

## RESEARCH ARTICLE OPEN ACCESS

# The Economic Value of Forecasts in Reducing Extreme Total Losses

David B. Stephenson 

Department of Mathematics and Statistics, University of Exeter, Exeter, UK

**Correspondence:** David B. Stephenson ([d.b.stephenson@exeter.ac.uk](mailto:d.b.stephenson@exeter.ac.uk))**Received:** 15 May 2025 | **Revised:** 15 September 2025 | **Accepted:** 14 October 2025**Funding:** The author received no specific funding for this work.**Keywords:** cost-loss ratio | expected loss | forecast value | loss variance | value-at-risk | VaR | verification

## ABSTRACT

A major aim of weather and other types of environmental forecasting is to provide early warning of extreme hazards that can then be used to take preventative actions to reduce loss. This study investigates what determines the loss distribution in the simplest context of repeatedly predicting/diagnosing the occurrence or not of a severe event/condition. Mathematical expressions for the expected total loss and variance of the total loss are derived in terms of the probability of event occurrence (the base rate), the cost-loss ratio and the hit rate ( $H$ ) and false alarm rate ( $F$ ) of the forecasting system. Expected loss and variance behave very differently as functions of hit and false alarm rate: expected loss is a linear function of  $F$  and  $H$  with a minimum at  $(F, H) = (0, 1)$  whereas variance is a non-linear function with a minimum at  $(F, H) = (1, 1)$ . For vanishingly rare events, expected loss can be less than that of taking no action only if  $F, H \rightarrow 0$  and the forecasting system issues warnings at a rate that is much less than the rate at which events occur. Rather than minimise expected losses, one might expect the value of a forecasting system to be its ability to reduce the risk of large losses. Large losses, quantified using the Value-at-Risk (VaR) measure, depend upon both the expected loss and the variance of the loss but become more determined by the variance as the base rate decreases. Minimum VaR occurs at a larger false alarm rate than does minimum expected value and so can be achieved with a forecasting system that issues warnings at a higher rate that is more comparable to the rate at which observed events occur.

## 1 | Introduction

A main aim of warning systems is to help diagnose when to take early action that can then mitigate larger potential losses from events/conditions. However, such warning systems are never perfect—they create false alarms (false positives) and misses (false negatives) that incur unnecessary prevention costs and losses from not being prepared, respectively. So an important question to ask is how valuable are such systems in reducing losses to less than those expected by never taking preventative action or taking preventative action all the time.

The simplest context in which to explore these ideas is to consider warnings of an event/condition that can be present or not (e.g., a storm on a given day). Expenses accrue either through the cost of making an intervention when there is a warning, or by suffering a loss if an event occurs when no warning has been issued. Since the early studies of Murphy (1966) and Murphy (1977), this idealised cost-loss ratio decision-making scenario has continued to serve as a helpful approach for considering the value of weather forecasts (e.g., Richardson 2000, 2001; Wilks 2001; Thornes and Stephenson 2001; Rodwell et al. 2020; Ben Bouallègue and Richardson 2022).

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *Meteorological Applications* published by John Wiley & Sons Ltd on behalf of Royal Meteorological Society.

Richardson (2012) presents a clear and comprehensive review of the approach.

These studies have focused on the expected loss of using a warning. However, forecast users generally accrue losses over a set of warnings and so the distribution of accumulated losses needs to be considered. Furthermore, there is generally greater utility in reducing large accumulated losses rather than small ones and so different risk measures than expected loss merit investigation. This study explores these ideas mathematically paying particular attention to what happens in the limit as the forecast event becomes increasingly rare, as is the case for more extreme events.

