	0.9	0.91	0.91	0.96	1.02	0.91	0.93	0.92	0.90	0.99	0.92	0.89	0.92
0.3	0.9	0.72	0.74	0.76	0.81	0.78	0.78	0.78	0.76	0.77	0.76	0.74	0.76
0.5	0.9	0.59	0.62	0.63	0.67	0.67	0.67	0.68	0.66	0.66	0.65	0.63	0.66

Notes. Relative mean square errors of baseline and augmented regressions obtained by Monte-Carlo simulations, as described in the text. In each experiment, 1000 replications were conducted.

Table 2-1: Regressions predicting excess bond returns using forward rates and additional regressors A1

Regression:	$n = 24$ Baseline	Aug.	$n = 36$ Baseline	Aug.	$n = 48$ Baseline	Aug.	$n = 60$ Baseline	Aug.	Average Baseline	Aug.
y_{12}	-0.75	-1.34	-1.41	-2.44	-2.13	-3.47	-2.73	-4.37	-1.76	-2.91
	(0.22)	(0.20)	(0.38)	(0.33)	(0.50)	(0.44)	(0.61)	(0.52)	(0.43)	(0.37)
f_{24}	0.12	1.04	-0.25	1.35	0.05	2.10	0.25	2.75	0.05	1.81
	(0.45)	(0.36)	(0.77)	(0.64)	(1.01)	(0.85)	(1.26)	(1.07)	(0.86)	(0.72)
f_{36}	1.04	1.18	2.63	2.85	2.74	3.01	2.99	3.31	2.35	2.59
	(0.38)	(0.32)	(0.64)	(0.58)	(0.86)	(0.80)	(1.10)	(1.02)	(0.74)	(0.68)
f_{48}	0.80	0.08	1.48	0.24	2.83	1.25	3.43	1.51	2.14	0.77
	(0.29)	(0.25)	(0.52)	(0.48)	(0.67)	(0.65)	(0.84)	(0.80)	(0.58)	(0.54)
f_{60}	-0.94	-0.88	-2.02	-1.89	-2.92	-2.75	-3.23	-3.00	-2.28	-2.13
	(0.37)	(0.28)	(0.63)	(0.48)	(0.84)	(0.66)	(1.07)	(0.84)	(0.72)	(0.56)
Wald ϕ		64.33		47.05		43.33		43.37		46.79
p-val (boot)		0.000		0.000		0.000		0.000		0.000
p-val (asy)		0.000		0.000		0.000		0.000		0.000
Wald Γ		120.88								
p-val (boot)		0.119								
p-val (asy)		0.000								

Notes. The baseline regressions show the estimated coefficients in regressions of excess 24-, 36-, 48- and 60-month bond returns on the term structure of forward rates and of the average of these four excess returns on the term structure of forward rates. The regressions are run on data from the first month of each quarter from 1968Q4 to 2006Q4. Asymptotic standard errors are shown in parentheses. All asymptotic standard errors and p-values are Newey-West with a lag length of 4. The augmented regressions control for the additional variable A1. The estimated coefficients on this additional variable are not reported, but the row Wald ϕ denotes the Wald statistic testing the hypothesis that it is equal to zero along with the associated asymptotic and bootstrap p-values, constructed as described in the text. The row Wald Γ denotes the Wald statistic testing the hypothesis that $\Gamma = 0$, again along with the associated asymptotic and bootstrap p-values. These are the same for each regression in the table because they depend only on $\{x_t\}$ and $\{w_t\}$.

Table 2-2: Regressions predicting excess bond returns using forward rates and additional regressors A2

