LEARNING CURVES IN CONTROL ROOMS: SKILL? RULE? KNOWLEDGE?

Romney B. Duffey
Atomic Energy of Canada Limited
Chalk River, Ontario K0J 1J0 Canada

John W. Saull
International Federation of Airworthiness
Little Beeches, Woodshill Lane, Ashurst Wood, East Grinstead, W. Sussex RH19 3RF UK

ABSTRACT

Reactor safety and risk are dominated by the potential and major contribution for human error in the design, operation, control, management, regulation and maintenance of the plant, and hence to *all* accidents. We need to determine the outcome (error) probability, or the chance of failure. Time and again we are faced with the same situation and question: how to predict the risk or chance of a mistake by an operating crew or team.

Conventional reliability engineering is associated with the failure rate of components, or systems, or mechanisms, not of human beings in and interacting with a technological system. The probability of failure requires a prior knowledge of the total number of outcomes, which for any predictive purposes we do not know or have. Analysis of failure rates due to human error based on the Learning Hypothesis allows a new determination of the dynamic human error rate in technological systems, consistent with and derived from the available world data. The basis for the analysis is that humans learn from experience, both prior to and during a transient or event, and consequently the accumulated experience defines the failure rate.

Our "best" equation for the probability of human error, outcome or failure rate, which has been validated against the full spectrum of the world's outcome data, allows for calculation and prediction of the probability of human error.

In nuclear probabilistic risk assessment, the modeling of nuclear plant operator actions and transient control behavior is extremely important, and is a requirement according to industry standards. The human error probability (HEP) often classified according to Skill, Rule and Knowledge based behavior. We examine the data and results observed in transients in both plants and simulators, available from France and the USA. We demonstrate that the human error probability (HEP) is dynamic, and that it may be predicted using the Learning Hypothesis and the minimum failure rate, and can be utilized for probabilistic risk analysis purposes.

INTRODUCTION

Using the learning hypothesis, analysis of the human rate of learning allows a new determination of the dynamic *probability* and human failure (error) rate in technological systems. The result is consistent with and derived from the available world data for modern technological systems [1,2,3,4,5,6]. Since the approach is based on the Learning Hypothesis, the resulting probability of outcome or error can be shown to agree with available data. The model is entirely consistent with the "theory of error correction" for human learning [7].

The two useful concepts are of the probability of failure (error), $p(\varepsilon)$ or *cumulative distribution function* (CDF) before or in less than any interval, ε; and the *probability density function* (PDF) as the differential probability of failure, $f(\varepsilon)$, in any small interval of experience, $d\varepsilon$. Usually or conventionally, elapsed calendar or component operating time, t, is taken as the independent variable, not the experience, ε, which is adopted here on the basis of the observed world outcome data.

THE PROBABILITY OF HUMAN RELIABILITY

The chance and risk of making or having errors changes with the accumulated experience. The reliability relationships and definitions in our new experience terminology are:

a) the hazard function is equivalent to the *failure or outcome rate* at any experience, $\lambda(\varepsilon)$, being the relative rate of change in the reliability with experience;
b) the *CDF or outcome fraction*, $F(\varepsilon)$, is just the observed frequency of prior outcomes, the ratio n/N, where we have recorded n, out of a total possible of N outcomes; and is identical to the observed *cumulative prior probability*, $p(\varepsilon)$;
c) the chance of an outcome in any small observation interval, is the PDF $f(\varepsilon)$, which is just the rate of change of the failure or outcome fraction with experience, $dp(\varepsilon)/d\varepsilon$.

Hence, the probability of the outcome or error occurring in or taking less than ε, is just the CDF, $p(\varepsilon) = n/N$, conventionally written as $F(\varepsilon)$. Relating this to the failure rate, via (a) through (c) above, gives:

$$p(\varepsilon) \equiv F(\varepsilon) = 1 - e^{-\int \lambda d\varepsilon} \qquad (1)$$

Therefore, the probability is a *double exponential* due to the exponential form of the Minimum Error Rate Equation (MERE) failure rate itself imposed on the probability expression. This form is related to or may be considered as an "extreme value distribution" function that has arisen quite naturally from the learning hypothesis. It creates a probability that is in the form of a "human bathtub" as experience increases. Substituting for the MERE hazard or failure rate, and carrying out the integration from an initial experience, ε_0, to any interval, ε, we obtain the probability as the double exponential:

$$p(\varepsilon) = 1 - \exp\{(\lambda - \lambda_m)/k - \lambda(\varepsilon_0 - \varepsilon)\} \qquad (2)$$

where, of course from the MERE,

$$\lambda(\varepsilon) = \lambda_m + (\lambda_0 - \lambda_m) \exp - k(\varepsilon - \varepsilon_0) \qquad (3)$$

and $\lambda(\varepsilon_0) = \lambda_0$ at the initial experience, ε_0, accumulated up to or at the initial outcome(s).