## 2 | Binary Warnings

### 2.1 | Repeated Warnings of Independent Events

Consider a warning system that issues binary predictions for a sample of  $n$  binary events (e.g., a month of daily storm forecasts, diagnosis of a medical condition in a sample of patients). For simplicity, assume that the probability of occurrence (the base rate)  $s$ , is constant and the events in the sample are independent of one another. The forecast user is assumed to take preventative action that costs  $C$  if the event is forecast to occur. If no preventative action has been taken but the event then occurs, the user makes a loss of  $L$ . For each of these events, there are four possible outcomes depending upon whether or not an event is forecast and whether or not it is subsequently observed (Table 1):

### 2.2 | Forecasting Systems

The type of forecasting system determines  $(F, H)$ . A perfect forecasting system has  $(F, H) = (0, 1)$  with  $p_H = s$ ,  $p_F = p_M = 0$ , and  $p_{CR} = 1 - s$ . A no-skill system that randomly issues warnings with probability  $f$  has  $(F, H) = (f, f)$  and  $p_H = sf$ ,  $p_F = (1 - s)f$ ,  $p_M = s(1 - f)$  and  $p_{CR} = (1 - s)(1 - f)$ . If warnings are always issued ( $f = 1$ ) then the loss for each outcome is  $C$ . If no warnings are ever issued ( $f = 0$ ) then the loss for an outcome is either  $L$  or  $0$  depending upon whether or not an event subsequently occurs. By varying a decision threshold parameter, forecasting systems can have a continuous range of  $(F, H)$  values going from  $(0, 0)$  to  $(1, 1)$  that defines a function  $H(F)$  known as the Receiver Operating Characteristic (ROC).

**TABLE 1** | The four possible outcomes that can occur.

Outcome	Warning	Event		Probability
		observed	Loss	
Hit	Yes	Yes	C	$p_H = sH$
False alarm	Yes	No	C	$p_F = (1 - s)F$
Miss	No	Yes	L	$p_M = s(1 - H)$
Correct rejection	No	No	0	$p_{CR} = (1 - s)(1 - F)$

Note: The base rate  $s$  is the probability of the event occurring, the hit rate  $H$  is the conditional probability that an observed event was forecast, and the false alarm rate  $F$  is the conditional probability that an observed non-event was forecast.

### 2.3 | Economic Loss

Three different losses are possible after each outcome: a loss  $L$  after a miss, a loss  $C$  after a hit or a false alarm, and a loss of  $0$  after a correct rejection (no event occurred and no warning was issued). We shall assume that the cost  $C > 0$  of preventative action is less than the loss  $L > 0$  incurred of not taking preventative action, and hence the cost-loss ratio  $r = C / L$  satisfies  $0 < r < 1$ . For more extreme events above higher event thresholds, vulnerability increases and so one might expect losses to increase with threshold, however, exposure (area of impact) also generally decreases and so  $L$  and  $r$  can either increase or decrease as threshold increases.

The total loss after  $n$  independent events is given by  $S = CN_H + CN_F + LN_M$  where  $N_H$ ,  $N_F$  and  $N_M$  are the number of hits, false alarms and misses respectively.  $S$  is a random variable that can take discrete irregularly-spaced values ranging from a minimum value of  $0$  when all outcomes are correct rejections up to a maximum of  $nL$  when all outcomes are misses.

The distribution of  $S$  depends upon the distributions of the counts  $N_H$ ,  $N_F$  and  $N_M$  and the cost  $C$  and loss  $L$ . The counts are multinomially distributed with probability  $P(N_H = n_H, N_F = n_F, N_M = n_M, N_{CR} = n_{CR})$  equal to

$$\frac{n!}{n_H! n_F! n_M! n_{CR}!} p_H^{n_H} p_F^{n_F} p_M^{n_M} p_{CR}^{n_{CR}}$$

Which is determined by the base rate  $s$ , and the hit and false alarm rates of the forecasting system. In what follows,  $S_{(F,H)}$  will be used to denote the total loss for a system having false alarm and hit rates  $(F, H)$ .