Regression:	$n = 24$ Baseline	Aug.	$n = 36$ Baseline	Aug.	$n = 48$ Baseline	Aug.	$n = 60$ Baseline	Aug.	Average Baseline	Aug.
y_{12}	-0.09	-0.56	-0.36	-1.26	-0.83	-2.02	1.07	-2.52	-0.59	-1.59
	(0.39)	(0.14)	(0.73)	(0.27)	(0.97)	(0.40)	(1.22)	(0.51)	(0.82)	(0.31)
f_{24}	-0.29	0.67	-0.81	1.04	-0.60	1.87	-0.57	2.48	-0.57	1.52
	(0.60)	(0.26)	(1.10)	(0.50)	(1.45)	(0.75)	(1.84)	(1.06)	(1.23)	(0.62)
f_{36}	0.42	0.36	1.56	1.43	1.38	1.16	1.19	0.88	1.14	0.96
	(0.40)	(0.15)	(0.79)	(0.33)	(1.13)	(0.56)	(1.39)	(0.78)	(0.92)	(0.44)
f_{48}	0.81	-0.33	1.50	-0.69	2.81	-0.09	3.45	-0.11	2.14	-0.30
	(0.23)	(0.12)	(0.44)	(0.22)	(0.60)	(0.35)	(0.75)	(0.45)	(0.50)	(0.27)
f_{60}	-0.49	-0.01	-1.31	-0.38	-2.03	-0.76	-2.08	-0.51	-1.48	-0.41
	(0.27)	(0.09)	(0.57)	(0.20)	(0.80)	(0.34)	(1.01)	(0.49)	(0.66)	(0.27)
Wald ϕ		419.60		520.30		383.26		321.19		415.27
p-val (boot)		0.000		0.000		0.000		0.000		0.000
p-val (asy)		0.000		0.000		0.000		0.000		0.000
Wald Γ		121.16								
p-val (boot)		0.329								
p-val (asy)		0.000								

Notes. As in Table 2-1, except using the additional variable A2 and the sample period 1981Q3 to 2006Q4.

Table 2-3: Regressions predicting excess bond returns using forward rates and additional regressors A3

Regression:	$n = 24$		$n = 36$		$n = 48$		$n = 60$		Average	
	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.
y_{12}	-0.28	-0.60	-0.84	-1.45	-1.60	-2.43	-2.43	-3.41	-1.29	-1.97
	(0.44)	(0.33)	(0.83)	(0.61)	(1.09)	(0.80)	(1.29)	(0.96)	(0.91)	(0.67)
f_{24}	-0.55	-0.15	-0.99	-0.22	-0.63	0.40	0.19	1.41	-0.50	0.36
	(0.77)	(0.54)	(1.37)	(0.95)	(1.81)	(1.24)	(2.17)	(1.50)	(1.53)	(1.05)
f_{36}	1.32	1.46	3.12	3.38	3.56	3.91	3.74	4.15	2.94	3.23
	(0.93)	(0.62)	(1.01)	(1.11)	(2.31)	(1.51)	(2.76)	(1.84)	(1.92)	(1.26)
f_{48}	0.71	0.53	0.87	0.66	1.90	1.43	1.74	1.19	1.34	0.95
	(0.45)	(0.30)	(0.59)	(0.62)	(1.22)	(0.92)	(1.54)	(1.25)	(1.00)	(0.76)
f_{60}	-0.92	-0.88	-1.80	-1.72	-2.50	-2.39	-2.35	-2.22	-1.89	-1.80
	(0.44)	(0.33)	(0.82)	(0.63)	(1.13)	(0.89)	(1.39)	(1.13)	(0.94)	(0.74)
Wald ϕ		31.52		29.27		26.45		20.89		26.00
p-val (boot)		0.005		0.006		0.006		0.008		0.006
p-val (asy)		0.000		0.000		0.000		0.000		0.000
Wald Γ		13.41								
p-val (boot)		0.567								
p-val (asy)		0.037								

Notes. The baseline regressions show the estimated coefficients in regressions of excess 24-, 36-, 48- and 60-month bond returns on the term structure of forward rates and of the average of these four excess returns on the term structure of forward rates. The regressions are run on monthly data from 1985:02 through 2006:12. Asymptotic standard errors are shown in parentheses. All asymptotic standard errors and p-values are Newey-West with a lag length of 18. The augmented regressions control for the additional variable A3. The estimated coefficients on this additional variable are not reported, but the row Wald ϕ denotes the Wald statistic testing the hypothesis that it is equal to zero along with the associated asymptotic and bootstrap p-values, constructed as described in the text. The row Wald Γ denotes the Wald statistic testing the hypothesis that $\Gamma = 0$, again along with the associated asymptotic and bootstrap p-values. These are the same for each regression in the table because they depend only on $\{x_t\}$ and $\{w_t\}$.