Using our "best" values for the learning rate constant, k=3, and for the minimum failure rate, we find that:

$$\lambda = 5.10^{-6} + \{(1/\varepsilon) - 5.10^{-6}\}e^{-3\varepsilon} \qquad (4)$$

The corresponding PDF $f(\varepsilon)$, is the probability that the error or outcome occurs in the interval dε, derived from the change in the CDF failure fraction with experience:

$$f(\varepsilon) = \lambda e^{-\int\lambda d\varepsilon} = \lambda(\varepsilon) \text{ x } (1 - p(\varepsilon)),$$

$$= \{(\lambda_m + (\lambda_0 - \lambda_m) \exp(-k(\varepsilon - \varepsilon_0))\} \text{ x } \{\exp((\lambda(\varepsilon) - \lambda_0)/k - \lambda_m(\varepsilon_0 - \varepsilon))\} \qquad (5)$$

The limits are clear: as experience becomes large, $\varepsilon \rightarrow \infty$, or the minimum rate is small, $\lambda_m << \lambda_0$, or the value of k varies, etc.

For comparison with standard "engineering reliability" models, we show the MERE result also as a Weibull chart or "probability plot" form in Figure 1.

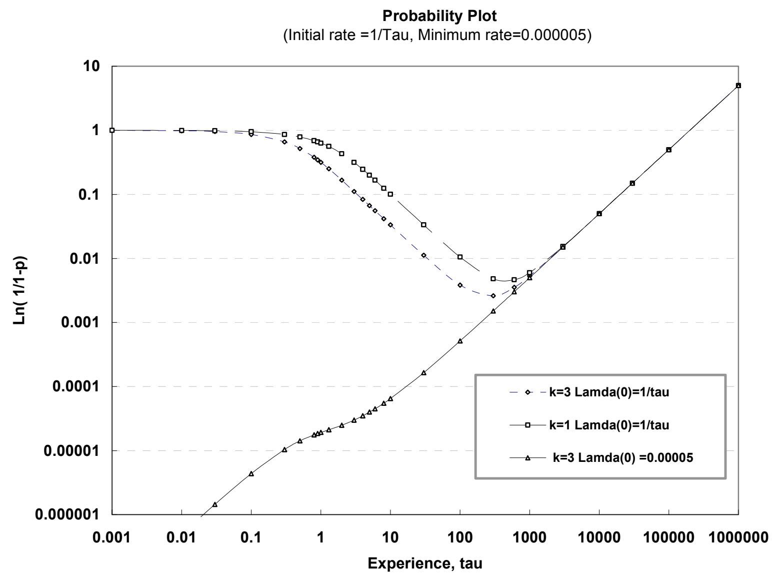

Figure 1: Human error probability shown in a Weibull form of plot.

In Figure 1 the natural logarithm $\ln(1/(1-p(\varepsilon))$ is shown plotted versus the accumulated experience, ε, measured in tau units for two learning rates, k=1 and k=3. The plot demonstrates that the influence of learning is to produce a distinct "kink" in otherwise straight (or conventional) Weibull log-log lines. For comparison only, we show an example with an assumed low constant initial rate of 0.000005, consistent with our minimum value derived from data.

THE HUMAN BATHTUB

The probability of human error, and its associated failure or error rate, we expect to be unchanged unless dramatic technology shifts occur. We can also estimate the likelihood of another event, and whether the MERE human error rate frequency gives sensible and consistent *predictions*. Using Bayesian reasoning, the posterior or future probability, p(P), of an error when we are at experience, ε, is,

$$\text{Posterior, } p(P) = \{\text{Prior, } p(\varepsilon)\} \times \{\text{Likelihood, } p(L)\} \quad (6)$$

where $p(\varepsilon)$ is the prior probability, and by definition both $|P,L| > \varepsilon$, our present accumulated experience.

The likelihood, p(L), is also a statistical estimate, and we must make an assumption, based on our prior knowledge, and often is taken as a uniform distribution. We can show that the likelihood is formally related to the number of outcomes for a given variation of the mean. Either:

a) the future likelihood is of the same form as experienced up to now; and/or
b) the future is an unknown statistical sample for the next increment of experience based on the differential probability, the PDF $f(\varepsilon)$.

