

Shrunken Interest Rate Forecasts are Better Forecasts

Reid Dorsey-Palmateer
Division of Research and Statistics
Federal Reserve Board
Washington, D. C. 20551

Gary Smith
Department of Economics
Pomona College
Claremont, California 91711

running title: Shrunken Interest Rate Forecasts

Abstract

Predicted changes in interest rates are imperfectly correlated with actual changes in interest rates. One statistical consequence may be that large predicted changes are more likely to be overestimates than underestimates of the magnitude of the change. If so, the accuracy of predicted interest rate changes can be improved by shrinking them toward a prior mean of zero. The application of this idea to interest rate forecasts by the Survey of Professional Forecasters found a consistent improvement in the accuracy of their predictions.

contact: Gary Smith, Department of Economics, Pomona College, Claremont, California 91711;
909.607.3135; gsmith@pomona.edu; Fax 909.621.8576

Shrunken Interest Rate Forecasts are Better Forecasts

Interest rate forecasts are potentially of great value to households, businesses, and policymakers, and there have been many studies of the accuracy of interest rate forecasts based on macroeconomic models, financial futures contracts, and the term structure (for example, Friedman 1979, 1980; Shiller 1990; Campbell and Shiller, 1991; Holden and Thompson 1996; Baghestani, Jung, and Zuchegno 2000; Den Butter and Jansen 2004).

Many continuous time models, such as Vasicek (1977) and Cox, Ingersoll, Ross (1985), assume that interest rates revert toward long-equilibrium values. Nowman (1998) estimated several continuous-time models with mean-reversion for a variety of countries; Hejazi and Li (2000) found mean-reversion in the forward-rate premiums for U. S. Treasury bills; and Chua, Koh, and Ramaswamy (2005) conclude that yield spreads revert to their historical averages.

A superficially similar but conceptually quite different argument is that predictions can often be improved by taking into account the regression to the mean that can be expected when people make imperfect predictions. Vergin (2001) and Lee and Smith (2002) identify successful betting strategies based on the presumption that bettors do not fully appreciate regression to the mean in the performance of National Football League teams. Keil, Smith, and Smith (2004) show that the accuracy of analysts' earnings forecasts for a cross section of companies can be improved by shrinking their forecasts toward the mean forecast. In this paper, we use interest rate forecasts from the Survey of Professional Forecasters to see whether the accuracy of predicted changes can be improved by shrinking them toward a prior mean of zero.

Regression Toward the Mean

Regression to the mean is frequently encountered in a sequence of cross-section data. For example, Galton (1886) observed regression toward the mean in his seminal study of the heights of parents and their adult children. Tall parents tend to have tall children, but the children of any

two parents are not all the same height because parental genes are not the only factor that determines their children's heights. Thus, adults who are 78 inches tall may be somewhat taller or shorter than their genetically predicted heights, but the former is more likely because there are many more people with genetically predicted heights below 78 inches than with genetic heights above 78 inches. Thus the observed heights of unusually tall parents usually overstate the genetic heights that they inherit from their parents and pass on to their children. The average heights of the children of unusually tall (or short) parents regress to the mean.

Similarly, observed athletic performances are an imperfect measure of skills. Schall and Smith (2000) looked at major league baseball players from 1901 through 1999 who had at least 50 times at bat or 25 innings pitched in two consecutive seasons. Of 4026 players who had batting averages of .300 or higher in any season, 80% did worse the following season. Of 3849 players who had earned run averages of 3.00 or lower in any season, 80% did worse the following season.

The educational testing literature provides a well-established framework for explaining regression to the mean in a sequence of cross-section data (Kelley 1947; Lord and Novick 1968). A person's *true score* μ is the statistical expected value of his or her score on a test. It is assumed that a person's observed score X on the test differs from the true score by an independent and identically distributed *error score* ϵ :

$$X = \mu + \epsilon \quad (1)$$

Looking at a test involving a group of students, there is a distribution of true scores and observed scores. Those who score the highest are likely to have positive error scores because it would be unusual for someone to score below his or her true score and still have the highest score on a test. Since a score that is high relative to the group is also likely to be high relative to this person's true score, this person's score on another test is likely to regress toward the mean.