Table 2-4: Regressions predicting excess bond returns using forward rates and additional regressors A4

Regression:	$n = 24$ Baseline	Aug.	$n = 36$ Baseline	Aug.	$n = 48$ Baseline	Aug.	$n = 60$ Baseline	Aug.	Average Baseline	Aug.
y_{12}	-0.17	-0.39	-0.74	-1.17	-1.57	-2.13	-2.54	-3.19	-1.25	-1.72
	(0.48)	(0.39)	(0.88)	(0.71)	(1.14)	(0.91)	(1.34)	(1.06)	(0.96)	(0.76)
f_{24}	-1.18	-0.57	-1.99	-0.83	-1.87	-0.34	-1.03	0.74	-1.52	-0.25
	(0.82)	(0.62)	(1.43)	(1.07)	(1.85)	(1.38)	(2.24)	(1.67)	(1.58)	(1.18)
f_{36}	2.46	1.96	5.34	4.39	6.74	5.49	7.63	6.17	5.54	4.50
	(0.96)	(0.65)	(1.74)	(1.15)	(2.26)	(1.49)	(2.70)	(1.79)	(1.90)	(1.26)
f_{48}	0.01	0.24	-0.68	−0.25	-0.78	-0.21	-1.90	-1.24	-0.84	-0.37
	(0.54)	(0.41)	(1.01)	(0.74)	(1.32)	(0.98)	(1.58)	(1.22)	(1.10)	(0.82)
f_{60}	-0.55	-0.85	-0.95	-1.51	-1.21	-1.95	-0.65	-1.51	-0.84	-1.45
	(0.46)	(0.44)	(0.96)	(0.90)	(1.36)	(1.29)	(1.74)	(1.67)	(1.12)	(1.07)
Wald ϕ		15.77		13.37		11.08		8.81		11.23
p-val (boot)		0.030		0.039		0.053		0.071		0.054
p-val (asy)		0.000		0.000		0.001		0.003		0.001
Wald Γ		26.042								
p-val (boot)		0.541								
p-val (asy)		0.000								

Notes. As in Table 2-3, except using the additional variable A4 and the sample period 1989:09 to 2006:12.

Table 2-5: Regressions predicting excess bond returns using forward rates and additional regressors A5

Regression:	$n = 24$ Baseline	Aug.	$n = 36$ Baseline	Aug.	$n = 48$ Baseline	Aug.	$n = 60$ Baseline	Aug.	Average Baseline	Aug.
y_{12}	-0.11	-0.20	-0.64	-0.82	-1.44	-1.68	-2.22	-2.51	-1.10	-1.30
	(0.63)	(0.37)	(1.17)	(0.67)	(1.50)	(0.84)	(1.77)	(0.98)	(1.26)	(0.71)
f_{24}	-1.38	-0.89	-2.27	-1.30	-2.20	-0.90	-1.57	0.00	-1.86	-0.77
	(0.92)	(0.60)	(1.62)	(1.02)	(2.06)	(1.32)	(2.46)	(1.59)	(1.76)	(1.12)
f_{36}	2.87	2.31	5.93	4.84	7.43	5.96	8.34	6.57	6.14	4.92
	(0.96)	(0.63)	(1.78)	(1.13)	(2.36)	(1.49)	(2.86)	(1.82)	(1.98)	(1.25)
f_{48}	-0.28	0.00	-1.13	-0.58	-1.33	-0.59	-2.56	-1.67	-1.32	-0.71
	(0.58)	(0.35)	(1.08)	(0.62)	(1.42)	(0.82)	(1.70)	(1.04)	(1.19)	(0.69)
f_{60}	-0.53	-0.89	-0.88	-1.59	-1.10	-2.06	-0.31	-1.46	-0.71	-1.50
	(0.58)	(0.44)	(1.25)	(0.93)	(1.76)	(1.34)	(2.24)	(1.75)	(1.45)	(1.10)
Wald ϕ		25.39		26.28		23.64		20.33		23.60
p-val (boot)		0.011		0.010		0.012		0.016		0.011
p-val (asy)		0.000		0.000		0.000		0.000		0.000
Wald Γ		16.756								
p-val (boot)		0.704								
p-val (asy)		0.010								

Notes. As in Table 2-3, except using the additional variable A5 and the sample period 1991:07 to 2006:12.