In the first case (a), we have that the future likelihood probability p(L) is the fraction or ratio of events remaining to occur out of the total possible number that is left.

For the second case (b), the future is an unknown statistical sample for the next increment of experience based on the PDF, $f(\varepsilon)$. This is called a "conditional probability", where the probability of the next outcome depends on the prior ones occurring, which was Bayes original premise.

The so-called generalized *conditional* probability or Likelihood, p(L), can be defined utilizing the CDF and PDF expressions. Described by Sveshnikov [8] as the "generalized Bayes formula", the expression given is based on the prior outcome having already occurred with the prior probability $p(\varepsilon)$.

This prior probability then gives the probability or Likelihood of the next outcome, p(L), in our present experience-based notation, as:

$$p(L) = \frac{(f(\varepsilon), \text{ probability that the outcome occurs in the interval } d\varepsilon)}{(p(\varepsilon), \text{ probability that the outcome occurred in or less than } \varepsilon)}$$

$$\equiv \{PDF/CDF\} = f(\varepsilon)/p(\varepsilon)$$

$$= 1/p(\varepsilon)\{dp(\varepsilon)/d\varepsilon\} = \lambda e^{-\int\lambda d\varepsilon}/(1 - e^{-\int\lambda d\varepsilon})$$

$$= \lambda\{(p(\varepsilon)-1)/p(\varepsilon)\} \qquad (7)$$

In mathematical notation [8], the PDF is a differential function, $f(\varepsilon) = dp(\varepsilon)/d\varepsilon$, being the probability of an outcome in any experience increment; and the CDF is an integral function for the observed outcome (failure) fraction, $F(\varepsilon)$, being the probability of the outcomes for all experience. Both functions can be evaluated using the (continuous random variable) MERE exponential solution for the outcome (failure) rate as we now show.

So, in the second case, for the next increment of experience, we may take the likelihood, p(L) as related to the PDF, $f(\varepsilon)$, and the CDF, $p(\varepsilon)$ by the expressions for the posterior probability:

$$p(P) = p(\varepsilon) \text{ x } (PDF, f(\varepsilon)/CDF, p(\varepsilon)) = f(\varepsilon) \qquad (8)$$

This implies that the likelihood is as we stated in (Eq.7). We can evaluate these Bayesian likelihood and posterior expressions using our "best" MERE values, obtaining the results shown in Figure 2.

It is clear from Figure 2 that the "human bathtub" prior probability, $p(\varepsilon)$, causes the likelihood to fluctuate up and down with increasing experience. The likelihood tracks the learning curve, then transitions via a bump or secondary peak to the lowest values as we approach certainty ($p \rightarrow 1$) at large experience. However, the posterior probability, p(P), just mirrors and follows the MERE failure rate, as we predicted, decreasing to a minimum value of $\sim 5.10^{-6}$, our ubiquitous minimum outcome rate, before finally falling away.

Hence, since the future probability estimate, the posterior p(P), is once again derivable from its (unchanged) prior value, $f(\varepsilon) = dp(\varepsilon)/d\varepsilon \sim \lambda(\varepsilon)$, derived from learning from experience, and thus *the past predicts the future*.

This purely deterministic view is predicated by assuming an unchanging homo-technological system (HTS) and learning rate, thus reflecting reality, and that the prior "collective" of outcomes is a true sample of the posterior ones. Therefore, as usual, the uncertainty is determined by the prior probability.

For *rare events*, $\lambda(\varepsilon) \sim n/\varepsilon$, and $p(\varepsilon) \ll 1$, so a sensible working estimate for the PDF is $f(\varepsilon) \sim \lambda(\varepsilon) \sim n/\varepsilon$, where $n \sim 1$. We can show how this estimate for the likelihood indeed corresponds to that derived from the learning theory with negligible learning (a rate constant $k \sim 10^3$), thus showing a consistent result.

For the special case of "perfect learning" when we learn from all the non-outcomes as well as the outcomes, the Poisson-type triple exponential form applies for low probabilities and small numbers of outcomes ($n \ll m$). Of course, the limit of "perfect learning" is when we have an outcome, so here $p(\varepsilon) = 1/\varepsilon$, and is the rare event case for n = 1. The Perfect Learning limit fails as soon as we have an event, as it should. But there is also a useful simple physical interpretation, which is that:

a) we learn from non-outcomes the same way we learn from outcomes;
b) the perfect learning ends as soon as we have just a single (rare) outcome; and
c) the influence of the finite minimum rate is then lost.

