

OPTIMAL PREDICTION UNDER ASYMMETRIC LOSS

PETER F. CHRISTOFFERSEN
International Monetary Fund

FRANCIS X. DIEBOLD
University of Pennsylvania and NBER

Prediction problems involving asymmetric loss functions arise routinely in many fields, yet the theory of optimal prediction under asymmetric loss is not well developed. We study the optimal prediction problem under general loss structures and characterize the optimal predictor. We compute the optimal predictor analytically in two leading tractable cases and show how to compute it numerically in less tractable cases. A key theme is that the conditionally optimal forecast is biased under asymmetric loss and that the conditionally optimal amount of bias is time varying in general and depends on higher order conditional moments. Thus, for example, volatility dynamics (e.g., GARCH effects) are relevant for optimal *point* prediction under asymmetric loss. More generally, even for models with linear conditional-mean structure, the optimal point predictor is in general nonlinear under asymmetric loss, which provides a link with the broader nonlinear time series literature.

1. INTRODUCTION

A moment's reflection yields the insight that prediction problems involving asymmetric loss structures arise routinely because a myriad of situation-specific factors may render positive errors more (or less) costly than negative errors. The potential necessity of allowing for asymmetric loss has long been acknowledged. Granger and Newbold (1986, p. 125), for example, noted that although "an assumption of symmetry about the conditional mean . . . is likely to be an easy one to accept, . . . an assumption of symmetry for the cost function is much less acceptable." Practitioners routinely echo this sentiment (e.g., Stockman, 1987).

This paper is a revised and shortened version of the first part of Christoffersen and Diebold (1994), which may be consulted for additional results, discussion, and examples. We benefited from constructive comments from the editor and two referees, and from Adolf Buse, Clive Granger, Hashem Pesaran, Dale Poirier, Enrique Sentana, Bob Stine, Jim Stock, and Ken Wallis. Helpful discussion was also provided by participants at the UCSD Conference on Multivariate Financial Econometrics, the NBER/NSF Forecasting Seminar, the World Congress of the Econometric Society, and numerous university seminars. Remaining inadequacies are ours alone. We thank the National Science Foundation, the Sloan Foundation, the University of Pennsylvania Research Foundation, and the Barbara and Edward Netter Fellowship for support. Address correspondence to: Francis X. Diebold, Department of Economics, University of Pennsylvania, Philadelphia, PA 19104-6297, USA.

In this paper we treat the prediction problem under general loss structures, building on the classic work of Granger (1969). In Section 2, we characterize the optimal predictor for non-Gaussian processes under asymmetric prediction-error loss. The results apply, for example, to important classes of conditionally heteroskedastic processes. In Section 3, we provide analytic solutions for the optimal predictor under two popular analytically tractable asymmetric loss functions. In Section 4, we provide methods for approximating the optimal predictor under more general loss functions. We conclude in Section 5.

2. OPTIMAL PREDICTION FOR NON-GAUSSIAN PROCESSES

Granger (1969) studied Gaussian processes and showed that under asymmetric loss the optimal predictor is the conditional mean plus a constant bias term. Granger's fundamental result, however, has two key limitations. First, the Gaussian assumption implies a constant conditional prediction-error variance. This is unfortunate because conditional heteroskedasticity is widespread in economic and financial data. Second, the loss function must be of prediction-error form; that is, $L(y_{t+h}, \hat{y}_{t+h}) = L(y_{t+h} - \hat{y}_{t+h}) = L(e_{t+h})$, where y_{t+h} is the h -step-ahead realization, \hat{y}_{t+h} is the h -step-ahead forecast (made at time t), and e_{t+h} is the corresponding forecast error. More general functions of realizations and predictions are excluded.

Let us begin by generalizing Granger's result to allow for conditional variance dynamics. We achieve this most simply by working in a conditionally Gaussian, but not necessarily unconditionally Gaussian, environment, with prediction-error loss. Subsequently, we shall allow for both conditional nonnormality and more general loss functions.