Table 2-6: Regressions predicting excess bond returns using forward rates and additional regressors A6

Regression:	$n = 24$		$n = 36$		$n = 48$		$n = 60$		Average	
	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.
y_{12}	-0.11	-0.22	-0.64	-0.86	-1.44	-1.75	-2.22	-2.61	-1.10	-1.36
	(0.63)	(0.37)	(1.17)	(0.70)	(1.50)	(0.91)	(1.77)	(1.09)	(1.26)	(0.76)
f_{24}	-1.38	-0.88	-2.27	-1.28	-2.20	-0.84	-1.57	0.07	-1.86	-0.73
	(0.92)	(0.63)	(1.62)	(1.12)	(2.06)	(1.46)	(2.46)	(1.77)	(1.76)	(1.24)
f_{36}	2.87	2.38	5.93	4.99	7.43	6.17	8.34	6.83	6.14	5.09
	(0.96)	(0.61)	(1.78)	(1.05)	(2.36)	(1.37)	(2.86)	(1.65)	(1.98)	(1.15)
f_{48}	-0.28	0.26	-1.13	-0.15	-1.33	-0.13	-2.56	-1.20	-1.32	-0.30
	(0.58)	(0.36)	(1.08)	(0.58)	(1.42)	(0.72)	(1.70)	(0.84)	(1.19)	(0.60)
f_{60}	-0.53	-1.00	-0.88	-1.74	-1.10	-2.12	-0.31	-1.45	-0.71	-1.58
	(0.58)	(0.36)	(1.25)	(0.75)	(1.76)	(1.04)	(2.24)	(1.29)	(1.45)	(0.85)
Wald ϕ		20.30		23.57		23.08		22.50		23.12
p-val (boot)		0.218		0.180		0.161		0.139		0.167
p-val (asy)		0.000		0.000		0.000		0.000		0.000
Wald Γ		429.29								
p-val (boot)		0.208								
p-val (asy)		0.000								

Notes. As in Table 2-3, except using the additional variables A4 and the sample period 1991:7 to 2006:12. In this case, the row Wald ϕ gives the Wald statistic testing the joint hypothesis that all the elements of ϕ are equal to zero.

Table 3: Out-of-Sample RRMSPE for excess bond returns from augmented regression relative to baseline

Additional Regressors		Bond Maturity			
		$n = 24$	$n = 36$	$n = 48$	$n = 60$
A1	Unrestricted	0.913	0.910	0.914	0.902
		0.059	*0.041*	*0.041*	*0.030*
	Restricted	0.948	0.925	0.900	0.892
		0.124	*0.060*	*0.030*	*0.016*
A2	Unrestricted	0.907	0.871	0.848	0.777
		0.023	*0.006*	*0.002*	*0.000*
	Restricted	0.815	0.815	0.826	0.851
		0.001	*0.000*	*0.001*	*0.002*
A3	Unrestricted	0.910	0.925	0.940	0.957
		0.059	*0.103*	*0.146*	*0.189*
	Restricted	0.960	0.948	0.938	0.931
		0.182	*0.147*	*0.147*	*0.130*
A4	Unrestricted	0.893	0.891	0.897	0.904
		0.119	*0.123*	*0.145*	*0.166*
	Restricted	0.908	0.903	0.894	0.895
		0.153	*0.140*	*0.140*	*0.141*
A5	Unrestricted	0.685	0.690	0.707	0.723
		0.008	*0.009*	*0.019*	*0.027*
	Restricted	0.730	0.719	0.699	0.698
		0.016	*0.015*	*0.017*	*0.018*
A6	Unrestricted	0.911	0.931	0.967	0.991
		0.316	*0.367*	*0.451*	*0.488*
	Restricted	0.973	0.965	0.958	0.956
		0.470	*0.439*	*0.437*	*0.433*

Notes. This table shows the root mean square prediction error for excess bond returns from the augmented model in which the augmenting regressors are predicted to be zero divided by the root mean square prediction error from the baseline model. Bootstrap p-values for one-sided tests testing the hypothesis of equality in root mean square prediction errors are shown in italics. The models are as described in the notes to Table 2. Models are either restricted or unrestricted: Restricted predictions impose that there is a single return forecasting factor with a different loading for each bond maturity n. Predictions are pseudo-out-of-sample with the first prediction made half-way through the estimation period.

Table 4-1: Regressions predicting excess stock returns using log dividend-price ratio and/or stochastically detrended short-term interest rate and additional regressors A1

Regression:	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.
dp_t	0.018	0.031			0.016	0.031
	(0.018)	(0.016)			(0.019)	(0.016)
$\tilde{r}_t^{RF}$			-0.749	-0.357	-0.550	-0.103
			(1.037)	(0.920)	(1.036)	(0.903)
Wald ϕ		38.104		39.676		39.352
p-val (boot)		0.000		0.000		0.000
p-val (asy)		0.000		0.000		0.000
Wald Γ		1.863		7.323		8.063
p-val (boot)		0.750		0.197		0.587
p-val (asy)		0.601		0.062		0.234