The same principle may apply to a group of interest rate forecasts made over a period of time.

At each point in time, we interpret \bar{Y} as the expected value of the change in an interest rate and X as the predicted change. Looking at a large group of forecasts made at different times, a natural assumption is that the distribution of \bar{Y} has a mean of zero. As with educational tests, large positive values of X are more likely to have positive \bar{Y} than to have negative \bar{Y} , while the reverse is true of large negative values of X . Therefore, the expected value of the change in interest rates may be closer to zero than is the forecast.

A Model of Regression Towards the Mean

Let the actual change Y in an interest rate at some point in time have a probability distribution with expected value \bar{Y} ,

$$Y = \bar{Y} + \epsilon, \quad E[\epsilon] = 0 \quad (2)$$

and assume that an expert forecast X differs from the expected value of the actual change by an independent and identically distributed error score ϵ :

$$X = \bar{Y} + \epsilon, \quad E[\epsilon] = 0 \quad (3)$$

All of these random variables and parameters are for a given time period but, for notational simplicity, the time subscript is suppressed. Even though the forecasts are unbiased, their accuracy can be improved by shrinking them toward a mean of zero!

If we look at a set of forecasts and expected values across time, the population covariance between X and \bar{Y} is equal to the variance of \bar{Y} :

$$\begin{aligned} \text{cov}[X, \bar{Y}] &= E[(X - E[X])(\bar{Y} - E[\bar{Y}])] \\ &= E[(\bar{Y} - E[\bar{Y}] + \epsilon - E[\epsilon])(\bar{Y} - E[\bar{Y}])] \\ &= E[(\bar{Y} - E[\bar{Y}])^2] \\ &= \sigma_{\bar{Y}}^2 \end{aligned}$$

In educational testing, a test's reliability is gauged by the squared correlation between scores and abilities, which equals the ratio of the variance of \bar{Y} to the variance of X :

$$\begin{aligned}
 R^2 &= \frac{\text{cov}[X, \hat{X}]^2}{\sigma_x^2 \sigma_{\hat{X}}^2} \\
 &= \frac{\sigma_{\hat{X}}^2}{\sigma_x^2} \\
 &= \frac{\sigma_{\hat{X}}^2}{\sigma_x^2}
 \end{aligned} \tag{4}$$

Here, we interpret R^2 as the reliability of the forecast. If the standard deviation of \hat{X} were 0, the forecasts would be error free and the reliability would be 1. As the standard deviation of \hat{X} becomes infinitely large, the reliability of the forecasts approaches 0.

If we knew the values of \hat{X} , we could use Equation 3 to make unbiased predictions of the forecasts. Table 1 shows some hypothetical data for six time periods. The forecasts are 0.5 below and 0.5 above the expected values of -1, 0, and 1. The forecasts are unbiased with an expected value of each forecast equal to the expected value of the change in the interest rate. Figure 1 shows that a least squares regression of X on \hat{X} has an estimated slope of 1.

However, we are interested in the reverse question: inferring the expected value of the change in the interest rate from the forecast change. Figure 2 shows that a least squares regression of \hat{X} on X has an estimated slope of 0.727. Since the least squares regression goes through the mean values of the variables, the least squares predicted deviation from the mean of the expected value of the interest rate change equals 0.727 times the deviation from the mean of the forecast change:

$$\hat{X} - \bar{\hat{X}} = 0.727(X - \bar{X})$$

In general, if we use a large sample to estimate

$$\hat{X} = \beta + \alpha X + \epsilon \tag{5}$$

by ordinary least squares, the slope will be close to

$$\begin{aligned}
 b &= \frac{\text{cov}[X, \Delta]}{\sigma_X^2} \\
 &= \frac{\sigma_\Delta^2}{\sigma_X^2} \\
 &= r^2
 \end{aligned}$$

The squared correlation between X and Δ should be used to shrink forecasts toward the mean:

$$\Delta_{\text{shrink}} = r^2(X - \bar{X})$$

In a large sample in which the means of X and Δ are close to 0,

$$\Delta = r^2 X \tag{6}$$

Thus, to predict the expected value of the change in an interest rate, we use the reliability of the forecast to shrink the expert forecast toward zero.