PROPOSITION 1. *If $y_{t+h}|\Omega_t \sim N(\mu_{t+h|t}, \sigma_{t+h|t}^2)$ is a conditionally Gaussian process and $L(e_{t+h})$ is any loss function defined on the h -step-ahead prediction error e_{t+h} , then the optimal predictor is of the form $\hat{y}_{t+h} = \mu_{t+h|t} + \alpha_{t+h|t}$, where $\alpha_{t+h|t}$ depends only on the loss function and the conditional prediction-error variance $\sigma_{t+h|t}^2 = \text{var}(y_{t+h}|\Omega_t) = \text{var}(e_{t+h}|\Omega_t)$.*

Proof. See Appendix.

The optimal predictor under conditional normality is not necessarily just a constant added to the conditional mean because the conditional prediction-error variance may be time varying. Conditionally Gaussian GARCH processes, for example, fall under the jurisdiction of Proposition 1 for $h = 1$. Thus, under asymmetric loss, conditional variance dynamics are important not only for interval prediction but also for *point* prediction. If loss is asymmetric but conditional, heteroskedasticity is ignored; the resulting point predictions will be suboptimal and may have dramatically greater conditionally expected loss in consequence.

The result of Proposition 1 that the "adjustment factor" depends only on the conditional variance depends crucially on conditional normality. We can dis-

pense with conditional normality and still obtain a sharp result, however, which is a straightforward extension of Proposition 1.

PROPOSITION 2. *If $y_{t+h}|\Omega_t$ has conditional mean $\mu_{t+h|t}$ and a vector of (possibly time varying) conditional moments of order two and higher $\lambda_{t+h|t}$, and $L(e_{t+h})$ is any loss function defined on the h -step-ahead prediction error e_{t+h} , then the optimal predictor is of the form $\hat{y}_{t+h} = \mu_{t+h|t} + \alpha_{t+h|t}$, where $\alpha_{t+h|t}$ depends only on the loss function and $\lambda_{t+h|t}$.*

Proof. See Appendix.

Note, however, that although Proposition 2 does not require a Gaussian process, it does require prediction-error loss. In Section 4, we will also relax that assumption. Note also that Proposition 2 holds for other measures of location, such as the conditional median.

3. ANALYTIC SOLUTIONS UNDER LINEX AND LINLIN LOSS

We examine two asymmetric loss functions (“linex” and “linlin”) for which it is possible to solve analytically for the optimal predictor. To maintain continuity of exposition, we work throughout this section with the conditionally Gaussian process $y_{t+h}|\Omega_t \sim N(\mu_{t+h|t}, \sigma_{t+h|t}^2)$.¹ For each loss function, we characterize the optimal predictor, $\hat{y}_{t+h} = \mu_{t+h|t} + \alpha_{t+h|t}$, and we compare its conditionally expected loss with that of two competitors, the conditional mean $\mu_{t+h|t}$ and the pseudo-optimal predictor $\hat{y}_{t+h} = \mu_{t+h|t} + \alpha_h$, where α_h depends only on the loss function and the unconditional prediction-error variance $\sigma_h^2 = \text{var}(e_{t+h})$. The optimal predictor acknowledges loss asymmetry and the possibility of conditional heteroskedasticity through a possibly *time-varying* adjustment to the conditional mean. The conditional mean, in contrast, is always suboptimal because it incorporates *no* adjustment. The pseudo-optimal predictor is intermediate in that it incorporates only a *constant* adjustment for asymmetry; thus, it is fully optimal only in the conditionally homoskedastic case $\sigma_{t+h|t}^2 = \sigma_h^2, \forall t, h$.