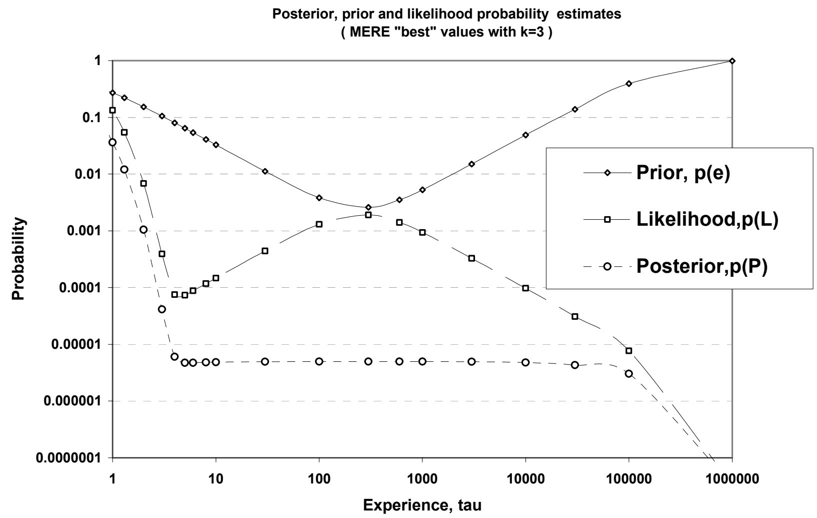

Figure 2: The estimate of the likelihood and posterior probabilities when learning.

COMPARISON TO DATA: THE PROBABILITY OF FAILURE AND HUMAN ERROR

We may now compare the above learning theory with the available data, where we are looking for the major impact of human error on the outcome rate. Fortunately or unfortunately, "human error" has an overriding (typically >60%) contribution to almost all HTS outcomes and events. This is true worldwide for the whole spectrum, all the way from transportation crashes, social system and medical errors, to large administrative failures and the whole gamut of industrial accidents. Such outcome rates generally follow the ULC (Duffey and Saull (2002)[4]), where a learning pattern is clearly evident, and well over 1000 data sources formed the basis for the estimated value of k~3 for the learning rate "constant" taken above.

To compare the failure rates to outcome data for the probability of human error now requires a further analysis step that we outline here. There are three data sets for catastrophic events with defined large human error contributions that are worth re-examining further:

1. the crash rate for global commercial airlines, noting most occur during maneuvering and approach for take-off and landing but as we have seen can also occur in flight;
2. the loss of the space shuttles, *Challenger* and *Columbia*, also on take-off and during the approach for landing; and
3. the probability of non-detection by plant operators of so-called latent (hidden) faults in the control of nuclear system transients.

Apparently disparate, these three all share the *common element of human involvement* in the management, design, safety "culture", control and operation of a large technological system; all are subject to intense media and public interest; and the costs of failure are extremely expensive and unacceptable in many ways.

For the first two cases, we calculate the outcome (fatal crash) probability for all and each i^{th} airline, where the probability, $p_i(\varepsilon)$ is given by $n_i(\varepsilon)/N$, where N is the total number of outcomes (~276 in the 30-year observation interval). To normalize the data over the observation interval of experience, we must adopt a measure for the maximum experience, which in the case of commercial airlines and *Concorde* is taken as the 720 Mh accumulated (τ_{max}) of actual flying from 1970–2000 (i.e., 1 tau ≡ 1Mh flying).

For the Shuttle [9], as a test we take for the experience normalization either:

1. the maximum Mh aircraft value, thus assuming the shuttle is just another type of commercial-type flying machine with the same human error but independent causes; or

2. just the total launch and re-entry amount for all the 113 shuttle missions, hence assuming that the failure rate and human error mechanisms are completely unrelated to that for commercial aircraft.

The comparison of the data to theory is shown in Figure 3 where the lines are the MERE calculated probability, p(ε) using the "best" values. The three lines use three bounding values for the minimum error rate to illustrate the sensitivity. Despite the scatter, a minimum rate of order ~5.10^{-6} is indeed an upper bound value, as we estimated before. The Shuttle data point sensitivity to the four orders of magnitude variation in choice of maximum experience shows the outcome probability is well matched with the aircraft data when only shuttle flights are considered, demonstrating that the minimum human error is independent but indeed is about the same magnitude (~5.10^{-6}). We note two points: the probability is now increasing with experience, as the minimum has been attained and passed; and the chance (probability) of a fatal crash for any airline is typically a maximum of between 1 and 10%. Thus, 90 to 99% of the airlines are the lucky ones and do not have one so far.

