

A Simple Recursive Forecasting Model

William A. Branch
University of California, Irvine

George W. Evans*
University of Oregon

February 1, 2005

Abstract

We compare the performance of alternative recursive forecasting models. A simple constant gain algorithm, used widely in the learning literature, both forecasts well out of sample and also provides the best fit to the Survey of Professional Forecasters.

JEL classification codes: E37, D84, D83.

Keywords: constant gain, recursive learning, expectations.

1 Introduction

There is now an extensive literature which, in place of rational expectations (RE), assumes that agents behave as econometricians when making forecasts (see [9] for a systematic treatment and references). In particular, a growing number of theoretical and applied macro models replace RE with a linear forecasting rule, the parameters of which are updated by a constant gain version of least squares ([22, 9, 10, 4, 14, 17, 19, 20, 2, 5, 3]). Although theoretically motivated, an important empirical issue remains open in this literature: is constant gain learning a plausible bounded rationality assumption? To address this issue we answer two questions: do constant gain algorithms perform well in out-of-sample forecasting in comparison to other recursive forecasting models, and, does constant gain learning provide a good description of actual forecaster behavior?

We begin by setting out a general recursive forecasting model that allows the parameters to follow a random walk. As noted by [15] and [22], depending on the covariance structure of the stochastic components of the parameter law of motion, specific adaptive learning algorithms arise as special cases, including constant gain least squares. Time-varying parameter models of the economy have been estimated extensively before, and recently by [6] and [23]. Our principal contributions are two-fold. First, we estimate the optimal gains for the constant gain least squares model and show that it provides the best out of sample forecasts,

*Email: gevans@uoregon.edu; Department of Economics, 1285 University of Oregon, Eugene, OR 97403-1285; Phone: (541) 346-4662; Fax: (541) 346-1243.

compared to the alternatives, for inflation and GDP growth. Second, we show that a simple version of this model, namely constant gain least squares with equal gain for each variable, provides the best fit to the Survey of Professional Forecasters (SPF).

2 Forecasting Model

We estimate forecasting models of inflation and output growth that are based on recursive parameter updating¹, and then evaluate their performance in out-of-sample forecasts. The importance of out-of-sample forecasts, to protect against data-mining and to provide a stringent test of alternative models, has been stressed, e.g., in [13], [24] and [21]. As a second way to assess the models, we compare the forecasts to those of the SPF.

2.1 A Unified Framework

We follow [15] and [22] in specifying a general state space forecasting model that will nest our alternative models. Let $y_t = (y_{1,t}, \dots, y_{n,t})'$ denote the variables of interest. The economic law of motion takes the form

$$\begin{aligned} y_{j,t} &= b'_{j,t}x_t + \varepsilon_{j,t}, \quad j = 1, \dots, n, \\ b_{j,t} &= b_{j,t-1} + \theta_{j,t} \end{aligned}$$

where $b_{j,t}$ is the $(K \times 1)$ parameter vector and x_t is the $(K \times 1)$ vector of explanatory variables. $\varepsilon_{j,t}$ and $\theta_{j,t}$ are assumed to be zero mean mutually independent random sequences with $E\varepsilon_{j,t}\varepsilon'_{j,t} = R_{2,t}^j$ and $E\theta_{j,t}\theta'_{j,t} = R_{1,t}^j$.

We are interested in Vector Autoregression (VAR) specifications so that $x_t = (1, y_{t-1})'$.² Conditional forecasts of $y_{j,t}$ are constructed as

$$y_{j,t|t-1} = \hat{b}'_{j,t-1}x_t.$$

The forecasting problem is selection of an appropriate procedure for constructing the sequence $\{\hat{b}_t\}$. The parameter vectors can be estimated by the Kalman Filter recursion,³

$$\hat{b}_{j,t} = \hat{b}_{j,t-1} + L_{j,t} \left(y_{j,t} - \hat{b}'_{j,t-1}x_t \right) \quad (1)$$

$$L_{j,t} = \frac{P_{j,t-1}x_t}{R_{2,t}^j + x'_t P_{j,t-1}x_t} \quad (2)$$

$$P_{j,t} = P_{j,t-1} - \frac{P_{j,t-1}x_t x'_t P_{j,t-1}}{R_{2,t}^j + x'_t P_{j,t-1}x_t} + R_{1,t}^j, \quad (3)$$

¹Some authors have argued that forecasts based on real-time data sets, rather than revised data, are preferable for matching survey data. To keep the analysis as close to the existing literature as possible we use historical data. For a discussion of real-time data in forecasting see [7] and [16].