### 2.4 | Data Example

The concepts will be illustrated by considering an example typical of an extreme hazard such as freezing days. We shall assume that the hazard event is rare with a base rate of  $s = 1/30$ , and forecasts are made for  $n = 30$  events (e.g., a month of daily occurrences). The hazard is assumed to cause a loss of  $L = 1$  when unprotected but costs  $C = 1/10$  to protect against, which gives a cost-loss ratio  $r = 1/10$  typical for users of meteorological warnings (Thornes and Stephenson 2001). As well as considering the behaviour for all  $(F, H)$  values, we shall also investigate in more detail what happens for a forecasting system that has a ROC curve described by  $H = \theta F / (1 + (\theta - 1)F)$  having constant odds ratio  $\theta = (H / (1 - H)) / (F / (1 - F)) = 10$  (Stephenson 2000). The conclusions are generic and not highly sensitive to the choice of example.

## 3 | The Loss Distribution

### 3.1 | Expectation of the Total Loss

The expectation of the loss over  $n$  cases is given by

$$\begin{aligned} E(S) &= E(CN_H + CN_F + LN_M) \\ &= nC(p_H + p_F) + nLp_M \\ &= nsL - ns(L - C)H + n(1 - s)CF \end{aligned} \tag{1}$$

The expectation of the loss is a linear function of  $H$  and  $F$  that decreases with  $H$  and increases with  $F$ . Hence, the lowest expected loss of  $nsC$  is obtained for perfect forecasts having  $(F, H) = (0, 1)$ . Figure 1a shows an example of how  $E(S_{(F,H)})$  varies with  $(F, H)$ .

For no-skill random forecasts  $E(S_{(f,f)}) = nLs + n(C - Ls)f = nL(s - sf + rf)$ . When  $s < r$ , lowest expectation of  $E(S_{(0,0)}) = nsL$  is obtained by taking no action ( $f = 0$ ), and when  $s > r$  lowest expectation  $E(S_{(1,1)}) = nC = nrL$  is obtained by always taking action ( $f = 1$ ). Following Richardson (2012), conditions can be derived for when expected loss is less than that of either taking no preventative action or taking action every time. Expected loss  $E(S)$  is less than the expected loss,  $nsL$ , of never taking preventative action if

$$\frac{H}{F} > \phi = \left(\frac{r}{1-r}\right)\left(\frac{1-s}{s}\right). \quad (2)$$

Hence,  $H/F > 1$  is a necessary (but not sufficient) condition of having smaller expected loss than never taking action when  $s < r$  and  $\phi > 1$ . Expected loss  $E(S)$  is less than the expected loss,  $nC$ , of taking preventative action every time if

$$\frac{1-H}{1-F} < \phi \quad (3)$$

Hence,  $(1-H)/(1-F) < 1$  is a necessary (but not sufficient) condition of having smaller expected loss than taking preventative action every time when  $s < r$ . Therefore, a necessary condition for the expectation  $E(S)$  to be less than both  $E(S_{(0,0)})$  and  $E(S_{(1,1)})$  when  $s < r$ , is that  $H/F > 1$ , or equivalently, the odds ratio  $\theta = (H/(1-H))/(F/(1-F)) > 1$ .

The ROC curve,  $H(F)$ , can be used to find which  $(F, H)$  values give the lowest expected loss. The hit rate increases when moving along the ROC curve away from the no action point  $(F, H) = (0, 0)$ , which causes the expected loss to decrease but then at some point this is compensated by increasing loss due to rising false alarm

rate. The minimum loss occurs when the derivative with respect to false alarm rate equals zero:

$$\begin{aligned} \frac{dE(S)}{dF} &= -ns(L-C)\frac{dH}{dF} + n(1-s)C = 0 \\ &\Rightarrow \frac{dH}{dF} = \phi. \end{aligned} \quad (4)$$