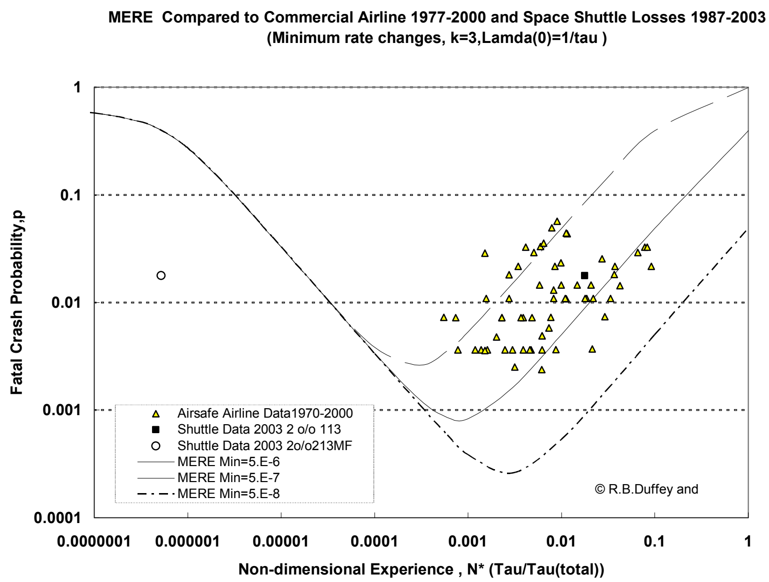

Figure 3: The probability of failure: MERE comparison to airline and shuttle data.

OPERATOR ACTIONS AND ERRORS IN NUCLEAR PLANTS

For the other human error case, we have the results of the probability of non-detection (i.e., human error) of latent faults for nuclear plant transients, which are also fairly regular events. The transient events examined by Baumont et al [10] were derived for a total reactor "fleet" of 58 units, which reported 900 outcomes spread over two years. The learning opportunity is the average experience per outcome for all the fleet, which is then given by:

Experience per outcome = (58 x 2 years x 365 days x 24 hours x 60)/900
$\equiv \tau_M$ = 68,000 minutes per event.

The data and the MERE error rate can be normalized to fitting the curves using this maximum experience (1 tau ≡ 1 transient hour) for the specific transients, close to the MERE minimum error interval. The initial probability was taken as unity, that is $p(\varepsilon_0) = 1$, for comparison purposes.

Initially, the operators had little or no chance of detection -- the latent fault remained undetected. As experience was gained and the event unfolded, the chance of finding the hidden fault increased dramatically. The data and the theory are in reasonable accord, despite the necessity of having to be renormalized. Perhaps the key two observations here are that: the shape of the data curve is indeed following the MERE failure rate prediction: and that the operators were indeed on the steepest (downward) or learning part.

This probability trend is also entirely consistent with what happened in the aircraft loss of fuel event over the Azores [11]. The latent fault was the hidden fuel-line leak, and it took about an hour, or ~ 10^{-5} tau on this same figure scale, for the crew to make the wrong diagnosis. Assuming the same human error forces are at work between the two industries, the probability of non-detection of this latent fault we would estimate at ~90%, consistent with the actual outcome.

All the outcomes we observe are the outcomes we should have expected, given the human involvement and the probability of human error.

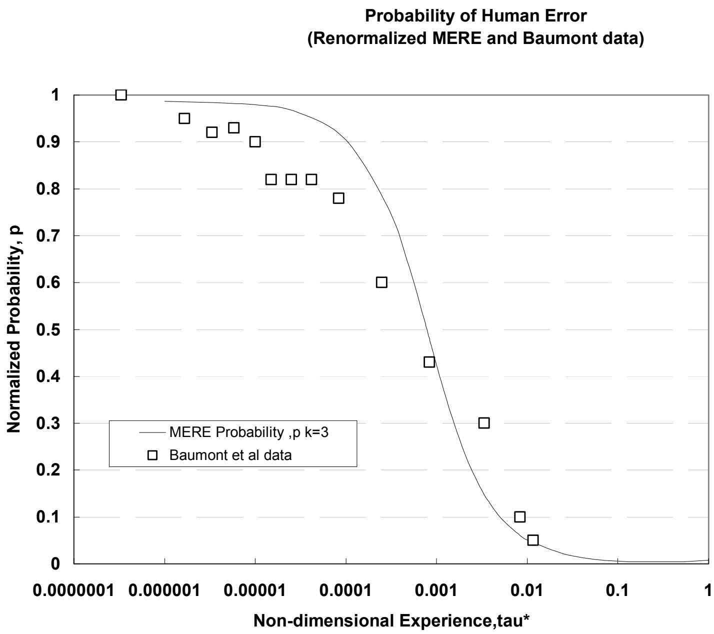

Figure 4: Comparison of MERE to human error probability data of Baumont et al. for French nuclear plants.