A Bayesian Interpretation

Recognizing that the error term represents the cumulative effects of a great many omitted variables and appealing to the central limit theorem, we assume that the error term Δ is normally distributed with mean 0 and standard deviation σ_Δ . A convenient conjugate prior for Δ is provided by a normal distribution with mean Δ_0 and standard deviation σ_0 . The mean of the posterior distribution for Δ is part way between the forecast and our prior mean:

$$\Delta = \frac{\sigma_0^2}{\sigma_0^2 + \sigma_\Delta^2} X + \frac{\sigma_\Delta^2}{\sigma_0^2 + \sigma_\Delta^2} \Delta_0 \tag{7}$$

If we had no information about interest rate changes in a particular period other than the analysts' forecast, we might set the prior mean for Δ equal to 0 and the prior standard deviation equal to the standard deviation of Δ over time. If so, Equation 7 becomes identical to Equation 6:

$$\begin{aligned}\hat{\beta} &= \frac{\sigma_{\beta}^2}{\sigma_{\beta}^2 + \sigma_{\epsilon}^2} X \\ &= \beta^2 X\end{aligned}$$

Implementing the Model

In practice, we don't observe β and consequently cannot use data for X and β to estimate β^2 . However, we do have data on Y , the actual change in the interest rate, and the covariance between X and Y is equal to the variance of β :

$$\begin{aligned}\text{cov}[X, Y] &= E[(X - E[X])(Y - E[Y])] \\ &= E[(\beta - E[\beta] + \epsilon)(\beta - E[\beta] + \epsilon)] \\ &= E[(\beta - E[\beta])^2] \\ &= \sigma_{\beta}^2\end{aligned}$$

The population correlation between X and Y equals the ratio of the variance of β to the variance of X , which Equation 4 shows to be the squared correlation between X and β :

$$\begin{aligned}\rho_{XY} &= \frac{\text{cov}[X, Y]}{\sigma_X \sigma_Y} \\ &= \frac{\sigma_{\beta}^2}{\sigma_X \sigma_Y} \\ &= \frac{\sigma_{\beta}^2}{\sigma_X^2} \\ &= \beta^2\end{aligned}\tag{8}$$

Thus we can use the sample correlation between X and Y to estimate β^2 , which we can then use to implement Equation 6. The correlation between the predicted and actual changes should be used to shrink each forecast toward 0.

It is reasonable that the appropriate shrinkage depends on the correlation between forecast and actual changes. If forecast and actual changes were perfectly correlated, we would not shrink the forecasts at all. If forecast and actual changes were uncorrelated, the forecast would be useless in

predicting interest rate changes. We would ignore the forecast and predict no change.

Data

The ASA/NBER Economic Outlook Survey was started by American Statistical Association and the National Bureau of Economic research in 1968. In 1990 the survey was taken over by the Federal Reserve Bank of Philadelphia and renamed the Survey of Professional Forecasters.

Approximately 35 professional forecasters are surveyed each quarter. The survey is distributed near the end of the first month of the quarter and returned by the middle of the second month.

Quarterly forecasts one-to-four quarters into the future are available for these three interest rates:

1. Treasury Bill Rate, three-month, secondary market rate, discount basis (unadjusted data, percent, daily). Source: Federal Reserve Release G.13 or H.15, or Federal Reserve Bulletin Table 1.35, Item 21. These forecasts are available from the third quarter of 1981.
2. Treasury Bond Rate, 10-year, constant-maturity, secondary market rate (unadjusted data, percent, daily). Source: Federal Reserve Release G.13 or H.15, or Federal Reserve Bulletin Table 1.35, Item 32. These forecasts are available from the first quarter of 1992.
3. Moody's AAA Corporate Bond Yield, (unadjusted data, percent, daily). Source: Moody's Investors Service, or Federal Reserve Release G.13 or H. 15, or Federal Reserve Bulletin Table 1.35, Item 40. These forecasts are available from the third quarter of 1981.