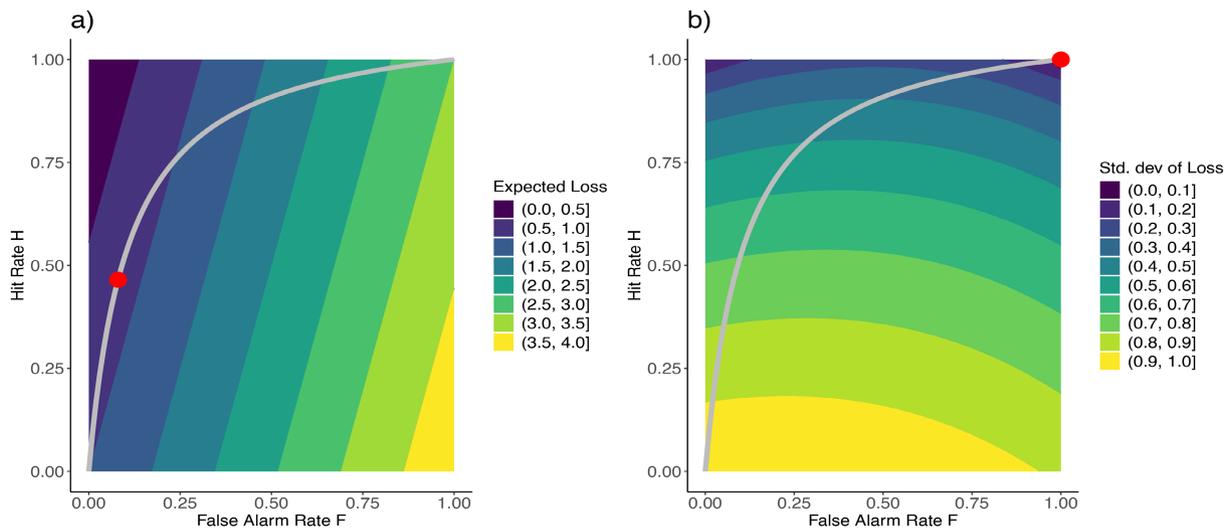
This derivative can then be used to find the optimal choice of  $F$  when  $H(F)$  is known. For example, when a forecasting system has a ROC curve that is well-approximated by  $H = \theta F / (1 + (\theta - 1)F)$  with constant odds ratio  $\theta = (H/(1-H))/(F/(1-F)) > 1$  (Stephenson 2000; Equation (13)) then  $\frac{dH}{dF} = \phi$  at  $F = (\theta - 1)^{-1}(\sqrt{\theta/\phi} - 1)$  which requires  $\theta > \phi > \theta^{-1}$  in order for  $0 < F < 1$ . The minimum expected loss is  $nsL - nsL(1-r)(\theta/(1+(\theta-1)F) - \phi)F$ . An example of this for a  $\theta = 10$  ROC curve is shown in Figure 1a. For rarer events as  $s \rightarrow 0$ ,  $\phi$  increases which necessitates a smaller optimal  $F$  because  $\frac{dH}{dF}$  decreases with  $F$ .

### 3.2 | Variance of the Total Loss

Total loss can be highly variable especially for rare events where there are very infrequent large loss  $L$  outcomes. The variance of the loss is given by

$$\begin{aligned} \text{Var}(S) &= \text{Var}(CN_H + CN_F + LN_M) \\ &= \text{Var}(CN_H) + \text{Var}(CN_F) + \text{Var}(LN_M) \\ &\quad + 2\text{Cov}(CN_H, CN_F) + 2\text{Cov}(CN_H, LN_M) + 2\text{Cov}(CN_F, LN_M) \\ &= C^2\text{Var}(N_H) + C^2\text{Var}(N_F) + L^2\text{Var}(N_M) \\ &\quad + 2C^2\text{Cov}(N_H, N_F) + 2CL\text{Cov}(N_H, N_M) + 2CL\text{Cov}(N_F, N_M). \end{aligned} \quad (5)$$

Using the property  $\text{Cov}(N_A, N_B) = np_A\delta_{AB} - np_Ap_B$  for counts from multinomial distributions (where  $\delta_{AB} = 1$  if  $A = B$  and 0 otherwise) then gives