Thus, the human error probability is indeed *dynamic*, and evolves with experience. In this difficult arena of coupled human and technological system behavior, these new results show a very reasonable level of concordance between the mere theory and the human error data, using the typical minimum error interval (100,000 to 200,000 hours). The method and approach we have validated here also allows for predictions to be made.

THE IMPACT OF LEARNING ON HUMAN RELIABILITY ANALYSIS IN NUCLEAR SAFETY

In probabilistic safety assessments, the use of human reliability analysis (HRA) is used to assess the probability of successful or conversely unsuccessful, human actions during transient and stressful decision-making. For nuclear power plants, which are another well-known HTS where operator and human actions are required, there is even an ASME engineering standard for risk assessment [12]. This includes the explicit treatment of human error in so-called "dependency analysis" during postulated reactor accident sequences.

The probability of human non-response (error) in a transient is handled in probabilistic safety analysis (PSA) in a number of ways, including static and dynamic terms and multipliers to include the effects and influence of cognition (understanding), implementation (action) and decision-making response (timing), as well as the human's state or basis for action (skill, rule or knowledge). Hannaman [13] has summarized the various empirical forms of the HEP functions. In general, the form taken is a summation of different error components:

$$p(t) = \Sigma_i P_i \qquad (9)$$

where P_1 is the error in detection and diagnosis, $P_2(t)$ is due to planning or non-response error, $P_3(t)$ the action error, etc., and these are also distinguished for skill, S, rule, R, and knowledge, K, based behavior. The effective timescale, t, is based on some median decision or diagnosis time, and it is stated and clear that there are not much statistical data to support the different model elements. So appeal must be made to simulator observations, where operating staff is put through simulated transients and their actions observed.

All the data and fits for the HEP values have been adjusted or normalized to some choice for the median decision or action timescale ($T_{½}$) based on observation of simulated events. Data on operator errors have been collected by Hannaman and others [13] for a range of some 200-simulator tests, and fitted with exponential and lognormal functions. The data and models reported show that the HEP estimates for the three classes of action are such that at any given instant the hierarchy is p(Skill)<p(Rule)<p(Knowledge, particularly for P_2, the non-response probability. So we asked Hannaman for access to the original data behind these unpublished functional fit curves. However, although unable to release the actual data, Hannaman [14] kindly supplied us with an Excel (xls.) data file containing the separate tabulations for the Skill, Rule and Knowledge probability functional fits.

To obtain the actual timeframe or risk opportunity *during the transient*, we must choose the appropriate needed units, τ, for the experience interval to calculate the MERE failure rate and probability. We adopt a simple multiplier, ξ, making the units for the experience interval, $\xi\tau$, which then cover the possible range of timescales for Skill, Rule and Knowledge based actions. The necessary assumption is that the behavior can indeed be grouped into these three classifications.

For comparison with the MERE result for the probability of non-response error (outcome) rates for any HTS in effect we have shifted the usual bathtub experience, ε, using the simple adjusted accumulated experience parameter, $\xi\varepsilon$:

$$p(\varepsilon) = 1 - \exp\{(\lambda - \lambda_m)/k - \lambda(\varepsilon_0 - \xi\varepsilon)\} \qquad (10)$$

We adopted without any adjustment the previously determined "best" values of k=3 and λ_m=5.10^{-6}, as derived from the comparisons to the world outcome data.

We initially found that the log-normal functions originally used and supplied by Hannaman [13] did not have a finite minimum probability. Hannaman [14] then included an arbitrary constant lowest value for each of the S,R and K cases; which contrasts with the MERE minimum determined naturally by the intersection of the falling learning and rising finite minimum failure rate portions to form the bathtub. The final comparisons obtained from these choices of experience measure and minima are shown in Figure 5, which contains the MERE probability prediction and the modified lognormal functions, p(S,K,R) [14]. The four MERE "human bathtub" curves cover the range of trends exhibited by the HEP functions, using arbitrarily chosen constant experience multiplier values of ξ= 0.1,1,10, and 100 on the experience units, τ.