3.1. Linex Loss

The linex loss function, introduced by Varian (1974) and used by Zellner (1986), is

$$L(x) = b[\exp(ax) - ax - 1], \quad a \in \mathbb{R} \setminus \{0\}, \quad b \in \mathbb{R}_+.$$

It is so named because when $a > 0$, loss is approximately linear to the left of the origin and approximately exponential to the right, and conversely when $a < 0$. The optimal h -step-ahead predictor under linex loss solves

$$\min_{\hat{y}_{t+h}} E_t\{b[\exp(a(y_{t+h} - \hat{y}_{t+h})) - a(y_{t+h} - \hat{y}_{t+h}) - 1]\}.$$

Differentiating and using the conditional moment-generating function for a conditionally Gaussian variate, we obtain $\hat{y}_{t+h} = \mu_{t+h|t} + (a/2)\sigma_{t+h|t}^2$. Similar calculations reveal that the pseudo-optimal predictor is $\hat{y}_{t+h} = \mu_{t+h|t} + (a/2)\sigma_h^2$, where $\sigma_h^2 = \text{var}(e_{t+h})$ is the unconditional h -step-ahead prediction-error variance.

Proposition 1 shows that the optimal predictor under conditional normality is the conditional mean plus a function of the conditional prediction-error variance. Under linex loss, the function is a simple linear one, depending on the degree of asymmetry of the loss function, as captured in the parameter a .² The reason is simple—when a is positive, for example, positive prediction errors are more devastating than negative errors, and so a negative conditionally expected error is desirable. The optimal amount of bias depends on the conditional prediction-error variance of the process; as it grows, so too does the optimal amount of bias to avoid large positive prediction errors. Effectively, optimal prediction under asymmetric loss corresponds to conditional-mean prediction of a transformed series, where the transformation reflects both the loss function and the higher-order conditional moments of the original series. For example, the optimal predictor of y_{t+h} under conditional normality and linex loss, $\hat{y}_{t+h} = \mu_{t+h|t} + (a/2)\sigma_{t+h|t}^2$, is the conditional mean of x_{t+h} , where $x_{t+h} = y_{t+h} + (a/2)\sigma_{t+h|t}^2$.³

Inserting the optimal, pseudo-optimal, and conditional mean predictors into the conditionally expected loss expression, we see that the conditionally expected linex losses are $ba^2\sigma_{t+h|t}^2/2$, $b[\exp(a^2(\sigma_{t+h|t}^2 - \sigma_h^2)/2) + a^2\sigma_h^2/2 - 1]$, and $b[\exp(a^2\sigma_{t+h|t}^2/2) - 1]$, respectively. By construction, the conditionally expected loss of the optimal predictor is less than or equal to that of any other predictor. Interestingly, however, it is not possible to rank the pseudo-optimal as superior to the conditional mean predictor. Tedious but straightforward algebra reveals that, for sufficiently small values of $\sigma_{t+h|t}^2$ (depending nonlinearly on the values of a and σ_h^2), the conditionally expected loss of the conditional mean will be smaller than that of the pseudo-optimal predictor. In very low volatility times, the conditionally optimal amount of bias is very small, resulting in a lower conditionally expected loss for the conditional mean than for the pseudo-optimal predictor, the bias of which is optimal in “average” times but too high in low-volatility times.

The situation is illustrated in Figure 1, in which we plot conditionally expected linex loss as a function of $\sigma_{t+h|t}^2$ for each of the three predictors. The conditionally expected loss of the optimal predictor is linear in $\sigma_{t+h|t}^2$ and is of course always lowest. The losses of the pseudo-optimal and the optimal predictors coincide when $\sigma_{t+h|t}^2 = \sigma_h^2 = 1$. As $\sigma_{t+h|t}^2$ falls below σ_h^2 , the loss of the conditional mean intersects the loss of the pseudo-optimal predictor from above. As $\sigma_{t+h|t}^2$ gets close to zero, the optimal predictor incorporates progressively smaller corrections to the conditional mean, and so the conditionally expected losses of the optimal and conditional mean predictors coincide.