**FIGURE 1** | Dependence of (a) expected loss  $E(S)$  and (b) standard deviation  $\sqrt{\text{Var}(S)}$  on  $(F, H)$  for the data example described in Section 2.4. ROC curves for odds ratio equal to 10 shown by grey lines with position of minima indicated by red circles.

$$\begin{aligned} \text{Var}(S) &= C^2 n p_H (1 - p_H) + C^2 n p_F (1 - p_F) + L^2 n p_M (1 - p_M) \\ &\quad - 2C^2 n p_H p_F - 2C n L p_H p_M - 2C n L p_F p_M \\ &= n C^2 ((p_H + p_F) - (p_H + p_F)^2) - 2n C L p_M (p_H + p_F) + n L^2 p_M (1 - p_M). \end{aligned} \tag{6}$$

The variance is a quadratic function of  $p_H + p_F$  and  $p_M$  and therefore of  $F$  and  $H$ . Unlike the expectation, variance does not simply decrease with  $H$  and increase with  $F$ . Figure 1b shows an example of how the square root of the variance depends upon  $(F, H)$ , and how it has a very different functional form to that of the expected loss shown in Figure 1a. For any given  $H$ , the variance increases with  $F$  until it reaches a maximum at

$$F_m = \phi^{-1} H + (1 - s)^{-1} \left( \frac{1}{2} - \frac{s}{r} \right), \tag{7}$$

and then decreases for  $F > F_m$ . Perfect forecasts have  $\text{Var}(S_{(0,1)}) = n C^2 s (1 - s)$  that does not depend on the loss  $L$  but is still greater than the zero variance  $\text{Var}(S_{(1,1)}) = 0$  obtained by taking action all the time.

Random forecasts have  $\text{Var}(S_{(f,f)}) = n(1 - f)(f(Ls - C)^2 + L^2 s(1 - s))$  and so lowest variance of 0 is obtained by taking action all the time ( $f = 1$ ) and taking no action  $f = 0$  gives  $\text{Var}(S_{(0,0)}) = n L^2 s(1 - s)$ . Maximum variance for random forecasts occurs at  $f = 1/2 - s(1 - s)/(2(s - r)^2)$ , which is always for  $f < 1/2$  but for some choices of  $s$  and  $r$  will be for  $f < 0$  and so then maximum variance will be achieved at  $f = 0$  (no action). In contrast to the behaviour for expected loss for a forecasting system with ROC curve  $H(F)$ , lowest variance of zero occurs at  $F = 1$ .

### 3.3 | Behaviour for Vanishingly Rare Events

Rarer more extreme events can lead to catastrophic losses and so it is of interest to consider what happens in the limit as the hazard event threshold increases and the base rate vanishes, that is  $s \rightarrow 0$ . Since loss  $L$  can depend on the magnitude of the hazard, the cost-loss ratio  $r$  can also change as  $s \rightarrow 0$  but eventually  $s < r$  because loss  $L$  has a finite upper bound and so  $r$  cannot vanish. When  $s < r$ , never taking action has a lower expected loss of  $nsL$  than the expected loss  $nC$  from taking action every time.

For a forecasting system to give a smaller expected loss than that of never taking action requires Equation (2) to be satisfied, which can be used to find bounds on the probability  $p = p_H + p_F = sH + (1 - s)F$  of issuing warnings:

$$\frac{sH}{r} > p > F \left( \frac{1 - s}{1 - r} \right). \tag{8}$$

In the limit as  $s \rightarrow 0$ ,  $sH/r \rightarrow 0$  and so  $p \rightarrow 0$  and  $F \rightarrow 0$ : the probability of issuing warnings and the false alarm rate vanish. Furthermore, the frequency bias  $\beta = (p_H + p_F)/(p_H + p_M) = p/s$  must be less than  $H/r$ . Typically, forecast and observed variables are asymptotically independent for extremes and so  $H \rightarrow 0$  as  $s \rightarrow 0$  for rarer events (Stephenson et al. 2008; Ferro and Stephenson 2012). This implies that the frequency bias  $\beta < H/r \rightarrow 0$  and warnings need

to be issued at a much lower rate than the base rate. The probability of misses  $p_M = s(1 - H) \rightarrow s$  so decreases slower than does  $p$ .