In each case, the quarterly values are averages of the monthly data and we use the mean of the professional forecasts. For each interest rate, we look at the predicted change: the difference between the predicted future value and the actual value during the quarter when the forecast is made. We work with the predicted change in interest rates because zero is a natural prior mean for a group of predicted changes made over time.

Our model suggests that analysts' predicted changes can be improved by using the correlation between predicted and actual changes to shrink each predicted change toward zero. To see

whether this is so, we use the historical data available at the time of the forecast to estimate the correlation between predicted and actual interest rate changes. For example, to adjust the analysts' forecasts made in the second quarter of 1998 of the change in Moody's AAA Corporate Bond Yield between the second and fourth quarters of 1998, we calculate the correlation between forecast and actual two-quarters ahead changes in this interest rate using data ending with the forecasts made in the fourth quarter of 1997. The adjusted forecast is then calculated by multiplying this correlation coefficient times the two-quarters-ahead forecast made in the second quarter of 1998.

One question is the minimum number of observations needed to estimate the correlation coefficient. The fewer the observations, the less confidence we have in the estimate; but the more observations, the fewer adjusted forecasts we have to test the model. We settled on a minimum of 10 observations as a reasonable compromise. We also looked at a minimum of 20 observations and this had no effect on the results other than reducing the number of forecasts.

Forecasting accuracy is measured in three ways: the number of occasions in which the adjusted or unadjusted forecasts are closer to the actual values, the mean absolute error for the adjusted and unadjusted forecasts, and the root mean square error for the adjusted and unadjusted forecasts.

Results

Table 2 compares the accuracy of the Survey of Professional Forecasters (SPF) forecasts and the adjusted forecasts. For a count of the number of forecasts that were closer to the actual value, the binomial distribution gives the p-value for a test of the null hypothesis that each method has a 0.5 probability of being closer to the actual value. A matched-pair test uses the difference between the adjusted and unadjusted forecast errors to test the null hypothesis that the expected value of the difference is zero; the t distribution gives the p-value. All of the p-values shown in

Table 2 are two-sided.

By every measure for every forecast variable, the adjusted forecasts are more accurate than the SPF forecasts, though the differences are most statistically persuasive for long-term bonds. As gauged by the number of forecasts that were closer to the actual values, seven of eight two-sided p-values for the long-term bonds are less than 0.05, and five are less than 0.01. All 12 matched-pair p-values are less than 0.001.

Summary

The statistical principle of regression to the mean suggests that the accuracy of predicted changes in interest rates might be improved by using the historical correlation between predicted and actual changes to shrink the predicted changes toward a prior mean of zero. The application of this idea to interest rate forecasts by the Survey of Professional Forecasters found a consistent improvement in the accuracy of their predictions.