3.2. Linlin Loss

The linlin loss function,

$$L(y_{t+h} - \hat{y}_{t+h}) = \begin{cases} a|y_{t+h} - \hat{y}_{t+h}|, & \text{if } (y_{t+h} - \hat{y}_{t+h}) > 0 \\ b|y_{t+h} - \hat{y}_{t+h}|, & \text{if } (y_{t+h} - \hat{y}_{t+h}) \leq 0, \end{cases}$$

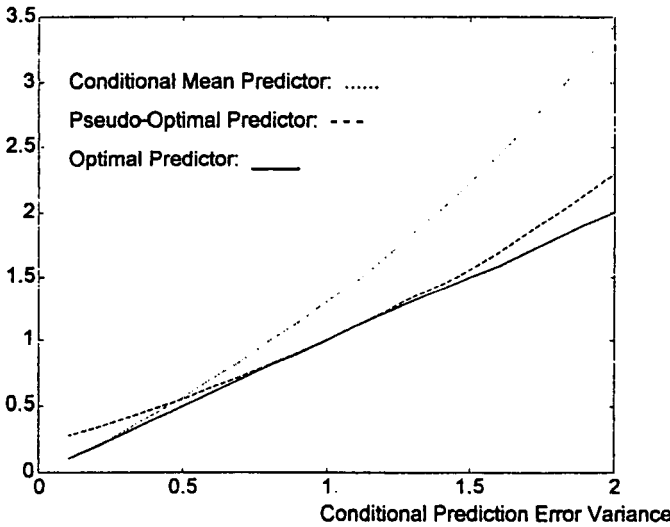


FIGURE 1. Conditionally expected linear loss of conditional mean, pseudo-optimal, and optimal predictors. The linear loss parameters are set to $a = 1$ and $b = 2$. The unconditional variance is fixed at 1.

so called because of its linearity on each side of the origin, was used by Granger (1969) and is the loss function underlying quantile regression. The optimal predictor solves

$$\min_{\hat{y}_{t+h}} \left[a \int_{\hat{y}_{t+h}}^{\infty} (y_{t+h} - \hat{y}_{t+h}) f(y_{t+h} | \Omega_t) dy_{t+h} - b \int_{-\infty}^{\hat{y}_{t+h}} (y_{t+h} - \hat{y}_{t+h}) f(y_{t+h} | \Omega_t) dy_{t+h} \right].$$

The first-order condition is

$$-a(1 - F(\hat{y}_{t+h} | \Omega_t)) + bF(\hat{y}_{t+h} | \Omega_t) = 0,$$

which is equivalent to

$$F(\hat{y}_{t+h} | \Omega_t) = \frac{a}{a + b},$$

where $F(y_{t+h} | \Omega_t)$ is the conditional cumulative distribution function (c.d.f.) of y_{t+h} and $f(y_{t+h} | \Omega_t)$ is the conditional density of y_{t+h} .

In the conditionally Gaussian case, we have from Proposition 1 that

$$\begin{aligned} F(\hat{y}_{t+h} | \Omega_t) &= \Pr(y_{t+h} \leq (\mu_{t+h|t} + \alpha_{t+h}) | \Omega_t) \\ &= \Pr\left(\frac{y_{t+h} - \mu_{t+h|t}}{\sigma_{t+h|t}} \leq \frac{\alpha_{t+h}}{\sigma_{t+h|t}} \mid \Omega_t\right) = \Phi\left(\frac{\alpha_{t+h}}{\sigma_{t+h|t}}\right) = \frac{a}{a + b}, \end{aligned}$$

where $\Phi(z)$ is the $N(0, 1)$ c.d.f. It follows that the conditionally optimal amount of bias is

$$\alpha_{t+h|t} = \sigma_{t+h|t} \Phi^{-1} \left(\frac{a}{a+b} \right),$$

so that⁴

$$\hat{y}_{t+h} = \mu_{t+h|t} + \sigma_{t+h|t} \Phi^{-1} \left(\frac{a}{a+b} \right).$$

Similar calculations reveal that the pseudo-optimal predictor is

$$\hat{y}_{t+h} = \mu_{t+h|t} + \sigma_h \Phi^{-1} \left(\frac{a}{a+b} \right).$$