The expected loss  $nL(s + r(1 - s)F - s(1 - r)H)$  has partial derivatives of  $nLr(1 - s)$  and  $-nLs(1 - r)$  with respect to  $F$  and  $H$ , respectively, and so becomes increasingly less sensitive to hit rate  $H$  than to  $F$  for smaller  $s$ . When  $p, p_M \rightarrow 0$ , the variance tends to  $nL^2(r^2 p + p_M) = nL^2(s + r^2(1 - s)F - s(1 - r^2)H)$ , which has partial derivatives of  $nL^2 r^2(1 - s)$  and  $-nL^2 s(1 - r^2)$  with respect to  $F$  and  $H$ , respectively, and so also becomes increasingly less sensitive to hit rate  $H$  for smaller  $s$ . From Equation (2),  $r^2(1 - s)F - s(1 - r^2)H < 0$  for  $F, H > 0$  and so the variance of the losses from the forecasting system are less than the variance of  $nL^2$  obtained when issuing no warnings. In the limit as  $s \rightarrow 0$ , the ratio of the variance to the expected loss tends to  $L$  and the ratio of the standard deviation of the loss to the expected loss tends to  $1/\sqrt{ns}$  and so increases.

### 3.4 | Skewness and Higher Moments of the Loss Distribution

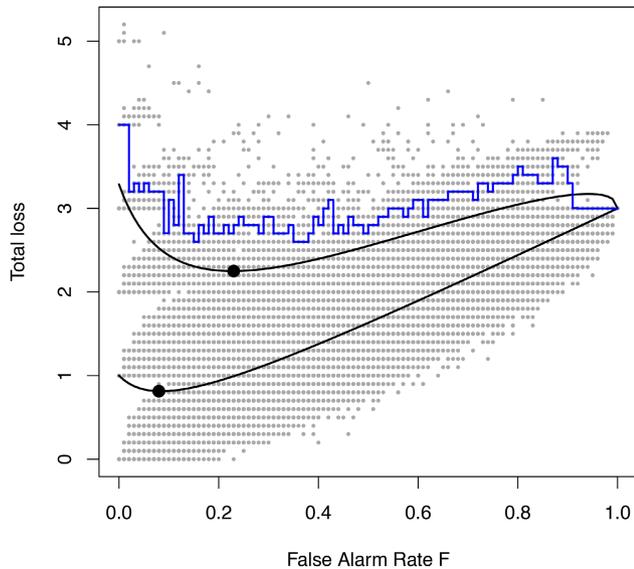
The loss distribution is not defined by only the expectation and variance, for example, it exhibits strong positive skewness particularly when  $ns$  is small. Such behaviour can be modelled using method of moments to fit shifted Gamma distributions to the losses (Section 6.2 of McNeil et al. 2015). The following will briefly show how such higher moments of the loss distribution can be derived. The expectation, variance and higher central moments can also be obtained by differentiation of the cumulant generating function  $C(t) = \log M(t)$ , where  $M(t) = E(e^{tS})$  is the moment generating function. The moment generating function is simply the  $n$ 'th power of the moment generating function for the losses of each of the  $n$  independent outcomes:

$$\begin{aligned} C(t) &= \log E(e^{tS}) \\ &= n \log(p_{CR} + (p_H + p_F)e^{tC} + p_M e^{tL}) \\ &= n \left( u - \frac{u^2}{2} + \frac{u^3}{3} - \dots \right) \end{aligned} \tag{9}$$

where  $u = (p_H + p_F)(e^{tC} - 1) + p_M(e^{tL} - 1)$ . Central moments can be found by taking derivatives of  $C(t)$  with respect to  $t$ , so for example, the variance equals  $C'' = n(u'' - (u')^2)$  and the 3rd central moment  $E((S - E(S))^3) = C''' = n(u''' - 3u''u' + 2(u')^3)$  evaluated at  $t = 0$ . In the limit as  $s \rightarrow 0$ ,  $u \rightarrow s(e^{tL} - 1)$  and so, for example,  $E((S - E(S))^3) \rightarrow nsL^3$ .