²For transparency, we focus on equation by equation estimation.

³The usual Kalman filter recursion is split into 'prediction' and 'transition' equations. It is straightforward to verify that for our assumed stochastic process the equations given represent the Kalman filter recursion.

where $\hat{b}_{j,t}$ has covariance matrix $P_{j,t} = E(b_{j,t} - \hat{b}_{j,t})(b_{j,t} - \hat{b}_{j,t})'$.

The properties of the estimator depend on the assumptions for $R_{1,t}^j, R_{2,t}^j$ and priors $\hat{b}_{j,0}, P_{j,0}$. [6] pursue a Bayesian estimation strategy where they characterize the posterior densities of b_j under very general assumptions on $R_{1,t}^j$. Our interest is in comparing the forecast performance of alternative assumptions on $R_{1,t}^j$ that yield simple recursive learning rules. We turn now to our specifications of $R_{1,t}^j$.

Recursive Least Squares. As noted by [15] and [22], under the restrictions

$$R_{1,t}^j = 0 \text{ and } R_{2,t}^j = 1, \text{ for } j = 1, \dots, n,$$

the Kalman Filter (1)-(3) is equivalent to recursive least squares (RLS),

$$\hat{b}_{j,t} = \hat{b}_{j,t-1} + \gamma_{j,t} R_t^{-1} x_t (y_{j,t} - \hat{b}_{j,t-1}' x_t) \quad (4)$$

$$R_t = R_{t-1} + \gamma_{j,t} (x_t x_t' - R_{t-1}). \quad (5)$$

where $\gamma_{j,t} = t^{-1}$. RLS is simply a recursive formulation of ordinary least squares (e.g. [9]).

Constant Gain Least Squares. As an alternative to RLS, [22] emphasizes the restrictions,

$$R_{1,t}^j = \frac{\gamma_j}{1 - \gamma_j} P_{j,t-1} \text{ and } R_{2,t}^j = 1 - \gamma_j,$$

and shows that it leads to the constant gain version of recursive least squares (4)-(5) in which $\gamma_{j,t} = \gamma_j$ where $0 < \gamma_j < 1$. The scalar γ_j , referred to as the “gain,” in this rule is constant over time in contrast to the “decreasing gain” t^{-1} of RLS. Constant gain least squares discounts past observations at geometric rate $1 - \gamma_j$ and is consequently more robust to structural change. We allow $\gamma_j \neq \gamma_{j'}$ for $j \neq j'$, which would be appropriate if variables have different rates of structural change.

Stationary Time-Varying Parameters. The restrictions on $R_{1,t}^j$ restrict the conditional distributions for the parameter vector. For RLS the parameters are non-stochastic and under constant gain least squares they may be heteroskedastic. As a third alternative, we follow [24] in specifying a time-varying parameter model estimated by the Kalman filter, with homoskedastic parameter error structure:

$$R_{1,t}^j = \lambda_j^2 \sigma_j^2 Q \text{ and } R_{2,t}^j = \sigma_j^2$$

where σ_j^2 is estimated as the variance from the residuals of an AR(1) model for y_j , $Q = (E x_t x_t')^{-1}$, and λ_j is a parameter for controlling the drift of the parameter estimates, which we again allow to differ across j . This approach is similar to [24] and [8] as formulated by [12]. For brevity we will refer to this as the “Kalman filter” learning rule.

A Simple Model. Finally, we also consider the special case of constant gain least squares in which we impose a common constant gain $\gamma_j = \gamma_{j'} = \gamma$ across variables. This model has the attraction of relative simplicity, which is, of course, one of the motivations of bounded rationality. More generally, parsimony is often recommended as one criterion for selecting between forecasting models ([13]).

3 Results

We focus on low order VARs. Although there are more sophisticated (and likely more accurate) ways to model the economy— for example, [6]— we are interested in simple, parsimonious forecasting models that are plausible descriptions of boundedly rational agents. We therefore examine the four formulations just presented using a bivariate VAR(1) in output growth and inflation.⁴ Output growth is measured in terms of quarterly real GDP calculated as $\log(GDP_t/GDP_{t-1})$. Inflation is measured similarly using the quarterly GDP deflator. These variables are chosen to stay close to the SPF, which asks a sample of professional forecasters, among other questions, for their expectation of one step ahead GDP and inflation.