Rather amazingly, the Skill, Rule and Knowledge curves generally agree with the bounds of MERE results using just these simple experience timescale adjustments. Although it is not perfect, the agreement between the trends and magnitudes are now very encouraging given:(a) the approximations necessarily adopted in order to make the comparisons possible in the absence of the original data; and (b) the fact that the original curves are based on arbitrary (and not theoretically-based) lognormal fits anyway.

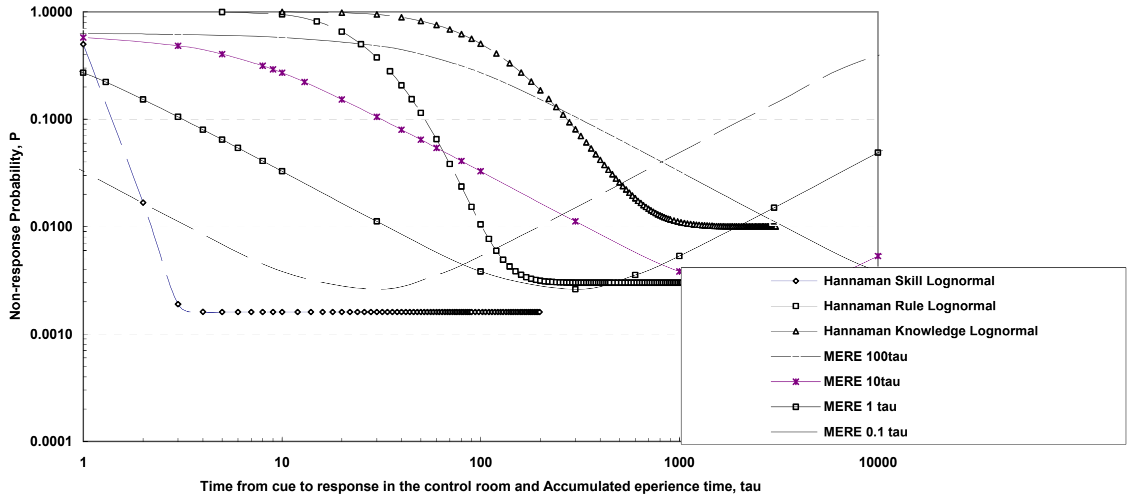

Figure 5: The Human Error Probability Comparison.

The agreement between these older data with the new MERE result is extremely encouraging given the simple shift choice made for experience, and constitutes an independent validation of the Learning Hypothesis. Hence, it clearly shows that we may use the Learning Hypothesis [4] to determine the dynamic HEP in actual transients and events, as well as the outcomes in entire HTS. This result is also consistent with the latest models developed for explaining the human learning process [7].

Therefore, for estimating the HEP in PSA analysis we recommend adopting the theoretically derived MERE "bathtub" expression,

$$p(\tau) = 1 - \exp\{(\lambda - \lambda_m)/k - \lambda(\tau_0 - \xi\tau)\}, \quad (11)$$

where

$$\lambda=\lambda_m + (\lambda_0 - \lambda_m)\exp - k(\xi\tau-\tau_0), \text{ and } 0.1<\xi<100. \quad (12)$$

Apparently, by shifting the experience parameter, τ, we can estimate p(Skill) using $\xi=0.1$; for p(Rule), $\xi=10$; and for p(Knowledge), $\xi=100$, thus underlining and emphasizing the performance premium obtained from having and acquiring skill. At any given experience, the transition to a Skill based response drastically reduces the probability of error.

IMPLICATIONS FOR GENERALIZED RISK PREDICTION

The implications of using this new approach for estimating HEP in HRA are potentially very profound. From this comparison and analysis, we may conclude that, within the uncertainties of such an analysis, the required standard HRA HEP models used in PSA can be fitted to the MERE form as derived from the Learning Hypothesis. Conversely, the MERE probability (the human bathtub) properly represents all the known data trends, such as they are, and hence can be used in PRA HEP estimation provided the correct measure is taken for experience.

The important point is that the same learning trend is evident, from actual detailed human (operator or team) actions observed during a specific transient, all the way to the outcomes for accidents and events observed during entire HTS operation. The only difference lies in the experience measure chosen as relevant to the Learning Hypothesis, which choice then naturally reflects the experience accumulated with whatever level of HTS and outcomes are under consideration.