References

- Baghestani, H., Jung, W., and Zuchegno, D. (2000) On the information content of futures market and professional forecasts of interest rates. *Applied Financial Economics*, 10, 679-684.
- Den Butter, F. A. G., and Jansen, P. W. (2004), An empirical analysis of the German long-term interest rate, *Applied Financial Economics*, 14, 731-741.
- Campbell, J., and Shiller, R. J. (1991) Yield spreads and interest rate movements: a bird's eye view, *Review of Economic Studies*, 58, 495-514.
- Chua, T. C., Koh, W. T. H., and Ramaswamy, K. (2005) Comparing returns of US treasuries versus equities: implications for market and portfolio efficiency, *Applied Financial Economics*, 15, 1213-1218.
- Cox, J. C., Ingersoll, J. E., Jr., and Ross, S. A. (1985) A theory of the term structure of interest rates, *Econometrica*, 53, 385-407.
- Friedman, B. M. (1979) Optimal expectations and the extreme information assumptions of "rational expectations" macromodels, *Journal of Monetary Economics*, 5, 23-41.
- Friedman, B. M. (1980) Survey evidence on the "rationality" of interest-rate expectations, *Journal of Monetary Economics*, 6, 453-465.
- Galton, F. (1886) Regression towards mediocrity in hereditary stature, *Journal of the Anthropological Institute*, 15, 246-263.
- Hejazi, W., and Li, Z. (2000) Are forward premia mean reverting?, *Applied Financial Economics*, 10, 343-350.
- Holden, K., and Thompson, J. L. (1996) Forecasting interest rates from financial futures markets. *Applied Financial Economics*, 6, 449-461.
- Keil, M., Smith, G., and Smith, M. H. (2004) Shrunken earnings predictions are better predictions, *Applied Financial Economics*, 14, 937-943.

- Kelley, T. L. (1947) *Fundamentals of Statistics*, Harvard University, Cambridge, MA, 409.
- Lee, M., and Smith, G. (2002) Regression to the mean and football wagers, *Journal of Behavioral Decision Making*, 15, 329–342.
- Lord, F. M., and Novick, M. R. (1968) *Statistical theory of mental test scores*, Addison-Wesley, Reading, MA.
- Mankiw, N. G., and Summers, L. H. (1984) Do long-term interest rates overreact to short-term interest rates?, *Brookings Papers on Economic Activity*, 1984: 1, 223-247.
- Nowman, K. B. (1998) Continuous-time short term interest rate models. *Applied Financial Economics*, 8, 401-407.
- Schall, T., and Smith, G. (2000) Baseball players regress toward the mean, *The American Statistician*, 54, 231-235.
- Shiller, R. J. (1990) The Term Structure of Interest Rates, in *Handbook of Monetary Economics*, Vol. 1, B. Friedman and F. Hahn, eds., North-Holland, Amsterdam, 627-723.
- Terregrossa, S. J. (1999) Combining analysts' forecasts with causal model forecasts of earnings growth. *Applied Financial Economics*, 9, 143-153.
- Vasicek, O. (1977) An equilibrium characterization of the term structure, *Journal of Financial Economics*, 5, 177-188.
- Vergin, R. C. (2001) Overreaction in the NFL point spread market, *Applied Financial Economics*, 11, 497-509.

Table 1 Hypothetical Expected Values μ and Forecasts X

μ	X
-1	-1.5
-1	-0.5
0	-0.5
0	0.5
1	0.5
1	1.5

Table 2 Measures of Forecasting Accuracy

	Number Closer to Actual Value			matched-pair p-value	MAE		RMSE	
	SPF	Adjusted	p-value		SPF	Adjusted	SPF	Adjusted
T-bill+1	37	51	.165	0.000	0.37	0.35	0.50	0.47
T-bill+2	35	50	.128	0.000	0.68	0.65	0.86	0.84
T-bill+3	37	47	.326	0.000	0.95	0.90	1.20	1.16
T-bill+4	36	46	.320	0.000	1.17	1.13	1.46	1.40
T-bond+1	9	37	.000	0.000	0.45	0.27	0.53	0.35
T-bond+2	14	30	.023	0.000	0.63	0.47	0.74	0.57
T-bond+3	15	27	.088	0.000	0.70	0.55	0.84	0.67
T-bond+4	12	28	.017	0.000	0.77	0.60	0.90	0.76
AAA+1	26	62	.000	0.000	0.43	0.28	0.53	0.35
AAA+2	25	61	.000	0.000	0.64	0.45	0.77	0.55
AAA+3	22	62	.000	0.000	0.79	0.58	0.97	0.73
AAA+4	28	54	.005	0.000	0.88	0.73	1.09	0.92