Now let us compute conditionally expected linlin loss for the optimal, pseudo-optimal, and conditional mean predictors. Recall the formulae for the truncated expectation,

$$E_t\{y_{t+h} | (y_{t+h} > \hat{y}_{t+h})\} = \frac{\int_{\hat{y}_{t+h}}^{\infty} y_{t+h} f(y_{t+h} | \Omega_t) dy_{t+h}}{1 - F(\hat{y}_{t+h} | \Omega_t)},$$

$$E_t\{y_{t+h} | (y_{t+h} < \hat{y}_{t+h})\} = \frac{\int_{-\infty}^{\hat{y}_{t+h}} y_{t+h} f(y_{t+h} | \Omega_t) dy_{t+h}}{F(\hat{y}_{t+h} | \Omega_t)},$$

and substitute them into the expected loss expression to obtain

$$E_t\{L(y_{t+h} - \hat{y}_{t+h})\} = a(1 - F(\hat{y}_{t+h} | \Omega_t)) [E_t\{y_{t+h} | (y_{t+h} > \hat{y}_{t+h})\} - \hat{y}_{t+h}] \\ - bF(\hat{y}_{t+h} | \Omega_t) [E_t\{y_{t+h} | (y_{t+h} < \hat{y}_{t+h})\} - \hat{y}_{t+h}].$$

But under conditional normality,

$$E_t\{y_{t+h} | (y_{t+h} > \hat{y}_{t+h})\} = \mu_{t+h|t} + \sigma_{t+h|t} \frac{\phi(\xi_{t+h|t})}{1 - \Phi(\xi_{t+h|t})}$$

and

$$E_t\{y_{t+h} | (y_{t+h} < \hat{y}_{t+h})\} = \mu_{t+h|t} - \sigma_{t+h|t} \frac{\phi(\xi_{t+h|t})}{\Phi(\xi_{t+h|t})},$$

where

$$\xi_{t+h|t} = \frac{\hat{y}_{t+h} - \mu_{t+h|t}}{\sigma_{t+h|t}},$$

and $\phi(\cdot)$ is the $N(0, 1)$ probability density function (p.d.f.). Substituting into the conditionally expected loss expression, we obtain (after some algebraic manipulation)

$$E_t\{L(y_{t+h} - \hat{y}_{t+h})\} = (a + b)\sigma_{t+h|t}\phi(\xi_{t+h|t}) - a(\hat{y}_{t+h|t} - \mu_{t+h|t}) + (a + b)\Phi(\xi_{t+h|t})(\hat{y}_{t+h|t} - \mu_{t+h|t}).$$

For the optimal predictor,

$$\xi_{t+h|t} = \Phi^{-1}\left(\frac{a}{a + b}\right),$$

yielding an expected loss of $(a + b)\sigma_{t+h|t}\phi(\Phi^{-1}(a/(a + b)))$. For the pseudo-optimal predictor,

$$\xi_{t+h|t} = \Phi^{-1}\left(\frac{a}{a + b}\right)\frac{\sigma_h}{\sigma_{t+h|t}},$$

yielding an expected loss of

$$(a + b)\sigma_{t+h|t}\phi\left(\Phi^{-1}\left(\frac{a}{a + b}\right)\frac{\sigma_h}{\sigma_{t+h|t}}\right) - a\Phi^{-1}\left(\frac{a}{a + b}\right)\sigma_h + (a + b)\Phi\left(\Phi^{-1}\left(\frac{a}{a + b}\right)\frac{\sigma_h}{\sigma_{t+h|t}}\right)\Phi^{-1}\left(\frac{a}{a + b}\right)\sigma_h.$$

For the conditional mean predictor, $\xi_{t+h|t} = 0$, yielding an expected loss of $(a + b)\sigma_{t+h|t}/\sqrt{2\pi}$. Qualitatively, the situation is identical to that shown in Figure 1 for the linex case.