Our approach is straightforward. We split the sample into three parts: the first, a pre-forecasting period, is the period in which prior beliefs are formed by estimating the VAR; the second part is an in-sample period in which the optimal parameters $\gamma_j, \lambda_j, j = 1, 2$ are determined (for RLS the gain sequence continues to be updated as t^{-1}); finally, the remainder is the out-of-sample forecasting period. We segment these samples so that the out-of-sample forecasting period corresponds to the SPF : 1981:3-2002:4. We choose a fairly long pre-forecasting period, 1947:2-1969:4, to avoid over-sensitivity of the initial in-sample estimates. The in-sample period is thus 1970:1-1981:2.

We envision professionals as forecasting throughout the sample. When they begin reporting their forecasts in 1981:3, their forecasting model has been shaped by experience, with the free parameters chosen optimally for fitting the in-sample period. Our metrics for whether a forecast method is a ‘reasonable’ model are mean square forecast error in out-of-sample forecasting and mean square deviations from the survey responses in the SPF.

3.1 Optimal In-sample Parameter Choices

To measure forecast success, we calculate the mean square forecast error

$$MSE(y_j) = \frac{1}{T} \sum_{t=t_0}^T (y_{j,t} - \hat{y}_{j,t})^2$$

where y_j is the j th component of y and $\hat{y}_{j,t}$ is the forecast of the j th component based on $t-1$ information. For constant gain least squares we determine the gain γ_j which minimizes the in-sample MSE by searching over all $\gamma_j \in (0, 1)$, $j = 1, 2$, with $t_0 = 1970.1$ and $T = 1982.2$. If the degree of structural change differs across series then we would expect the optimal gain to be different. Evidence of different structural change in inflation and GDP growth is provided by [6] and [1]. We choose the in-sample estimates of λ_j , for the Kalman filter rule, in the same way. The results for the “in-sample” period 1970:1-1981:2 are as follows:

⁴Our key qualitative results hold also for a VAR(4).

1970:1 - 1981:2	γ	λ
GDP growth	0.007	0.001
Inflation	0.062	0.021

The estimates of both γ and λ are larger for inflation than for GDP growth, presumably because the degree of structural change in GDP growth during the 1970s was considerably less than for inflation. (The theoretical link between degree of structural change and equilibrium gain in a self-referential model is investigated in [11].) An estimate of the optimal gain γ was also given in [20]. However, their estimate is based on calibration of a structural model. Our focus is on out-of-sample forecasting performance, to which we now turn.

3.2 Forecast Comparison

We now compare how the forecasting models perform in out-of-sample forecasting over 1981:3-2002:4. The top panel of Table 1 presents out-of-sample MSE results for RLS, constant gain and the Kalman filter, where the learning rule parameters have been set at their in-sample 1970:1-1981:2 estimates.⁵ Table 1 demonstrates that the optimal constant gain algorithm outperforms both RLS and the Kalman filter over the period 1981:3-2002:4. One possible explanation is that both parameter drift and conditional heteroskedasticity are present in the data, and that constant gain least squares captures both of these features. The top panel of Table 1 also shows that the forecasting gains are non-trivial, from about 2.9-4.3% in the case of GDP growth and 5.6-12.5% for inflation forecasting.

INSERT TABLE 1 HERE

These results have two implications for theoretical and applied models. First, constant gain least squares, with each gain optimally tuned to track the variable of interest, does appear plausible as a model of bounded rationality given its satisfactory out-of sample performance. Second, in line with [6], the results suggest that both parameter drift and stochastic volatility are components of the stochastic structure of the economy.

3.3 Comparisons to the Survey of Professional Forecasters

We now compare the forecasts from the optimal constant gain, RLS, Kalman filter, and simple constant gain models to those from the SPF. The learning literature often motivates bounded rationality on the basis of assuming that forecasters behave like econometricians, and thus one test of the plausibility of constant gain models of adaptive learning is their ability to fit professionals' forecasts.

⁵The [24] approach sets $R_{1,t}^j = \lambda^2 \sigma^2 Q$ where $Q = (Ex_t x_t')^{-1}$. Stock and Watson compare $\lambda = .0025, .0075, .015, .01, .02$. We instead set λ_j according to in-sample forecasting success.