### 3.5 | Extreme Losses

Rather than using a forecasting system to minimise expected losses, one may be interested in using it to reduce extreme losses. One way to quantify this is to consider the popular risk measure known as Value-at-Risk, which is simply the  $1 - \alpha$  quantile of the loss distribution  $s_\alpha = F_S^{-1}(1 - \alpha)$  where  $F_S$  is the cumulative distribution function of  $S$  (Section 2.2, McNeil et al. 2015). The probability of loss  $S$  breaching  $s_{(1-\alpha)}$  is  $\alpha$ .



**FIGURE 2** | Simulated losses (grey dots) and 0.99 empirical quantile (blue line) for values of false alarm rate and hit rate along the ROC curve of the data example. Smooth curves show the expected loss (lower curve) and the Gaussian estimate of the 0.99 quantile (upper curve) with black circles marking their minima.

An example of Value-at-Risk is shown in Figure 2 (blue line) where 1000 losses have been simulated for each of 100 values of  $(F, H)$  from along the ROC curve of the data example. For large  $n$ , the loss distribution becomes well-approximated by the Normal distribution (Central Limit Theorem), which allows Value-at-Risk to be approximated by  $s_\alpha \approx E(S) + \sqrt{\text{Var}(S)}\Phi^{-1}(1 - \alpha)$ , where  $\Phi^{-1}(1 - \alpha)$  is the  $1 - \alpha$  quantile of the standard normal distribution. This gives the upper smooth black curve shown in Figure 2, which closely follows the empirical quantile (blue line) albeit slightly below due to the positive skewness of the loss distribution. The minimum value for this quantile occurs at larger false alarm rate ( $F = 0.23$ ) than where the minimum occurs in expected value ( $F = 0.08$ ). This shift becomes more pronounced for more extreme quantiles and for smaller base rates where the quantile is increasingly controlled by the variance of the losses rather than by the expected loss.

## 4 | Conclusions

This study has investigated properties of the loss distribution in the simplest context of deterministically forecasting the occurrence or not of a rare severe event. Mathematical expressions for the expected loss and variance of the loss have been derived in terms of the base rate and the cost-loss ratio, as well as the hit and false alarm rates of the forecasting system. Expected loss is a linear function of false alarm with a minimum at zero false alarm rate, whereas the variance of the loss is a non-linear function that can decrease with false alarm rate. The limits of these measures for vanishingly rare events have been presented and it has been shown that for forecasts to have value (i.e., expected loss less than that of never making a warning), a forecasting system must issue warnings at a rate that is much less than the rate at which events occur. For increasingly rare events, the expected loss becomes increasingly determined by the false alarm rate rather than the hit rate. Large losses, quantified using the

Value-at-Risk (VaR) measure, depend upon both the expected loss and the variance of the loss but become more determined by the variance for rarer events. Minimum VaR occurs at a larger false alarm rate than does minimum expected value and so can be achieved with a forecasting system that issues warnings at a higher rate that is more comparable to the rate at which observed events occur.

Following common practice in forecast verification, several simplifying assumptions have been made about the properties of the events. It has been assumed that the base rate is constant, which is likely to be a good approximation for sets of forecasts issued only over relatively short forecasting periods (e.g., seasons or less). It has also been assumed that successive events (and successive forecasts) are independent of one another. This assumption is more questionable for real-world hazards that do have serial dependence for example daily extremes in heat waves. However, including serial dependence will generally inflate the variance in the total loss making it even more important in determining extreme losses.