4. APPROXIMATING THE OPTIMAL PREDICTOR

The preceding analytic results rely on simple loss functions. In general, however, it is not possible to solve analytically for the optimal predictor. In this section, we develop an approximately optimal predictor via series expansions. The approach is of interest because it frees us from two potentially restrictive assumptions—conditional normality and prediction-error loss.

For the moment, maintain the conditional normality assumption and assume that the optimal predictor exists and is unique,

$$\hat{y}_{t+h} = G(\mu_{t+h|t}, \sigma_{t+h|t}^2),$$

where $G(\cdot, \cdot)$ is at least twice continuously differentiable. Then we can take a second-order Taylor series expansion around the unconditional (and time invariant) moments μ_h and σ_h^2 ,

$$\hat{y}_{t+h} \approx G(\mu_h, \sigma_h^2) + G'(\mu_h, \sigma_h^2)\begin{pmatrix} \mu_{t+h|t} - \mu_h \\ \sigma_{t+h|t}^2 - \sigma_h^2 \end{pmatrix} + \frac{1}{2}(\mu_{t+h|t} - \mu_h, \sigma_{t+h|t}^2 - \sigma_h^2)G''(\mu_h, \sigma_h^2)\begin{pmatrix} \mu_{t+h|t} - \mu_h \\ \sigma_{t+h|t}^2 - \sigma_h^2 \end{pmatrix}.$$

Rewrite this as

$$\hat{y}_{t+h} \approx \beta_0 + \beta_1\mu_{t+h|t} + \beta_2\sigma_{t+h|t}^2 + \beta_3(\mu_{t+h|t})^2 + \beta_4(\sigma_{t+h|t}^2)^2 + \beta_5(\mu_{t+h|t}\sigma_{t+h|t}^2) \equiv y_{t+h}^*(\beta),$$

where $\beta = (\beta_0, \beta_1, \dots, \beta_5)'$ and $\beta_i = H_i(\mu_h, \sigma_h^2)$, $i = 0, 1, \dots, 5$. Because the function $G(\cdot, \cdot)$ is generally unknown, so too are the $H(\cdot, \cdot)$ functions. But $\mu_{t+h|t}$ and $\sigma_{t+h|t}^2$ are known, and the minimization that defines $\hat{\beta}_N$ can be done over a very long simulated realization of length N ; that is,

$$\hat{\beta}_N = \operatorname{argmin}_{\beta \in B} \sum_{t=1}^N L(y_{t+h}, y_{t+h}^*(\beta)).$$

Under regularity conditions given in the Appendix, the following proposition is immediate.

PROPOSITION 3. *As $N \rightarrow \infty$, $y_{t+h}^*(\hat{\beta}_N) \rightarrow y_{t+h}^*(\beta_0)$, where $y_{t+h}^*(\beta_0)$ is the best predictor within the $y_{t+h}^*(\cdot)$ family, with respect to the metric $L(\cdot, \cdot)$.*

Proof. See Appendix.

A number of remarks are in order. First, the h -step-ahead conditional expectation and the corresponding conditional variance may be computed conveniently by using the Kalman filter recursions. Second, if loss is in fact of prediction-error form, $L(e_{t+h})$, one may set $\beta_1 = 1$ and $\beta_3 = \beta_5 = 0$ a priori, due to Proposition 1. Third, it is clear that higher-order expansions in $\mu_{t+h|t}$ and $\sigma_{t+h|t}^2$ may be entertained and may lead to improvements. Fourth, conditional nonnormality may be handled with expansions involving more than the first two conditional moments (e.g., involving conditional skewness and kurtosis). Fifth, and related, parametric economy can be achieved in conditionally non-Gaussian cases by using the autoregressive conditional density framework of Hansen (1994). Hansen's framework exploits parametric conditional mean and variance functions but allows for higher-order conditional dynamics by letting the normalized variable $z_{t+h}(\theta) = (y_{t+h} - \mu_{t+h|t}(\theta)) / \sigma_{t+h|t}(\theta)$ follow a distribution with possibly time-varying "shape" parameters, such as a t -distribution with time-varying degrees of freedom (and variance standardized to 1). Sixth, in both the conditionally Gaussian and conditionally non-Gaussian cases, one is of course not limited to series expansions; other nonparametric functional estimators may be used.