For comparison with the SPF we calculate the Mean Square Comparison Error (MSCE)

$$MSCE(y_j) = \frac{1}{T} \sum_{t=1}^T (y_{j,t}^{SPF} - \hat{y}_{j,t})^2$$

where y_j^{SPF} is the median survey forecast of the j th component of y . We take the same optimal parameters $\gamma_j, \lambda_j, j = 1, 2$ estimated for the in-sample period 1970:1-1981:2, and compare whether the forecasts over 1981:3-2002:4 are consistent with the SPF over this period. We also assess the simple model in which a common constant gain γ is imposed for both variables. For the simple constant gain model we choose γ to be average of the optimal gains for the two variables in the in-sample period, i.e. $\gamma = .0345$.

The bottom panel of Table 1 reports the results. RLS is clearly dominated by the other forecast methods that allow for parameter drift. For GDP growth the optimal constant gain method fits more closely than the Kalman filter model, whereas the opposite is true for inflation forecasts. Most strikingly, however, the simple model with common gain $\gamma = .0345$ fits best across all models for *both* variables. This provides significant support for constant gain learning as a model of actual forecasting.

3.4 Discussion

Figure 1 provides additional information on the performance of constant gain least squares in the 1981:3 - 2002:4 period. The left two curves show the MSE, for constant gain forecasts of GDP growth and inflation, and the right two curves show the corresponding MSCE, measuring the closeness of fit of constant gain forecasts to the SPF forecasts. In each case the curves plot mean square error as a function of different values of the gain.

INSERT FIGURE 1 HERE

Several features stand out. First, the optimal gains for forecasting both variables, given by the gains at which the curves are minimized, are higher during 1981:3 - 2002:4 than for the 1970:1-1981:2 period, and the optimal gain in the later as well as the earlier period is again higher for inflation than for GDP growth. However, given the flatness of the MSE for inflation forecasts, a common gain of $\gamma = 0.05$ produces reasonable forecasts.⁶ Second, for explaining SPF forecasts, the best fitting gain for inflation is substantially below those for fitting actual data in the period, and a common constant gain in the range $\gamma = .03$ to $\gamma = .04$ provide the best fit for the SPF.

Overall our results provide empirical support for the recent “learning theory” approach in which agents are modeled as employing constant gain least squares forecasting rules, with the gain chosen to reflect the degree of structural change. Our results complement those of [20], who use a calibrated structural macroeconomic model to fit and analyze the role of stabilization policy in the US, and of [18] who estimates econometrically a DSGE model with learning and provides evidence that constant gain learning is a major source of persistence.

⁶If the two MSEs are weighted equally the optimal common gain for this period is $\gamma = .047$.

4 Conclusion

Constant gain least squares is becoming widely used in learning theory approaches to macroeconomics that replace RE with econometric forecasting rules. Using constant gain is a simple way of modifying least squares to adapt to ongoing structural change. We have examined the empirical performance of this learning rule, compared to recursive least squares and a more conventional time-varying parameter model, in forecasting US GDP growth and inflation. Our results are very encouraging. Constant gain least squares, with the gain optimally chosen in-sample for each variable, performs best in out-of-sample forecasting. An even simpler version, in which a common constant gain is used for both variables, provides the best fit for the Survey of Professional Forecasters.

While the most accurate forecasting model will depend on the stochastic nature of the economy, constant gain least squares does appear plausible as a model of adaptive learning for economic agents. Since the degree of structural change itself seems to change over time, it would be valuable in future work to explore empirical models in which agents gradually adapt the gain over time in response to observed data.⁷

References

- [1] Branch, W. A., and G. W. Evans, 2004, "Model Uncertainty and Endogenous Volatility," mimeo.
- [2] Bullard, J. and I.-K. Cho, 2002, "Escapist Policy Rules," Working Paper 20002-022C, FRB St. Louis, forthcoming *Journal of Economic Dynamics and Control*.
- [3] Bullard, J. and J. Duffy, 2004, "Learning and Structural Change in Macroeconomic Data," Working Paper 2004-016A, Federal Reserve Bank of St. Louis.
- [4] Cho, I.-K., N. Williams, and T. J. Sargent, 2002, "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1-40.
- [5] Cho, I.-K. and K. Kasa, 2002, "Learning Dynamics and Endogenous Currency Crises," mimeo., University of Illinois.
- [6] Cogley, T., and T. J. Sargent, "Drifts and Volatilities: Monetary Policy and Outcomes in the Post WWII U.S.," forthcoming *Review of Economic Dynamics*.
- [7] Croushore, D. and T. Stark, 2002, "Forecasting With a Real-time Data Set for Macroeconomists," *Journal of Macroeconomics*, 24.
- [8] Doan, T., R. B. Litterman, and C. A. Sims, 1984, "Forecasting and Conditional Projection Using Realistic Prior Distributions," *Econometric Reviews*, 3, 1-100.