MAE: mean absolute forecasting error

RMSE: root mean square error

SPF: Survey of Professional Forecasters

Adjusted: forecasts adjusted for regression to the mean

The p-values are two-sided

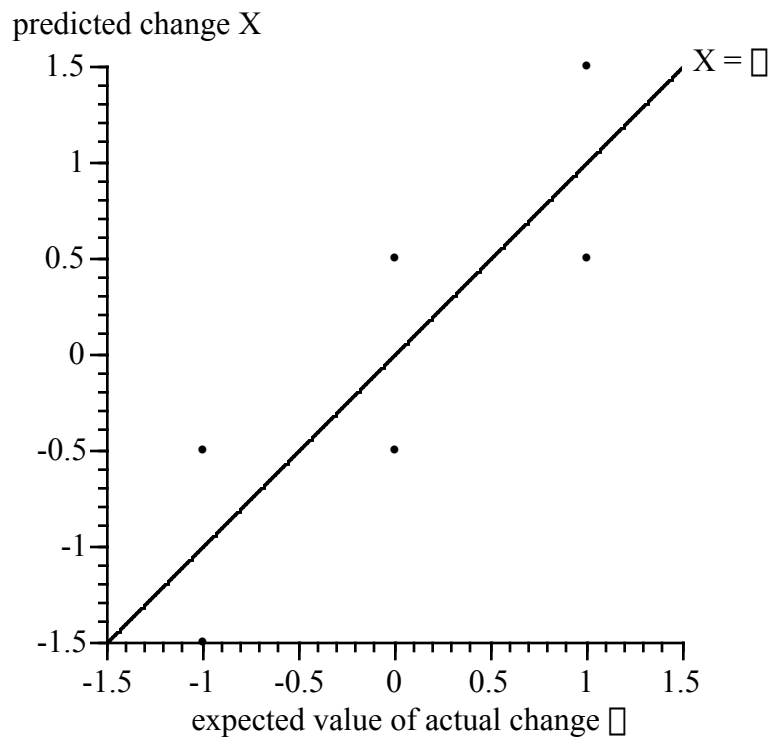


Figure 1 X as a function of \square

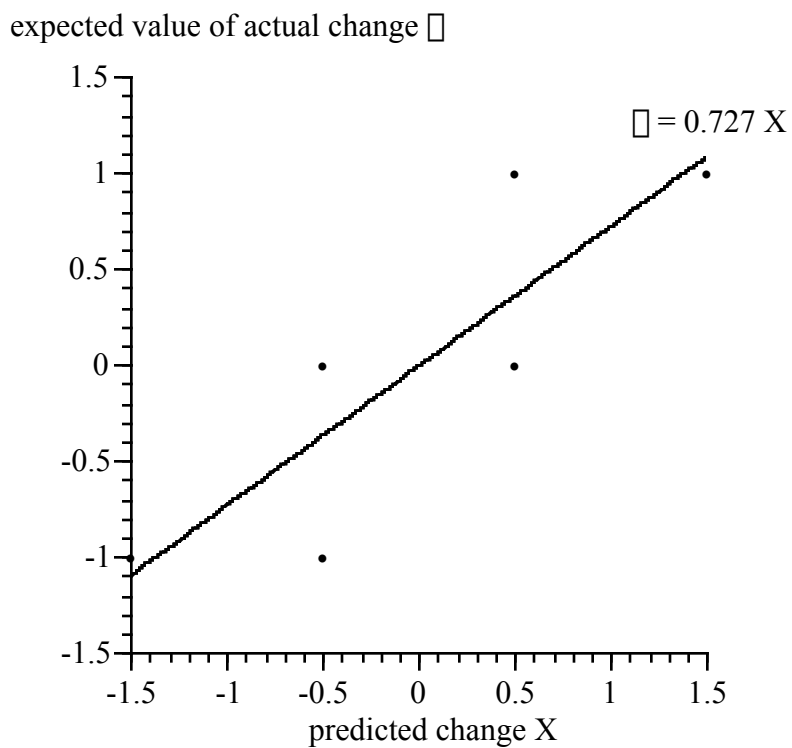


Figure 2 \square as a function of X