Notes. The baseline regressions show the estimated coefficients in regressions of 3 month excess stock returns (CRSP value-weighted returns less the Fama-Bliss riskfree rate) on the corresponding log dividend-price ratio and/or the stochastically detrended short term interest rate. The regressions are run on data from the first month of each quarter from 1968Q4 through 2006Q4. Asymptotic standard errors are shown in parentheses. The augmented regressions control for the additional variable A1. The estimated coefficients on this additional variable are not reported, but the row Wald ϕ denotes the Wald statistic testing the hypothesis that it is equal to zero along with the associated asymptotic and bootstrap p-values, constructed as described in the text. The row Wald Γ denotes the Wald statistic testing the hypothesis that $\Gamma = 0$, again along with the associated asymptotic and bootstrap p-values. These are the same for each regression in the table because they depend only on $\{x_t\}$ and $\{w_t\}$.

Table 4-2: Regressions predicting excess stock returns using log dividend-price ratio and/or stochastically detrended short-term interest rate and additional regressors A2

Regression:	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.
dp_t	0.013	0.020			0.013	0.020
	(0.015)	(0.014)			(0.015)	(0.014)
$\tilde{r}_t^{RF}$			-1.336	-0.894	-1.343	-0.897
			(0.668)	(0.615)	(0.646)	(0.596)
Wald ϕ		23.77		16.68		18.66
p-val (boot)		0.002		0.004		0.003
p-val (asy)		0.000		0.001		0.000
Wald Γ		12.78		24.49		34.85
p-val (boot)		0.718		0.002		0.217
p-val (asy)		0.005		0.000		0.000

Table 4-3: Regressions predicting excess stock returns using log dividend-price ratio and/or stochastically detrended short-term interest rate and additional regressors A3

Regression:	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.
dp_t	0.034	0.036			0.034	0.037
	(0.018)	(0.019)			(0.019)	(0.019)
$\tilde{r}_t^{RF}$			0.221	0.270	0.386	0.476
			(1.005)	(0.987)	(0.983)	(0.956)
Wald ϕ		0.608		0.270		0.681
p-val (boot)		0.545		0.681		0.517
p-val (asy)		0.436		0.603		0.409
Wald Γ		0.125		1.237		1.540
p-val (boot)		0.782		0.365		0.633
p-val (asy)		0.724		0.266		0.463

Notes. The baseline regressions show the estimated coefficients in regressions of 3 month excess stock returns (CRSP value-weighted returns less the Fama-Bliss riskfree rate) on the corresponding log dividend-price ratio and/or the stochastically detrended short term interest rate. The regressions are run on monthly data from 1985:02 through 2006:12. Asymptotic standard errors are shown in parentheses. The augmented regressions control for the additional variable A3. The estimated coefficients on this additional variable are not reported, but the row Wald ϕ denotes the Wald statistic testing the hypothesis that it is equal to zero along with the associated asymptotic and bootstrap p-values, constructed as described in the text. The row Wald Γ denotes the Wald statistic testing the hypothesis that $\Gamma = 0$, again along with the associated asymptotic and bootstrap p-values. These are the same for each regression in the table because they depend only on $\{x_t\}$ and $\{w_t\}$.

Table 4-4: Regressions predicting excess stock returns using log dividend-price ratio and/or stochastically detrended short-term interest rate and additional regressors A4

Regression:	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.
dp_t	0.035	0.033			0.038	0.036
	(0.025)	(0.025)			(0.025)	(0.025)
$\tilde{r}_t^{RF}$			0.752	0.845	1.006	1.077
			(1.126)	(1.119)	(1.072)	(1.069)
Wald ϕ		1.943		2.783		2.132
p-val (boot)		0.242		0.157		0.223
p-val (asy)		0.163		0.095		0.144
Wald Γ		4.575		0.986		4.765
p-val (boot)		0.387		0.407		0.508
p-val (asy)		0.032		0.321		0.092

Notes. As in Table 4-3, except using the additional variable A4 and the sample period 1989:09 to 2006:12.

Table 4-5: Regressions predicting excess stock returns using log dividend-price ratio and/or stochastically detrended short-term interest rate and additional regressors A5

Regression:	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.
dp_t	0.066	0.065			0.065	0.065
	(0.025)	(0.025)			(0.024)	(0.025)
$\tilde{r}_t^{RF}$			1.060	1.076	0.981	0.990
			(1.206)	(1.206)	(1.113)	(1.111)
Wald ϕ		0.043		0.222		0.063
p-val (boot)		0.845		0.679		0.818
p-val (asy)		0.835		0.637		0.802
Wald Γ		2.001		0.193		2.067
p-val (boot)		0.467		0.737		0.660
p-val (asy)		0.157		0.660		0.356

Notes. As in Table 4-3, except using the additional variable A5 and the sample period 1991:07 to 2006:12.