5. SUMMARY AND CONCLUDING REMARKS

This paper is part of a research program aimed at allowing for general loss structures in estimation, model selection, prediction, and forecast evaluation. Recently, a number of authors have made progress toward that goal, including Weiss (1996) on estimation, Phillips (1996) on model selection, and Diebold and Mariano (1995) on forecast evaluation. We have focused on prediction and have analyzed the optimal prediction problem under asymmetric loss. We have computed the optimal predictor analytically in two leading tractable cases and showed how to compute it numerically in less tractable cases. Christoffersen and Diebold (1996) presented extensions, and they provided an illustration in the context of point prediction of GARCH processes under asymmetric loss.

A key emergent theme is that the conditionally optimal forecast is biased and that the conditionally optimal amount of bias is time varying in general and de-

depends on higher-order conditional moments. Thus, even for models with linear conditional-mean structure, the optimal predictor is in general nonlinear, thereby providing a link with the broader nonlinear time series literature.

Interestingly, some important recent work in dynamic economic theory is very much linked to the idea of prediction under asymmetric loss discussed here. Building on Whittle (1990), Hansen, Sargent, and Tallarini (1993) set up and motivated a general-equilibrium economy with “risk-sensitive” preferences, resulting in equilibria with certainty-equivalence properties.⁵ Thus, the prediction and decision problems may be done sequentially—but prediction is done with respect to a distorted probability measure that yields predictions different from the conditional mean.

NOTES

1. As will be made clear, however, although conditional normality is crucial to our derivation of the optimal predictor under linex loss, it may readily be discarded under linlin loss.

2. As $a \rightarrow 0$ the conditionally optimal amount of bias approaches zero. Quadratic loss obtains as $a \rightarrow 0$ because if a is small one can replace the exponential part of the loss function by the first two terms of its Taylor series expansion, yielding the approximation $L(x) \propto x^2$.

3. Because y_{t+h} is conditionally normal with $E[y_{t+h}|\Omega_t] = \mu_{t+h|t}$, x_{t+h} is conditionally normal with $E[x_{t+h}|\Omega_t] = \mu_{t+h|t} + (a/2)\sigma_{t+h|t}^2 = \hat{y}_{t+h}$. This insight is very much related to the computations associated with the risk-sensitive optimal control of Whittle (1990) and Hansen et al. (1993); see also the summary and concluding remarks at the end of this paper.

4. With linlin loss (in contrast to linex loss), it is very easy, even for non-Gaussian conditional distributions, to find the optimal predictor—just draw the conditional c.d.f. and read the value on the x -axis corresponding to $a/(a+b)$. More formally, $\hat{y}_{t+h} = F^{-1}(a/(a+b)|\Omega_t)$, and so \hat{y}_{t+h} is simply the $(a/(a+b))$ th conditional quantile. When $a = b$, of course, \hat{y}_{t+h} is the conditional median.