⁷[11] provide a suitable algorithm for updating the gain in a univariate setting.

- [9] Evans, G. W., and S. Honkapohja, 2001, *Learning and Expectations in Macroeconomics*, Princeton University Press, Princeton, NJ.
- [10] Evans, G. W., and S. Honkapohja, 2004, "Policy Interaction, Expectations and the Liquidity Trap," *Review of Economic Dynamics*, in press.
- [11] Evans, G. W., and G. Ramey, 2003. "Adaptive Expectations, Underparameterization and the Lucas Critique," University of Oregon Working Paper No. 237, forthcoming *Journal of Monetary Economics*.
- [12] Hamilton, J. D., 1994, *Time Series Analysis*, Princeton University Press Princeton, NJ.
- [13] Harvey, A., 1990, *The Econometric Analysis of Time Series*, 2nd Ed., Philip Allan, London.
- [14] Kasa, K., "Learning, Large Deviations, and Recurrent Currency Crises," *International Economic Review*, 45, 141-173.
- [15] Ljung, L., and T. Soderstrom, 1983, *Theory and Practice of Recursive Identification*, MIT Press.
- [16] Kozicki, S., and P. A. Tinsley, 2002, "Comments on 'Forecasting with a Real-time Data Set for Macroeconomists,'" *Journal of Macroeconomics*, 24.
- [17] McGough, B., 2004, "Shocking Escapes," mimeo., Oregon State University, forthcoming *Economic Journal*.
- [18] Milani, F., 2004, "Expectations, Learning and Macroeconomic Persistence," mimeo., Princeton University.
- [19] Orphanides, A., and J. C. Williams, 2005, "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," Chapter 5 in *Inflation Targeting*, ed. Ben S. Bernanke and Michael Woodford, NBER and University of Chicago Press, Chicago.
- [20] Orphanides, A., and J. C. Williams, 2003, "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations," FRBSF 2003-24.
- [21] Pesaran, M. H., and A. Timmermann, 1995, "Predictability of Stock Returns: Robustness and Economic Significance," *Journal of Finance*, 50, 4, 1201-1228.
- [22] Sargent, T. J., 1999, *The Conquest of American Inflation*, Princeton University Press, Princeton, NJ.
- [23] Sims, C. A., and T. Zha, 2004, "Were There Regime Switches in U.S. Monetary Policy?," mimeo.
- [24] Stock, J. H., and M. W. Watson, 1996, "Evidence on Structural Instability in Macroeconomic Time Series Relations," *Journal of Business and Economic Statistics*, 14, 11-30.
- [25] Williams, N., 2002, "Escape Dynamics in Learning Models," mimeo., Princeton University.

Out-of-sample period: 1981:3-2002:4		
Mean Square Forecast Error:		
	GDP Growth	Inflation
RLS	.3507	.0532
Constant Gain	.3408	.0503
Kalman filter	.3558	.0570
log(RLS/Constant Gain)	2.86%	5.61%
log(Kalman filter/Constant Gain)	4.31%	12.50 %
Mean Square SPF Error:		
	GDP Growth	Inflation
RLS	.1354	.0316
Constant Gain	.1204	.0266
Kalman filter	.1258	.0259
Simple Model	.1181	.0250
log(RLS/Simple Model)	13.67%	23.43%
log(Const. Gain/Simple Model)	1.93%	6.20 %
log(Kalman filter/ Simple Model)	6.32%	3.54%

Table 1. Comparisons of fit across forecasting models. Learning rule parameters have been set at their 1970:1 -1981:2 in-sample estimates, with the Simple Model gain set at .0345. The top panel of the table is the mean square forecast error in out-of-sample forecasting of actual inflation and GDP. The bottom panel is the mean square deviations of the forecasting models from the median response of the SPF.




Figure 1. Mean square comparison and forecast errors as a function of the gain. Left panels forecast out-of-sample 1983:3-2002:4, right panels forecast the SPF.