Table 4-6: Regressions predicting excess stock returns using log dividend-price ratio and/or stochastically detrended short-term interest rate and additional regressors A6

Regression:	Baseline	Aug.	Baseline	Aug.	Baseline	Aug.
dp_t	0.066	0.068			0.065	0.068
	(0.025)	(0.024)			(0.024)	(0.023)
$\tilde{r}_t^{RF}$			1.060	1.356	0.981	1.276
			(1.206)	(1.176)	(1.113)	(1.057)
Wald ϕ		6.753		7.128		7.460
p-val (boot)		0.230		0.197		0.189
p-val (asy)		0.080		0.068		0.059
Wald Γ		0.348		2.293		2.905
p-val (boot)		0.973		0.723		0.948
p-val (asy)		0.951		0.514		0.821

Notes. As in Table 4-3, except using the additional variables A6 and the sample period 1991:7 to 2006:12. In this case, the row Wald ϕ gives the Wald statistic testing the joint hypothesis that all the elements of ϕ are equal to zero.

Table 5: Out-of-Sample RMSPE for excess stock returns from augmented regression relative to baseline

Additional Regressors	Predictors		
	dp_t	$\tilde{r}_t^{RF}$	dp_t & $\tilde{r}_t^{RF}$
A1	0.960	0.947	0.917
	0.010	*0.001*	*0.002*
A2	0.974	0.978	0.975
	0.093	*0.000*	*0.089*
A3	0.999	1.000	0.998
	0.357	*0.398*	*0.284*
A4	1.028	0.995	1.033
	0.955	*0.189*	*0.967*
A5	1.014	0.999	1.014
	0.876	*0.390*	*0.862*
A6	1.013	0.996	1.005
	0.736	*0.339*	*0.555*

Notes. This table shows the root mean square prediction error for excess stock returns from the augmented model in which the augmenting regressors are predicted to be zero divided by the root mean square prediction error from the baseline model. Bootstrap p-values for one-sided tests testing the hypothesis of equality in root mean square prediction errors are shown in italics. The models are as described in the notes to Table 4.

Table A1: Simulated Effective Size of Nominal 5 Percent Tests of $\phi = 0$, $\Gamma = 0$, and RRMSPE= 1

	$\phi = 0$					$\Gamma = 0$	RRMSPE=1			
	n=24	n=36	n=48	n=60	Average		n=24	n=36	n=48	n=60
T=263										
Asymptotic	23.9	23.2	21.6	20.5	21.7	62.5				
Bootstrap	4.9	5.2	5.0	4.7	4.8	2.3	4.5	4.1	4.0	4.2
T=208										
Asymptotic	26.2	25.1	24.9	24.2	24.9	69.4				
Bootstrap	5.0	5.0	4.8	4.6	4.7	3.3	4.9	4.7	5.0	4.8
T=186										
Asymptotic	29.0	28.9	28.0	27.6	27.8	75.3				
Bootstrap	5.4	5.4	5.8	5.8	5.6	2.8	5.2	5.0	4.9	4.4

Notes. In these simulations, $\{y_t\}$ and $\{x_{t-h}\}$ were fixed at their values in the excess bond return prediction application using sample periods beginning in 1985:02, 1989:09 and 1991:07, respectively. These correspond to the regressions with three alternative economic news indexes and give the sample sizes of T=263, T=208 and T=186, respectively. Normal random draws of the economic news index with mean zero and standard deviation 8 were taken and were then cumulated from t-12 to t to form the augmenting variable $\{w_t\}$, making this a Gaussian MA(12). The Wald test statistics testing the hypothesis $\phi = 0$ (for the different excess return series) and $\Gamma = 0$ were then computed and compared to the asymptotic and bootstrap critical values. The table reports the fraction of simulations in which the test rejected using asymptotic and bootstrap critical values (nomional significance level: 5 percent). The table also reports the percentage of times that the bootstrap p-value for testing the hypothesis that RRMSPE=1 in the out-of-sample forecasting experiment is less than 5 percent, using the unrestricted regression (there are no asymptotic p-values for this test).

Fig. 1: 12-Month Moving Sum of Economic News Indexes Ending in Month Shown