The findings here are not only relevant to forecasting environmental hazards but also to the diagnosis of other types of binary events such as medical screening for rare diseases (e.g., Smith et al. 2000). The lower bound in Equation (8) can be used to show that the probability that disease is present given a positive test result (the *positive predictive value*),  $sH/p$ , has a maximum value of  $(H/F)s/(1-s)(1-r)$  which tends to  $(H/F)s(1-r)$  in the limit of vanishing incidence as  $s \rightarrow 0$ . Hence, for small  $r$ , the positive predictive value is bounded from above by the incidence rate times the odds ratio as  $F, H \rightarrow 0$ . Measures such as Value-at-Risk presented here are of potential interest in quantifying errors in cross-classification in other disciplines.

## Author Contributions

**David B. Stephenson:** conceptualization, investigation, writing – original draft, methodology, validation, visualization, writing – review and editing, software, formal analysis.

## Acknowledgements

The author has nothing to report.

## Conflicts of Interest

The author declares no conflicts of interest.

## Data Availability Statement

Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

## References

- Ben Bouallègue, Z., and D. S. Richardson. 2022. “On the ROC Area of Ensemble Forecasts for Rare Events.” *Weather and Forecasting* 37, no. 5: 787–796.
- Ferro, C. A. T. F., and D. B. Stephenson. 2012. “Deterministic Forecast Events and Warnings.” In *Forecast Verification: A Practitioner’s Guide in Atmospheric Science*, edited by I. T. Jolliffe and D. B. Stephenson, 2nd ed., 249. Wiley.

- McNeil, A. J., R. Frey, and P. Embrechts. 2015. *Quantitative Risk Management: Concepts, Techniques and Tools—Revised Edition*, 720. Princeton University Press.
- Murphy, A. H. 1966. “A Note on the Utility of Probabilistic Predictions and the Probability Score in the Cost-Loss Ratio Decision Situation.” *Journal of Applied Meteorology* 5: 534–537.
- Murphy, A. H. 1977. “The Value of Climatological, Categorical and Probabilistic Forecasts in the Cost-Loss Ratio Situation.” *Monthly Weather Review* 105: 803–816.
- Richardson, D. S. 2000. “Skill and Relative Economic Value of the ECMWF Ensemble Prediction System.” *Quarterly Journal of the Royal Meteorological Society* 126: 649–668.
- Richardson, D. S. 2001. “Measures of Skill and Value of Ensemble Prediction Systems, Their Interrelationship and the Effect of Ensemble Size.” *Quarterly Journal of the Royal Meteorological Society* 127: 2473–2489.
- Richardson, D. S. 2012. “Economic Value and Skill.” In *Forecast Verification: A Practitioner’s Guide in Atmospheric Science*, edited by I. T. Jolliffe and D. B. Stephenson, 2nd ed., 249. Wiley.
- Rodwell, M. J., J. Hammond, S. Thornton, and D. S. Richardson. 2020. “User Decisions, and How These Could Guide Developments in Probabilistic Forecasting.” *Quarterly Journal of the Royal Meteorological Society* 146: 3266–3284.
- Smith, J. E., R. L. Winkler, and D. G. Fryback. 2000. “The First Positive: Computing Positive Predictive Value at the Extremes.” *Annals of Internal Medicine* 132, no. 10: 804–809.
- Stephenson, D. B. 2000. “Use of the “Odds Ratio” for Diagnosing Forecast Skill.” *Weather and Forecasting* 15, no. 2: 221–232.
- Stephenson, D. B., B. Casati, C. A. T. Ferro, and C. A. Wilson. 2008. “The Extreme Dependency Score: A Non-Vanishing Measure for Forecasts of Rare Events.” *Meteorological Applications* 15, no. 1: 41–50.
- Thornes, J. E., and D. B. Stephenson. 2001. “How to Judge the Quality and Value of Weather Forecast Products.” *Meteorological Applications* 8: 307–314.
- Wilks, D. S. 2001. “A Skill Score Based on Economic Value for Probability Forecasts.” *Meteorological Applications* 8: 209–219.