This new probability estimate is based on the failure rate describing the Universal Learning Curve (ULC), utilizing the validation from the outcome data of the world's HTS. Thus we have *seamlessly* linked all the way from individual human actions to the observed outcomes in entire systems throughout experience space, utilizing the same values everywhere for the learning rate constant, k, and the minimum error rate, λ_m. The implication is also that separate accounting or addition is then *not* needed for the assumed cognition, P_1, and execution probabilities, P_3. They are already implicitly included in the integral MERE probability estimate, $p(\tau)$, which also automatically exhibits the minimum possible attainable probability as a natural function of the accumulated experience, τ.

Hence, for the first time, we are also able to make predictions of the probability of errors and outcomes for *any* assumed experience interval in *any* HTS.

CONCLUSIONS: THE PROBABLE HUMAN RISK

Analysis of human errors and the rate of learning allow a new determination of the dynamic human error rate in nuclear and other technological systems, consistent with and derived from the available world data. The basis for the analysis is the "learning hypothesis" that humans learn from both prior and current experience, and consequently the accumulated experience defines the failure rate. The exponential failure rate solution of the MERE then defines the probability of human error as a function of experience, which forms the shape of the "human bathtub" curve. This estimate allows for calculation and prediction of the probability of human error in any system.

Comparisons to observed human error data, all the world's data, to individual nuclear plant operator transient control behavior, show accord with the Learning Hypothesis and the "human bathtub" result. The results demonstrate that the HEP is dynamic, and that it may be predicted using the MERE learning hypothesis and the minimum failure rate, and can be utilized for probabilistic risk analysis purposes with the appropriate choice of the experience measure.

ACKNOWLEDGEMENT

The authors thank G. William Hannaman for his valuable comments and critique of the ideas presented here, and for making available the tabulated curve fits.

REFERENCES

[1] Duffey, R.B. and Saull, J.W., 2005, "The Probability of System Failure and Human Error and the Influence of Depth of Experience", Proceedings of International Topical Meeting on Probabilistic Safety Analysis (PSA'05), San Francisco, CA, September 11-15.

[2] Duffey, R.B., 2004, "A New General Accident Theory", Paper # 0759, Proceedings of International Conference on Probabilistic Safety Assessment and Management (PSAM 7), Berlin, Germany, June 14-18.

[3] Duffey, R.B. and Saull, J.W., 2004, "Error States – Proactive Learning, Forgetting and Elimination using Safety Management", Paper #0760, Proceedings of International Conference on Probabilistic Safety Assessment and Management (PSAM 7), Berlin, Germany, June 14-18.

[4] Duffey, R.B. and Saull, J.W., 2002, "Know the Risk", First Edition, Butterworth and Heinemann, Boston, USA.

[5] Duffey, R.B., Saull, J.W., 2001, "Errors in Technological Systems", Proceedings of World Congress on Safety of Modern Technical Systems, TUV Foundation, Saarbruecken, Germany, September 12-14.

[6] Duffey, R.B. and Saull, J.W., 1999, "On a Minimum Error Rate in Complex Technological Systems", Proceedings of the 52nd Annual International Air Safety Seminar on Enhancing Safety in the 21st Century, Rio de Janeiro, Brazil, pp.289-305.

[7] Ohlsson, S., 1996, "Learning from Performance Errors", Psychological Review, Vol 103, No 2, pp.241-262.

[8] Sveshnikov, A.A., 1968, " Problems in Probability Theory, Mathematical Statistics and the Theory of Random Functions", Dover, New York, pp.48, 49 and 80.

[9] Columbia Accident Investigation Board (CAIB), 2003, "Report Volume 1", August, and "Working Scenario", CAIB/NASA/NAIT, July, and "Volumes II through VII", October (available at http://www.caib.us).

[10] Baumont, G., Bardou, S. and Matahri, N., 2000, "Indicators of Plant Performance During Events Identified by Recuperare Method", Paper presented at Specialist Meeting on Safety Performance Indicators, Madrid, Spain, October 17-19.

[11] Gabinete de Prevenção e Investigação de Acidentes com Aeronaves (GPIAA), 2001, Government of Portugal, Ministério das Obras Públicas, Transportes e Comunicações, Aviation Accidents Prevention and Investigation Department, (ACCID) "All Engines-out Landing Due to Fuel Exhaustion Air Transit Airbus A330-243 marks C-GITS Lajes, Azores, Accident Investigation Final Report 22 / AC, Portugal 24 August.

[12] ASME Books, RA-S-2002, "Probabilistic Risk Assessment for Nuclear Power Plant Applications", ISBN #0791827453

[13] Hannaman, G.W., 2005, "Selecting Probability Models for HRA", Proceedings of PSA Conference, ANS, San Francisco.

[14] Hannaman, G.W., 2006, private communication.