5. See also Whittle (1979).

REFERENCES

- Amemiya, T. (1985) *Advanced Econometrics*. Cambridge, Massachusetts: Harvard University Press.
- Christoffersen, P.F. & F.X. Diebold (1994) Optimal Prediction under Asymmetric Loss. National Bureau of Economic Research Technical Working paper 167, Cambridge, Massachusetts.
- Christoffersen, P.F. & F.X. Diebold (1996) Further results on forecasting and model selection under asymmetric loss. *Journal of Applied Econometrics* 11, 561–572.
- Diebold, F.X. & R.S. Mariano (1995) Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253–265.
- Granger, C.W.J. (1969) Prediction with a generalized cost of error function. *Operational Research Quarterly* 20, 199–207.
- Granger, C.W.J. & P. Newbold (1986) *Forecasting Economic Time Series*, 2nd ed. Orlando: Academic Press.
- Hansen, B.E. (1994) Autoregressive conditional density estimation. *International Economic Review* 35, 705–730.
- Hansen, L.P., T.J. Sargent, & T.D. Tallarini (1993) Pessimism, Neurosis, and Feelings about Risk in General Equilibrium. Manuscript, University of Chicago.
- Phillips, P.C.B. (1996) Econometric model determination. *Econometrica* 64, 763–812.
- Stockman, A.C. (1987) Economic theory and exchange rate forecasts. *International Journal of Forecasting* 3, 3–15.
- Varian, H. (1974) A Bayesian approach to real estate assessment. In S.E. Feinberg & A. Zellner (eds.), *Studies in Bayesian Econometrics and Statistics in Honor of L.J. Savage*, pp. 195–208. Amsterdam: North-Holland.

Weiss, A.A. (1996) Estimating time series models using the relevant cost function. *Journal of Applied Econometrics* 11, 539–560.
 Whittle, P. (1979) Why predict? Prediction as an adjunct to action. In O.D. Anderson (ed.), *Forecasting*, pp. 177–183. Amsterdam: North-Holland.
 Whittle, P. (1990) *Risk-Sensitive Optimal Control*. New York: Wiley.
 Zellner, A. (1986) Bayesian estimation and prediction using asymmetric loss functions. *Journal of the American Statistical Association* 81, 446–451.

APPENDIX

Proof of Proposition 1. We seek the predictor that solves

$$\min_{\hat{y}_{t+h}} E_t\{L(y_{t+h} - \hat{y}_{t+h})\} = \min_{\hat{y}_{t+h}} \int_{-\infty}^{\infty} L(y_{t+h} - \hat{y}_{t+h})f(y_{t+h}|\Omega_t) dy_{t+h}.$$

(Here and throughout, $E_t(x)$ denotes $E(x|\Omega_t)$.) Without loss of generality, we can write $\hat{y}_{t+h} = \mu_{t+h|t} + \alpha_{t+h|t}$ and $y_{t+h} = \mu_{t+h|t} + x_{t+h}$, so that

$$\operatorname{argmin}_{\hat{y}_{t+h}} E_t\{L(y_{t+h} - \hat{y}_{t+h})\} = \mu_{t+h|t} + \operatorname{argmin}_{\alpha_{t+h|t}} \int_{-\infty}^{\infty} L(x_{t+h} - \alpha_{t+h|t})f(x_{t+h}|\Omega_t) dx_{t+h}.$$

Because $f(x_{t+h}|\Omega_t)$ depends on $\sigma_{t+h|t}^2$ but not on $\mu_{t+h|t}$, so too does the $\alpha_{t+h|t}$ that solves the minimization problem depend on $\sigma_{t+h|t}^2$ but not on $\mu_{t+h|t}$. ■

Proof of Proposition 2. Precisely parallels that of Proposition 1. ■

Proof of Proposition 3. Following Amemiya (1985), we require three conditions:

1. $\beta_0 \in B$ is a compact subset of \mathbb{R}^k .
2. $L_N(\beta) = \sum_{t=1}^N L(y_{t+h}, y_{t+h}^*(\beta))$ is continuous in $\beta \in B$ for all $y = (y_{1+h}, \dots, y_{N+h})$ and is a measurable function of y for all $\beta \in B$.
3. $N^{-1}L_N(\beta)$ converges to a nonstochastic continuous function $L(\beta)$ in probability uniformly in $\beta \in B$ as $N \rightarrow \infty$, and $L(\beta)$ attains a unique global minimum at β_0 .

Under the conditions, $\hat{\beta}_N = \operatorname{argmin}_{\beta \in B} L_N(\beta)$ converges in probability to β_0 by the argument of Amemiya (1985, p. 107). Thus, $y_{t+h}^*(\hat{\beta}_N)$ converges in probability to $y_{t+h}^*(\beta_0)$ by continuity of $y_{t+h}^*(\beta_N)$. ■